CSSLab

Computational Social Science Lab

Annual Report 2025

Annenberg School for Communication, University of Pennsylvania

Penn Engineering, University of Pennsylvania

Wharton, University of Pennsylvania

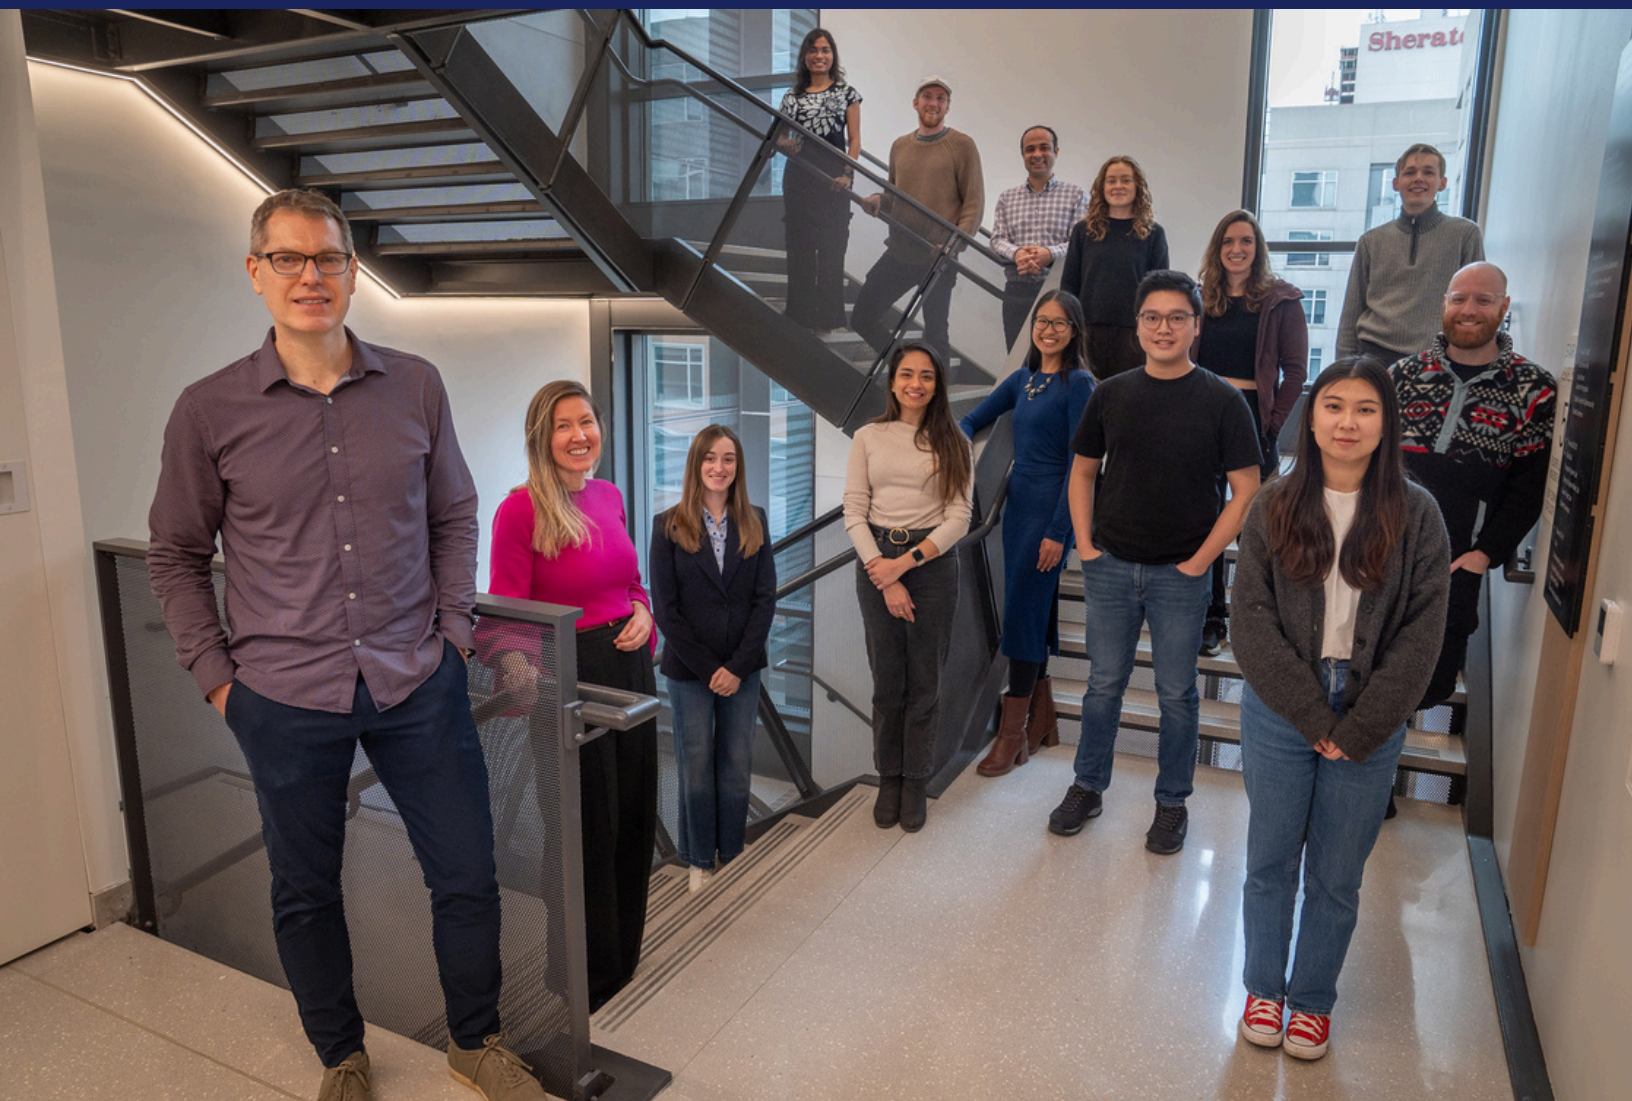

RESEARCH UPDATES:

The CSSLab is excited to announce new research in its four research areas: PennMAP, Integrative Experiments, NOMAD, and Common Sense.

PennMAP (Penn Media Accountability Project)- PennMAP was launched to increase media transparency and accountability. It also hosts the "Living Journal," a set of dashboards based on PennMAP research displaying data that updates in real-time, providing users with live insights into the information ecosystem.

Narrative License and the presence of inappropriate causal claims in Social Science articles- PhD students Calvin Isch, Timothy Dorr, and Neil Fasching have been studying narrative license- the tendency of some researchers to make claims that deviate from their research findings. For example, researchers may incorrectly imply that correlation equals causation from their findings or exaggerate statistical results, misrepresenting them. Even after the publication stage, narrative license can occur on the communications side with press releases. This misleading practice has drawn concerns, so the researchers wanted to study how prevalent narrative license is in social sciences literature. Isch, Dorr, and Fasching focused on English-written social science studies that were published between December 31, 2014 and January 1st, 2025, yielding 17,448 articles to analyze. The researchers also utilized a large-language (LLM) model to identify causal language and causality in these articles and found that this language has increased over time, specifically among papers on COVID-19. Overall, inappropriate causal language was found across all social science disciplines as well as in both low- and high-impact journals. LLMs have the potential to detect Narrative License, which would reduce the prevalence of misleading claims in research and improve academic communication.

Media bias in portrayals of mortality risks and health behavior (Calvin Isch)- Silent illnesses, or chronic diseases, contribute to 70% of deaths in the US annually and six in ten Americans suffer from at least one chronic condition. Despite the statistics, coverage of this public health crisis is disproportionately overshadowed by sensational risks, including terrorism, homicide, and traffic accidents- incidents that are far more likely to grab a reader's attention. Isch further explored this bias and imbalance in media coverage, finding that the mainstream media tends to amplify sensational risks while underrepresenting chronic risks, highlighting a disparity between risks covered by the media and the mortality risks that threaten Americans the most. Using natural language processing techniques, Isch collected monthly data on 14 different mortality risks using keyword searches on 823,406 major US news outlet articles that were published between 1999 to 2020. Heart disease, the leading cause of death in the US (accounts for 36% of overall deaths), receives one article per 323 deaths; in contrast, terrorism which accounts for 0.00008% of overall deaths, has 36 articles written about it per death, highlighting that deadlier risks receive less coverage compared to their sensational counterparts. Overall there is a weak connection between mortality

rates and media coverage. Isch's analysis also looked for any mention of health interventions as well as article tone to determine if bias manifested in other ways in articles. It was found that any mitigation strategies discussed fell under three main categories: policy, behavioral, and technological. Articles on chronic diseases usually emphasized behavioral and technological solutions, while those on sensational risks focused more on collective policy solutions. Studying these trends is significant because of the potential for the media to influence public-health policy, highlighting the need for a shift in prioritization of public health policies and how mortality risks are presented in the media.

Tracking the production and absorption of political narratives in 2024 (Baird Howland, David Rothschild, Duncan Watts)- This research focuses on a new data analysis system which tracks political narratives across news publications across the political spectrum, revealing tone and partisanship. They carried this out by collecting the top 20 articles from six publishers from January 1, 2024 to Election Day on November 5, 2024. This system combines BERTopic and LLM to track the most important narratives of the year. The motivation behind this research is from previous literature which discusses how political narratives shape people's beliefs about different policies, with narratives defined as mental models which people use to understand the world around them, including how they interpret the news; this causes people to interpret the same news events differently, which further contributes to political polarization. So the researchers classified these narratives by analyzing six major news publishers- *Breitbart*, *Fox News*, *New York Times*, *Washington Post*, *The Guardian*, and *HuffPost*. Then the data analysis system clustered articles by group in terms of similar narratives to identify recurring themes; this process was refined until a set of narratives was constructed. GPT was then used to label these narratives by partisanship, tone, and discipline (consistency of framing and messaging), and it was found that most narratives were negative, where in *Fox News*, 47% of their articles were dominated by the same 20 narratives. Out of all the narratives, there is a pro-Republican bias, with *New York Times* surprisingly having a Republican lean in their articles and *Fox News* and *Breitbart* having narratives that are highly consistent with their ideological stance. Additionally, 15,000 people were surveyed in the weeks leading up to the 2024 Election, and the results showed that Republicans are more likely than Democrats to produce and absorb narratives; and that public opinion is not carried out by publishing individual articles but rather repeated narratives which have much stronger impact on readers.

Exploring Political Attitudes under Targeted Media Consumption with Large Language Models- Josh Nguyen, Upasana Dutta, Samar Haider, Neil Fasching, Bryan Li, Duncan Watts carried out large-scale AI-run experiments in order to determine whether selective media exposure has negative consequences or does not have any polarization effects at all. For this study, LLMs substituted human subjects which are more costly and challenging to test. They utilized training data for this study so that the LLMs can internalize biases as well as the political opinions of the people who may consume the kind of content present in the data. To gather the data, the researchers selected news content related to immigration and crime published by four outlets: *HuffPost*, *CNN*, *Fox News*, and *Breitbart*. Using different copies of GPT-4 to mimic real participants, the models were given 20 articles with different combinations of content from publishers and then were given statements of different political lean in which they were asked to indicate their agreement on. The subjects tended

to lean more liberal on the two topics when they were exposed to articles all from the same publishers, with their ratings aligning with the political stance of that publisher. The researchers also estimated treatment effects where the impact of media exposure on opinions was measured using statistical models, revealing how different news publishers influence attitudes on immigration and crime.

Reinforcement Learning from Human Feedback: Adaptation in Polarized Environments- Reinforcement learning from human feedback is a tool used to align LLMs with human preferences, and this relationship is notable to researchers given that AI has both positive and negative impacts on political discourse and deliberation. Sikata Sengupta, Amir Ghasemian, and Duncan Watts developed an AI system that learns personalized biases by allowing users to label responses as more left- or right- leaning relative to their own views, balancing personalization and neutrality in order to yield coherent AI responses. This approach enables the AI algorithm to adjust their responses to better reflect individual perspectives while maintaining overall neutrality. To test the efficacy of this, the researchers plan to carry out Deliberation Experiments to see how users interact with AI during conversations.

Bad Medicine: How low-quality hoaxes jump into the mainstream- David Rothschild carried out research on the reach of conspiracy theories, specifically COVID hoaxes. Rothschild created a dataset by collecting information from government health sources and fact-checkers about COVID conspiracy theories (156 total) that were debunked within the first three months of the pandemic, focusing on low-traction medical hoaxes that either were not discussed in the mainstream media or barely made in into low-quality media. Using media and search data and combining the use of LLMs with human-in-the-loop review to classify debunking content, Rothschild found that mainstream media sources weren't much better at debunking medical hoaxes compared to low-quality news outlets. Of the 156 hoaxes analyzed in this study, only a few actually gained significant interest and subsequently received extensive media coverage. Articles on hoaxes that received more attention tended to highlight rather than focusing on debunking the hoax itself. Using postal-code data, it was found that senior citizens were more likely to "engage" with medical hoaxes, meaning they searched for specific hoaxes multiple times instead of just once. Instead of focusing on what makes a medical hoax more likely to reach the mainstream media, Rothschild's work identifies the factors that contribute to the engagement and spread of medical hoaxes.

Same Economy, Different Stories: A Decade of Partisan Economic Reporting in American Media- Exploring how Americans' perceptions on the economy are influenced by the political affiliation of the news they watched, David Rothschild analyzed 200,000 articles that were published on the landing pages of mainstream media outlets (*Wall Street Journal*, *New York Times*, *Washington Post*, *Fox News*, *Breitbart*, and *HuffPost*) from 2015-2025, with a focus on articles published around the 2024 election cycle. Rothschild's analysis reveals that publishers have many choices when it comes to how to cover the economy and how to frame stories on the economy, causing further divisions among voters. The analysis specifically focuses the selection and framing of economic statistics in news articles where outlets tend to cover stories on the economy depending on their

political lean (ie. right leaning outlets publishing more negative articles on the economy when Biden was president), with this effect being especially strong among *Fox News* and *Breitbart*. Rothschild also looked at how they chose to cover various topics- with some publishers (*Bretibart* and *New York Times*) covering inflation extensively even after periods where inflation has cooled down. Overall, right-leaning outlets published more negative coverage of the economy compared to their left-leaning counterparts. This analysis is relevant to the current information ecosystem because it shows that during election seasons, the mainstream media puts out economic narratives that fit their political beliefs, resulting in further polarization among public opinion.

Feelings Don't Care About Facts: Persuasive Effect of Framing Outweighs Factuality- Concerns about fake news are misplaced, as factually accurate but biased information is more harmful to readers than misinformation because of its ability to impact public perceptions and potentially mislead. For example, negativity bias is considered to be less of a threat than misinformation, so this focus on fake news may be because of assumptions that fake news has a stronger persuasive effect than stories that don't include falsehoods. But there is a lack of research which supports this, so Amir Tohidi, Samar Haider, Duncan Watts, David Rothschild designed a study where they employed LLMs to generate synthetic articles from a set of facts and quotes extracted from real articles with different combinations of tone and credibility. Then 503 participants that were recruited from Prolific were asked to read a combination of the articles and complete a survey about how the articles impacted their feelings and opinions. From the results, biased information did have a significant impact on readers' opinions, with similar effects of that to being exposed to fake news; these findings challenge previous claims that fake news is more powerful at shaping public opinion than biased information. The researchers therefore advocate for exploring other ways that the public can be misled, beyond just examining misinformation in the media.

Rightwing narratives outnumber mainstream narratives in US political media eight years straight- Unlike stories which have a clear ending and beginning, narratives are open-ended frameworks by which events are interpreted to reinforce existing political ideologies or political movements. In order to better understand media narratives, Baird Howland, David Rothschild, and Duncan Watts examined "narrative phrases" (politically charged labels that evoke narratives, such as "migrant criminal" and "critical race theory") and how the use of narrative phrases compares between right-leaning and left-leaning outlets. In order to better identify media narratives, it is important to look for repetition, such as repetitive coverage that uses the same stylized elements (ie. imagery) that reinforce an existing narrative. Using 342,000 news articles, the researchers filtered for common phrases that appeared frequently and showed repeated bursts of attention over time. These phrases matched to one out of 24 news topics if at least 20% of the articles that mentioned the phrase were also about the topic, yielding a total of 18 topics with a large volume of associated narrative phrases. The results showed that right-leaning outlets have significantly more narrative phrases in their coverage compared to mainstream publishers, with "Elections, "Crime," Civil Rights, Minority Issues, and Civil Liberties" and "Education" dominating right-wing narratives. As a whole, right-wing publishers have a "quantitative advantage" over their left-leaning counterparts as they

cover a broader range of stories and topics more frequently, which consequently results in more support for far-right movements.

"Genocide" vs. "War with Hamas": documenting the Western media narrative on Gaza- Following the October 7th attacks, media coverage in the Israel-Gaza conflict is regarded as highly-polarized. To better understand the competing narratives present in the mainstream media, Baird Howland and Timothy Dorr examined coverage of the war from 89 news publishers from four different geopolitical groups (U.S., Western Allied Nations, Middle East, and other countries from the rest of the world) from November 3, 2024 to February 20, 2024 to determine if the media environment has an impact on how the conflict has been covered. To identify how different outlets cover the War, the researchers used a keyword extractor to identify conflict-related phrases from articles, where similar phrases were clustered together, yielding a set of labels relevant to coverage of the war. These labels reflect how outlets described the conflict, whether they emphasized Hamas or Hamas terrorism or even acknowledge a "genocide" in Palestine. Based on the results, 46 out of 53 US publishers chose a pro-Israel stance, emphasizing "Hamas War" and "Hamas Terror." For non-US Western outlets, the majority of them also aligned with the Israeli narrative, and 50% of Middle Eastern outlets showed a preference for Palestinian perspectives. But compared to US outlets, outlets from the rest of the world emphasized violence against Palestinians, prompting further research into whether historical contexts are woven into these narratives.

Uncovering Anti-Democratic Norm Justifications (Timothy Dorr)- While open-ended surveys yield deeper insights into social science research, carrying them out is more expensive and time-consuming. Consequently, quantitative survey research relies on top-down structures, where categories of responses are already predetermined by the researcher, and researchers may miss potential answers that don't fit into the responses. To address this limitation, Timothy Dorr introduces a new method which uses LLMs to analyze open-ended responses more efficiently by categorizing them using 13,000 responses about democratic attitudes in the US. Both human and LLM coding were used, and classification prompts were used to help the LLMs interpret responses more accurately. To test this mode, Dorr had the model analyze over 13,000 responses from a survey where participants gave open-ended responses on their perspectives on democracy, where a contradiction was found: people who expressed strong support for democracy have the tendency to also support policies which undermine democracy. The open-ended question for them to answer was to reconcile and explain their contradictory position, and the LLM classified six different responses on why they support anti-democratic policies, with most justifications viewing political institutions as corrupt. This comprehensive method will lead to more descriptive findings and quantifiable results into research on political attitudes.

Divergence Between Predicted and Actual Perception of Climate Information- With the American political system being strongly influenced by moderate voters, policies focused on mitigating climate action might be hindered if the party's constituents have little or no concern over the issue. Previous research has shown that despite a highly polarized political climate, bipartisan support is necessary for effective climate change action. As actions are often influenced by

perceptions of how others will take action, people are less likely to take climate change action if they believe that climate change skeptics are unlikely to take action on climate change themselves. Using a crowdsource approach, Amir Tohidi studied the relationship between expected and actual persuasiveness of climate change information. Based on results from participants of varying climate change stances (advocates, moderates, skeptics), there was a significant divergence between anticipated and actual shifts in opinion before and after reading climate change articles, as each group thought that news articles would not be persuasive enough in to change climate skeptics' opinions. But despite a change in attitude among climate skeptics, this didn't translate into actions, such as making climate change donations or adopting more climate-friendly behaviors. This gap highlights that cross-partisan dialogue is necessary in order to better engage skeptics.

News-like Information Ecosystem on YouTube- With the rise of digital platforms is another type of news source, known as "news-like" content, which differs from traditional news in that it is not produced by journalistic institutions and is typically found on social media. News-like content is user-generated; for example, it can be created by influencers with growing platforms on X. Another key difference between news-like information and traditional news is shareability, where news-like content has greater engagement potential and often includes multimedia elements such as short videos and audio which appeal more to younger audiences. These features have narrowed the gap between what is considered news vs. entertainment. To compare the production and consumption of news and news-like content, Amir Ghasemian, Homa Hosseinmardi, Upasana Dutta, Jennifer Allen, David Rothschild, and Duncan Watts analyzed content on 1,690 YouTube channels to determine if news-like content dominates traditional news content. The researchers found that though traditional outlets produced more content overall, news-like content received significantly more views. In regards to information sources of news-like content, content creators typically referenced other platforms (mostly X) or cited other people. Because of the prevalence of news-like content, news consumption studies should broaden to include this newer category so that researchers have a more holistic view of emerging trends in the information ecosystem.

Trapped by Personalization: LLM-Assisted Recommendations and Their Hidden Impact on Human-AI Ecosystems- Recommender systems influence decision-making processes across multiple areas, including entertainment, healthcare, and social interactions. These effects are especially prominent on social media, where systems can prolong engagement for users which can lead to cognitive challenges. Rec systems also help with personalizing content, and over this time this can affect user behavior and contribute to societal polarization. To study recommendation effects, Amir Ghasemian, Homa Hosseinmardi, Upasana Dutta, and Duncan Watts utilized counterfactual bots to mimic real users by having them watch videos based on real viewing histories; then they prompted LLMs to generate recommendations in order to better understand their influence on content exposure. Comparing YouTube's recommendation system with LLM-assisted recommendation systems, the LLMs produced more personalized recommendations (especially for right-leaning users), but they also recommended more extreme content. These findings highlight the need for responsible LLM systems that have instruction-based prompts to help with guiding the recommendation system's behavior so that they can recommend less harmful content.

Characterizing controversial far-right communications on Telegram- Fringe communities, or social groups that share controversial or extreme content, have resorted to migrating to alternative platforms to share information more privately. One of these platforms is Telegram, which has been around for over 10 years and has around 700 million users. Upasana Dutta carried out a study which studies the contents of messages shared on Telegram and whether it is a small percentage of users accounting for the majority of messages shared on Telegram or if this message sharing is more spread out among users. Dutta analyzed over 200,000 messages from the platform from far-right groups and then applied a BERT-based topic modeling approach to identify the most common topics discussed in the messages. The MAGA Trumpists and QAnon channels mainly focused on political figures and geopolitical events (ie. Kamala Harris, Ukraine), while White supremacist groups mainly discussed identity politics (ie. pro-white). And then to determine the amount of harmful content on the platform, Dutta used an openAI moderation model that classified speech as either harassment, hate, or violence. The Proud Boys group exhibited the highest amounts of harassment while the MAGA Trumpists and conspiracy theory groups contained the most violence-related content. In addition to gaining insights into the content, Dutta also examined how information was exchanged between the platforms on the channel and found that a very small number of channels were responsible for most of the messages exchanged.

Russian News on Social Media Networks: Evidence from Facebook and VK- Under authoritarian regimes, news outlets have less journalistic freedom due to increased censorship which has led to fears of retaliation. This has prompted more interest into the dynamics between authoritarian governments and the media outlets under them, and how these outlets navigate content production. But less is known about the demand for this content rather than the type of content produced. Using media data from Facebook and VKontakte, Sam Wolken explores the relationship between news production and consumption in Russia. Article links on these social media platforms that originally came from 19 Russian news outlets were extracted and analyzed by their engagement and classified by topic, such as their tone towards Putin and if there were any mentions of democracy or other foreign countries (ie. Ukraine). Wolken found that around 50% of the content from these news outlets was political, with negative content about Putin and the Russian government receiving the most engagement; this increase was especially pronounced after Russia's invasion of Ukraine in 2022.

Integrative Experiments: This research area includes several subgroups- Group Synergy, PGG, Team Communication, and Deliberation Lab. Projects in these areas are focused on how people make decisions in group settings and how social dynamics influence outcomes.

Rich Talk: How Communication Predicts Cooperative Behavior- Communications is often described as "cheap talk" because it is considered non-binding and therefore meaningless. But on the other hand, communication is clearly meaningful as it increases cooperation; however, researchers are still trying to understand why communication is so effective. Using natural language processing, Emily Hu analyzed free-form conversations in order to identify the individual features that predict

cooperative behavior. By examining 836 public goods games (where participants are given a set of coins and can decide how many to contribute to a public fund or keep for themselves after chatting with each other via text), Hu was able to analyze each game round and determine if specific conversational features could predict how people contribute in these games. Having the chat itself improved cooperation between participations, but the main predictor of cooperation ended up being "contribution norms," which refers to how many times a group member may attempt to establish norms within the game by making promises or recommendations to the other participants. Hu's work integrates existing theories on why communication encourages cooperation and highlights the use of natural language processing in analyzing large amounts of conversational data.

Integrative Experiments Identify How Punishment Impacts Welfare in Public Goods Games- A persistent challenge continues to be identifying the factors that influence whether people choose what is best for them individually versus what is best for the group. Public goods games are a good model to study this as it reflects real-world decisions people make while also providing a framework to carry out replicable experiments. Mohammed Alsobay, David Rand, Duncan Watts, and Abdullah Almaatouq specifically focus on how punishment affects cooperation in groups and under what circumstances does punishment result in the most cooperation. The researchers implemented a design where they tested hundreds of public goods game configurations to determine the conditions where punishment leads to the most cooperation using a predictive model. In general, punishment did increase contributions, but this did not translate into group efficiency. Based on the predictive model that the researchers applied, punishment only improved game outcomes under certain conditions, where the model outperformed human predictions. It was found that factors such as participants being able to discuss with others have larger effects on how punishment impacts games, while the effects of game length and contribution framing were more complicated in that they were dependent on other factors present.

Harmony in Discord: A Multimodal Exploration of Constructive Behaviors for Managing Disagreements- While conflict is usually seen as harmful, there is evidence that disagreement can lead to problem-solving, as seen in the case of cross-partisan dialogue. Emily Hu and James Houghton use machine learning to explore "conflict expressions," or behaviors exhibited by people when they are disagreeing with the person they are speaking to. The researchers carried out an online video study with 94 participants, where they extracted features (visual, vocal, linguistic) from the videos and determined which ones allowed for the differentiation of "destructive" vs "constructive" conflicts throughout the study. Conversations were grouped into high and low conflict pairs, with high-conflict pairs showing more interruptions and back-and-forth in the earlier stages of the discussion and most of the conversation having a negative tone. These findings will eventually be validated by humans, which will better help in studying how to help high-conflict groups have more productive conversations.

Perceptions of Fairness in Machine Learning- As algorithms become more involved in decision-making, their ability to impact fairness may be greater than that of human decisions. Existing research on algorithmic fairness seeks to create algorithms that cause the least harm to

people, focusing on fairness in terms of how outcomes are allocated and or how decisions are made. While some algorithms appear to be fair by creating equal outcomes across different groups, there can be instances of people within the same group (ie. race) receiving different predictions. To further examine this, Bethany Hsiao and Duncan Watts created a study where participants were shown how a medical algorithm treats Black versus White patients and they were asked to set treatment thresholds for both groups. The results showed that it is difficult to achieve differing definitions of fairness simultaneously, as treating both groups the same may be more unfair for one group compared to the other. Additionally, what is considered fair can be highly dependent on the context as well as the type of information people receive before coming to a decision.

NOMAD (Network for Open Mobility Analysis and Data): Previously known as COVID-Philly, NOMAD is an infrastructure project that focuses on democratizing access to large-scale GPS mobility data and providing standardized analytical tools for advancing human mobility research and epidemiology. The project seeks to address challenges that limit researchers' and policymakers' ability to effectively utilize human mobility data for research and policy design. By removing these barriers, NOMAD also supports equity by facilitating global research by smaller and diverse institutions, as well as applications in historically data-poor regions.

A Synthetic Dataset and Framework to Test Pre-processing Algorithms of Human Mobility Data- GPS mobility data is widely used in a number of fields including epidemic modeling and transportation planning, but this data isn't free from biases. This bias typically stems from how GPS data is processed, as algorithms used to detect stops and identify important locations can yield errors due to a number of factors including sparse data or complex urban environments. To address these limitations, Thomas Li, Francisco Barreras, and Duncan Watts introduce "Garden City," an open-source tool that simulates realistic GPS movement data that serve as a proxy for ground-truth. In this model, agents move through a simulated city and have designated work and home locations which yield detailed trajectories; then biases are introduced into this model to mimic the challenges found in real-world GPS data. By using Garden City, researchers can test GPS-processing algorithms in order to gain more insights into how different data conditions can affect the quality and accuracy of GPS mobility data. Additionally, this would lead to more robust methods in processing real GPS data.

Common Sense:

Exploring Common Sense Around the World- To build off on their previous common sense research, Mark Whiting, Amirhossein Nakhaei, Josh Nguyen, and Duncan Watts built a public platform called The common sense project, where participants can take survey to measure their common sense and be assigned a commonsensicality score. Since the launch of the project site, the researchers have collected ratings from 54,028 participants from 30 countries and the platform has

also been translated to 10 additional languages: Arabic, Bengali, Spanish, French, Hindi, Japanese, Portuguese, Russian, and Chinese. The development of this research platform has allowed the researchers to explore common sense on a global scale and the researchers will continually open participation to better measure common sense on both the individual and group levels. On a broad level, most countries exhibit similar levels of commonsensicality. The researchers also focused on different demographic traits, ie. gender and marital status, and found more variations between these groups; these results suggest that there is potential for studying larger groups to further examine these metrics.

Evaluating Commonsense Knowledge in Human and Large Language Model-Simulated Populations- Because there is no ground truth when it comes to common sense, research on common sense tends to focus on measuring how much people agree with each other on certain statements within groups. This study by Josh Nguyen, Mark Whiting, and Duncan Watts looks into whether large language models (LLMs) are able to reflect the diversity of common sense beliefs by humans, as LLMs have shown to be able to imitate human thinking in areas such as public opinion and moral judgement. The researchers had the LLM models evaluate several common sense statements and whether they thought others would agree with each statement. The researchers calculated a commonsensicality score based on the LLMs's responses and compared them to human scores to see how closely LLMs could match human judgements. While the models only moderately reflected human judgement, they exhibited some of the same tendencies as humans, including the preference for fact over opinion, descriptions of physical over social realities, and literal expressions over figures of speech.

COLLABORATIONS:

The CSSLab is proud to collaborate with the following data providers, research platforms, and application partners who make our research possible:

- **Polygraph** (industry)- Polygraph is a data visualization studio that specializes in creating sophisticated visualizations and visual storytelling. Their team assists with all stages of a project, from ideation to research, design, narrative, code and publishing, playing a large role in developing the front-end of the Media Bias Detector.
- **Pollfish** (industry)- The CSSLab has also collaborated with Pollfish as part of the Media Bias Detector project; we used pollfish to create a weekly tracking survey which measured participants' responses to how often they heard about political topics in the last week.
- **Amazon Web Services** (platform)- Hosts three data repositories: Nielsen, TVEyes, and PeakMetrics. The CSSLab's data is stored and processed on this platform.
- **Nielsen** (industry)- Nielsen is a data analytics company that provides U.S. media consumption data of multiple mediums, including TV and mobile.

- **TVEyes** (industry)- TV data provider that monitors the production of 2,500 TV and radio stations globally in real time and provides closed-caption transcripts.
- **PeakMetrics** (industry)- Provides web publisher data including a continuously updating collection of news websites.
- **Harmony Labs**- Media research lab that provides preprocessed web data and TV data.
- **Safegraph** (industry)- Provides location and mobility data for regions based in the US.
- **Center for Innovation and Sustainability in Local Media** (community)- Has been helping with the "mapping the media landscape" project by providing data on local and regional newspapers across the country, which has helped the news publishers database grow.
- **NewsGuard** (industry) / **AllSides** (industry) / **Media Bias / Fact Check** (industry)- All three of these partnerships provide us with political bias and credibility ratings for several thousand news outlets.
- **Empirica** (platform)- Where high-throughput experiments take place; experiment designers can focus on the user interface and how participants interact during the experiment while the platform takes care of data storage and other background tasks.
- **Wharton Behavioral Lab** (platform within Penn)- This lab has provided other resources for carrying out high-throughput experiments.
- **Open Science Foundation** (platform)- Serves as a hub for researchers to manage and share their data with others.
- **The COVID-19 High Performance Computing Consortium** (platform)- Led by the White House Office of Science and Technology, the U.S. Department of Energy, and IBM, this platform provided IT resources to our lab for epidemiological research.
- **City of Philadelphia Office of Innovation and Technology** (community)- The City of Philadelphia is an application partner that oversees technology projects in the City and we have partnered with them in building a set of dashboards that will track human mobility patterns at the community scale.
- **Spectus** (industry)- Provided location and mobility data that made up the initial datasets for human mobility dashboards.
- **Prolific** (industry)- Technology company that connects researchers with a large pool of verified participants, allowing for large-scale data collection for studies, surveys, and experiments.

OUTPUT:

DATA DASHBOARDS

1. **Media Bias Detector** (launched on June 25, 2024)- Funded by Wharton alumnus, Richard Jay Mack, (W'89), the CSSLab's Media Bias Detector analyzes and classifies news coverage by ten news publishers in close to real-time. The main features of the dashboard are the Coverage and Events sections, where users can explore trends in political lean and tone among different news topics between publishers or within the same publisher. Users can also view the most covered news stories of the day to see what the mainstream media chooses to emphasize or not cover, with these visualizations revealing greater insights into the editorial decisions that journalists make every day.

Media Bias Detector Blogs: The Media Bias Detector also has a blog page, which consists of editorials written by various CSSLab members. These blogs draw on data from the dashboard and provide analyses designed to help users engage more with the Media Bias Detector.

- The Trump Effect: Shift in Coverage on the Ukraine War (March 21, 2025)- Delphine Gardiner
- Urban Disasters Through the Lens of the Mainstream Media (February 18, 2025)- Delphine Gardiner
- Media Coverage (or lack thereof) of Trump's First Week in Office Part 3: Executive Orders (February 10, 2025)- Delphine Gardiner
- Understanding the Impact of Tariffs on Consumers and Businesses (February 6, 2025) -David Rothschild
- Media Coverage (or lack thereof) of Trump's First Week in Office Part 2: Elon Musk's Salute (February 6, 2025)- Delphine Gardiner
- Media Coverage (or lack thereof) of Trump's First Week in Office Part 1: Trump's Inauguration (January 30, 2025)- David Rothschild
- New York Times: Biden's age is 10.9x more important than Trump's Project 2025 (December 19, 2024)- David Rothschild, Amir Tohidi, Duncan Watts
- A comparison of media coverage on Trump's age vs. Biden's age (December 18, 2024)- Delphine Gardiner
- 2024 Election Recap: What the Media Told Us (December 6, 2024)- Delphine Gardiner, David Rothschild
- The Biggest News Stories Before Election Day (November 11, 2024)- Delphine Gardiner
- Sanewashing in the media: the coverage that normalizes Trump's bizarre behavior (November 1, 2024)- Delphine Gardiner

- Worried about the Russians dividing America? The call is coming from inside the house (September 29, 2024)- Jennifer Allen
- How the Media Bias Detector Shows Differences in Public Discourse Between the Biden vs Trump and the Harris vs Trump Debates (September 25, 2024)- David Rothschild, Delphine Gardiner
- What the Media Bias Detector is telling us going into the upcoming Trump v. Harris debate (and what it will tell us coming out of it) (September 10, 2024)- David Rothschild
- A Closer Look at Media Coverage of Kamala Harris's Endorsement of Tim Walz (August 9, 2024)- Delphine Gardiner
- Death in the news: The mortality risks we face and the ones media organizations talk about (August 1, 2024)- Calvin Isch
- Investigating the Impact of Media Bias on News Readers (July 26, 2024)- Delphine Gardiner, Amir Tohidi, Samar Haider
- Examining Media Coverage of Misinformation (July 16, 2024)- Delphine Gardiner
- Trump v Biden: How the Media Covered the Presidential Debate (July 3, 2024)- Samar Haider
- Don't Judge a News Story by its Headline (June 24, 2024)- Amir Tohidi, Samar Haider
- David's Personal Highlights of the Media Bias Detector! (June 24, 2024)- David Rothschild
- Understanding Media Bias: What Our Dashboard Shows (June 12, 2024)- Delphine Gardiner
- Joe Biden's (but not Donald Trump's) age: A case study in the New York Times' inconsistent narrative selection and framing (June 12, 2024)- David Rothschild, Duncan Watts, Jennifer Wang

2. **News Consumption** (launched on June 10, 2024)- The CSSLab's second dashboard is a set of four data visualizations that highlight Americans' media consumption habits, with a focus on echo chambers. The interactive visualizations include: Overall News Consumption (where Americans get their news and who is the typical consumer of news), Partisan Echo Chambers (how many people primarily consume partisan news), Echo Chambers Map (Echo chambers broken down by state as well as the most popular TV programs in each state), and Changing TV Audiences (changes in news consumption habits over time). This dashboard is an extension of two previous CSSLab publications: "Evaluating the fake news problem at the scale of the information ecosystem" and "Quantifying partisan news diets in Web and TV audiences" The papers describe how misinformation and highly partisan political content fit into the broader context of media consumption, and the dashboard extends this work by updating the papers' data and allowing for the exploration of dimensions not available in the static version of the papers.

CSSLAB BLOG POSTS: Written by Delphine Gardiner, the CSSLab's communications specialist and in-house writer

- Duncan Watts Presents Fake News, Echo Chambers, and Algorithms at Penn Engineering ASSET Seminar (April 25, 2025)
- Researcher Spotlight: CSSLab Members Present Work at Annenberg Workshop (April 25, 2025)
- News On Climate Change Is More Persuasive Than Expected, Study Finds (March 24, 2025)
- From Cracks to Gardens: Creating a Thriving Social Media Through Research (March 10, 2025)
- New Study Challenges YouTube's Rabbit Hole Effect on Political Polarization (February 18, 2025)
- How the Media Distorts Perceptions on Chronic Disease Risks (February 5, 2025)
- Violent Language in Films Has Increased Since the 1970s: A New Study (January 8, 2025)
- Research Assistant Spotlight: Priya D'Costa Presents at NeurIPS (December 18, 2024)
- Commonsensicality: A New Platform to Measure Your Common Sense (October 4, 2024)
- CSSLab 2024 End-of-Summer Research Seminar Recap (September 13, 2024)
- CSSLab Establishes Virtual Deliberation Lab to Reduce Affective Polarization (September 13, 2024)
- The Team Communication Toolkit: Emily Hu's Award-Winning Project (August 13, 2024)
- Overcoming the Challenges of GPS Mobility Data in Epidemic Modeling (July 29, 2024)
- The CSSLab Launches News Consumption Dashboard (June 10, 2024)
- Homa Hosseinmardi and Sam Wolken Speak at Annenberg Workshop (May 20, 2024)

PUBLICATIONS:

Media Bias Detector: Designing and Implementing a Tool for Real-Time Selection and Framing Bias Analysis in News Coverage (*CHI '25: Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, April 25, 2025)- The authors- Jenny Wang, Samar Haider, Amir Tohidi, Anushkaa Gupta, Yuxuan Zhang, Chris Callison-Burch, David Rothschild, and Duncan Watts- built the Media Bias Detector dashboard, a tool for people to learn more about how bias manifests in news articles from the top publishers. To assess the impact of the tool, the researchers interviewed 13 experts from journalism, communications, and political science to gain detailed feedback on the effectiveness of the dashboard and identify areas for improvement. The interviewees completed tasks utilizing the platform and were then asked to evaluate the Bias Director's ability to help users recognize bias and compare editorial choices across publishers. A few participants noted that while this tool didn't necessarily change their own perceptions on bias, the Media Bias Detector could be a helpful tool for others to better understand media bias. Some interviewers also noted the need for greater transparency such as being able to access the articles that the tool labels and

classifies. To determine the broader applicability of the platform, the researchers also surveyed 150 participants on media bias, and what made the Media Bias Detector stand out to them was the ability for users to view data from many different perspectives.

Divergence Between Predicted and Actual Perception of Climate Information (*PNAS Nexus*, March 18, 2025)- Climate change demands urgent and effective action, but media-driven polarization, especially among strong climate skeptics, remains a barrier. To address this challenge, Amir Tohidi and coauthors explored the relationship between expected and actual persuasiveness of climate change news articles. The authors first conducted a survey experiment in which they collected participants' predictions about the persuasiveness of a set of climate change related articles and the authors found that these articles increased concern among skeptics contrary to most people's expectations that they would be ineffective. In the second part of the experiment, 1,000 climate skeptics were randomly assigned to either read a climate change article or control article (on an unrelated topic); those who read the articles related to climate information had a positive change in attitude towards climate change. However, this concern did not translate into a willingness to support climate policies or take personal action. Tohidi et al.'s research is the first study that assesses the relationship between one's climate change stance and their expectations on the persuasiveness of climate change news articles. The paper also highlights the need for bipartisan support and dialogue to drive climate action, as well as addressing misconceptions about climate skeptics. While skeptics are more receptive to change than expected, the attitude-behavior gap observed in skeptics remains a challenge.

A research agenda for encouraging prosocial behavior on social media (*Nature Human Behavior*, March 10, 2025)- Social media platforms were initially praised for their ability to connect the world, but today these same platforms are criticized for spreading hate and misinformation. These dynamics are explored by Timothy Dorr and Duncan Watts, who examine the tension between positive and negative outcomes on social media platforms. Dorr identified several categories of prosocial behavior that emerge on social media platforms: connecting new communities, collective problem solving, and expanding philanthropy. As social media is already flooded with negative content, there is a lot of research potential in how to make social media better for users and how to create more positive outcomes. Research on prosocial behavior is still in its infancy, but the authors emphasize that progress will require meaningful collaborations between academics and industry partners. By understanding the causal factors that facilitate these prosocial interactions, we can work towards platforms that not only minimize harm but also maximize social benefit, leveraging social media's unprecedented connective capabilities to create spaces where communities can thrive and address shared challenges.

Short-term exposure to filter-bubble recommendation systems has limited polarization effects: Naturalistic experiments on YouTube (*PNAS*, February 18, 2025)- YouTube's role in political polarization is highly contested, with debates over the relationship between the platform's algorithm and user choice in content consumption. To test this, Emily Hu, Dean Knox, and coauthors built a

YouTube-like interface for participants to watch videos and interact with the platform. By running experiments on this custom-built interface, the authors found that algorithmic recommendations had little impact on one's political beliefs and behaviors. The researchers evaluated the effect of the algorithm by comparing the viewer's choices and their post-viewing opinions, with the video content focusing on two policy issues: gun control and minimum wage. In total, the experiments generated over 130,000 recommendations (designed to reflect real recommendations from YouTube). Regardless of the recommendation algorithm, participants tended to choose more videos that aligned with their existing political beliefs; in addition, the rabbit holes were found to not be extremizing. The authors' work stands out for its realism, as they used real videos and human viewers on a fully instrumented YouTube-like platform. Additionally, they gave participants more agency during the study and allowed them to interact naturally with the platform, which more closely mirrors how people engage with recommendation algorithms in the real world.

Media bias in portrayals of mortality risks: Comparison of newspaper coverage to death rates (*Social Science & Medicine*, January 2025)- Chronic diseases contribute to 70% of deaths in the US annually, but coverage of this public health crisis is disproportionately overshadowed by news on sensational risks, such as terrorism and homicide. Calvin Isch further explored this bias and imbalance in media coverage, where he found that the mainstream media tends to amplify sensational risks while underrepresenting chronic risks, highlighting a disparity between risks covered by the media and the mortality risks that threaten Americans the most. Using natural language processing techniques, Isch collected monthly data on 14 different mortality risks using keyword searches on 823,406 US news articles. Isch's analysis showed that deadlier risks (ie. chronic conditions) receive less coverage compared to their sensational counterparts; overall there is a weak connection between mortality rates and media coverage. Isch also found that any health intervention strategies discussed fell under three main categories: policy, behavioral, and technological. Articles on chronic diseases usually emphasized behavioral and technological solutions, while those on sensational risks focused more on collective policy solutions. Terrorism was the most policy-focused risk; in contrast, coverage of diabetes is the least likely to focus on policy, instead emphasizing technological or behavioral solutions. In terms of tone, sensational risk coverage was visibly more negative compared to coverage on chronic disease which was more neutral, which can impact how policymakers prioritize certain risks over others. Further exploring the media's impact on perceptions will help researchers better understand how these perceptions shape public attitudes and policy-making.

Trends of Violence in Movies During the Past Half Century (*Jama Pediatrics*, December 30, 2024)- Violent entertainment has made it into the public discourse due to rising concerns about the graphic nature of highly popular video game franchises. But what about violence in films? After the R Rating was established by the Motion Picture Association of America (MPAA) film rating system in 1968, there was an increase in violent content in films thereafter. The relationship between media and culture is bidirectional. On one hand, movies can shape cultural tastes, beliefs, attitudes, and behaviors. On the other hand, movies often reflect societal realities, including the prevalence of

violence. To better understand this bidirectional dynamic, a crucial question arises: how is the portrayal of violence in media changing over time? Amir Tohidi and coauthors found that violent language in the film industry has increased over time. The researchers used movie dialogue from a total of 166,534 films that were released between 1970 to 2020, extracting murderous words, including *kill* and *murder*. They then used Natural language processing techniques to match the murderous dialogue to the characters; the analysis revealed that 6.97% of these films contained violent language, showing that dialogue involving characters murdering others is increasing over time- not just limited to the crime genre but across all film categories. As this kind of content becomes more commonplace, it becomes a greater concern for many people including parents, healthcare workers, and policymakers.

BordIRlines: A Dataset for Evaluating Cross-lingual Retrieval Augmented Generation (*Proceedings of the First Workshop on Advancing Natural Language Processing for Wikipedia*, November, 2024)- Large Language Models generally perform well when it comes to helping with academic research and technologies, but they have a tendency to sometimes hallucinate information or reinforce existing biases. To address these limitations, retrieval-augmented generation (or RAGs) have been developed to ground these LLMs in real-world sources by retrieving information from a dataset and then giving responses to a prompt. To test the robustness of RAGs, Bryan Li, Samar Haider, and coauthors developed a framework known as BORDIRLINES. They specifically focused on geopolitical bias and also incorporated cross-lingual data to determine if language has an impact on how RAGs respond and handle sensitive bias issues such as territorial disputes. The BORDERLINES dataset of territorial disputes consisted of 720 queries for 251 disputed territories, using articles drawn from Wikipedia. Using Crimea (Ukrainian territory currently annexed by Russia) as an example, the model responded with "Russia" if the query was in Russian but responded with "Ukraine" if the query was in English or Ukrainian. LLMS were found to have a Western cultural bias when asked to make assumptions about various groups of people. But when presented with data consisting of articles with competing information from multiple languages, RAGs will give mixed responses, indicating that a lack of consistency remains a challenge. The findings from this study highlight the need for more robust strategies so that RAG systems can produce more culturally aware, unbiased responses across diverse geopolitical contexts.

The exciting potential and daunting challenge of using GPS human-mobility data for epidemic modeling *(Nature Computational Science*, June 19, 2024)- Epidemic modeling is a framework for evaluating the location and timing of disease transmission events. Global Positioning System (GPS) data, aggregated into massive human mobility datasets, has an unprecedented precision that can revolutionize human mobility science. Because this data is commercial, researchers interested in building epidemic models based on this data face limitations related to data privacy, quality and multiple methodological choices that might bias any findings. Francisco (Paco) Barreras and Duncan Watts discuss the nature of these challenges, determining that to improve the usefulness of this class of data, it is necessary to identify best practices from existing studies on human mobility, identify "sharp edges" in the data processing, and increase transparency and accessibility so that findings with

important policy implications will be more robust. For next steps, Barreras and Watts recommend building repositories where relevant documentation and datasets can be found in a central location to improve the replicability and accessibility of data to researchers. Then the standardization of methodologies will establish guidelines for processing data and promote transparency. The authors hope that these suggestions can encourage community-level dialogue on obtaining the most robust and reliable results for developing epidemic models not just for the sake of academic pursuits but also to improve policymakers' responses to public health crises.

This Land is {Your, My} Land: Evaluating Geopolitical Biases in Language Models (*Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, June 16, 2024)- Though LLMs are useful for generating content in creative contexts, it has been a challenge for them to generate responses close to real-world knowledge, especially in regards to politically sensitive topics. This can create issues where different "facts" will be given depending on the language of the prompt, which could potentially reinforce divisions in viewpoints across cultures, also known as geopolitical bias. Bryan Li, Samar Haider, and Chris Callison-Burch introduce a multilingual dataset called BORDERLINES, and it was used to test an LLM by evaluating its responses to territorial dispute questions in order to measure the model's factual recall, geopolitical bias, and consistency. The researchers also prompted the models to take on different personas (ie. nationalist persona or a neutral one). The BORDERLINES dataset consisted of two parts: a list of disputed territories and questions associated with it, with each territorial dispute containing a multilingual query set (MQS), or queries asking the same question in different languages. Results show that all models used in this study produced insistent answers and that larger models underperformed compared to their smaller counterparts. With regards to altering the model to encapsulate different personas, the nationalism persona increased the model's bias, whereas the neutral persona had the potential to reduce geopolitical bias. Using the BORDERLINES dataset, the researchers showed that multilingual LLMs respond inconsistently to territorial dispute questions depending on the language of the prompt and contextual clues.

Misunderstanding the harms of online misinformation (*Nature*, June 5, 2024)- Despite claims that algorithms drive misinformation, recent research shows that it is largely user-driven. Duncan Watts, David Rothschild, and coauthors reviewed 154 papers on social media misinformation, finding that misinformation exposure is very low and concentrated within fringe communities of people drawn to extreme content. Exaggerated media framing and misleading statistics have led many Americans to believe that misinformation is more prevalent, even though fake news only makes up 0.1% of an average American's news diet. Additionally, these claims rely on correlations, but the rise of polarization is not primarily caused by social media. Watts and Rothschild also argued that exposure, not engagement, is a more accurate indication of the prevalence of misinformation, as people are more selective about what they choose to share publicly, usually because of worrying about what others think of them. Though misinformation is

rare, its impact is disproportionately larger and can lead to "real-world harm," such as increases in hate crimes and civil unrest. To mitigate the harms of misinformation, increasing collaboration between researchers and social media companies while protecting user privacy would allow researchers to thoroughly study the effects of social media to develop more robust interventions and greater platform accountability.

Quantifying the impact of misinformation and vaccine-skeptical content on Facebook (*Science*, May 31, 2024)- Compared to misinformation on COVID-19 vaccines, unflagged vaccine-skeptical content (factually accurate but misleading information) often gets overlooked, thus exacerbating hesitancy towards vaccines. To address this gap in research, Jennifer Allen, Duncan Watts, and David Rand developed a framework to predict the impact of misinformation and vaccine-skeptical content on Facebook on people's vaccination behaviors. Using crowdsourcing and machine learning, the researchers found that unflagged content decreased vaccination intention by 2.28 percentage points per user, making it 46 times more impactful than flagged misinformation, which only reduced intentions by 0.05 points. On the surface, misinformation appears to be more persuasive, but flagged misinformation on the COVID-19 vaccine actually received far less views. Many of these unflagged URLs were linked back to mainstream news, ranging from low-credibility to high-credibility with the *New York Post* and *Fox News* notably having the most negative impact on vaccination outcomes. To combat misleading content, the authors proposed 'physiological inoculation'- exposing people to common deceptive techniques that will help them better recognize factually correct but misleading content and become critical of what they see on social media.

RESEARCHER SITES

This year, CSSLab members launched three project websites:

1. **Deliberation Lab**- With funding from the Templeton World Charity Foundation, James Houghton (Post-doctoral Researcher) and Duncan Watts established The Deliberation Lab, an open platform to automatically coordinate participants, facilitate discussions, and collect standardized data for online experiments all in the same place. By standardizing research methods and data collection procedures, the Deliberation Lab promotes the core values of open science—transparency and replicability. This approach allows researchers to systematically compare conversations across different populations, topics, and demographic mixes to better understand what makes conversations succeed or fail. It will also enable the CSSLab to continue supporting the crucial work of cross-partisan dialogue, aiding researchers who are committed to fostering a world where conversations across lines of difference are both positive and productive.

2. **The common sense project**- The CSSLab invites participants to measure their own common sense by taking a survey on a new platform, The common sense project. This initiative builds on Mark Whiting's (Postdoctoral Fellow) and Duncan Watts' paper, A framework for quantifying individual and collective common sense, published in 2024 in the *Proceedings of the National Academy of Sciences (PNAS)*. In the paper, about 2,000 raters were given a range of claims and asked if they agreed with the claim and whether they thought others also agreed with it. The claims were derived from a wide range of sources, including media, non-fiction literature, and statements generated by both humans and AI. Participation in this platform is open to the general public; participants can visit the site and obtain a score for their own common sense by performing the same tasks as the raters described in Watts' and Whiting's paper.

3. **Team Communication Toolkit**- Emily Hu (Wharton OID PhD student), launched her award-winning Team Communication Toolkit at the Academy of Management Conference in August of 2024. This toolkit allows researchers to analyze text-based communication data among groups and teams by providing over a hundred research-backed conversational features, eliminating the need to compute these features from scratch. The Team Communication Toolkit site allows users to access an open-sourced downloadable Python package for generating conversation features from text data. Emily Hu hopes that this toolkit will greatly benefit the social science community by reducing the cost of extracting features from conversational data. Researchers will now have access to more than a hundred features right at their fingertips, allowing them to conduct analyses more efficiently, as well as test the robustness of their study's findings.

ANNUAL NEWSLETTER

We also produced the 2024 CSSLab Annual Newsletter, which was sent to over 300 of our contacts on December 20, 2024. The newsletter comprised six sections summarizing the CSSLab's research output and lab members' accomplishments throughout the year:

1. Featured
    a. New Media Bias Detector Dashboard
2. Researcher Spotlight
    a. Research Assistant Spotlight: Priya D'Costa Presents at NeurIPS
    b. CSSLab 2024 End-of-Summer Research Seminar Recap
    c. Accelerating the Path to a Master's Degree
3. CSSLab in the News
    a. Commonsensicality: A New Platform to Measure Your Common Sense
    b. The mechanics of collaboration

    c. No, the NFL isn't using facial recognition software to monitor fans
4. Research Funding Successes
    a. University of Pennsylvania launches Penn Center on Media, Technology, and Democracy
    b. CSSLab Establishes Virtual Deliberation Lab to Reduce Affective Polarization
5. Community Engagement
    a. Introducing the Penn AI Council
    b. Investigating the Impact of Media Bias on News Readers
    c. Homa Hosseinmardi and Sam Wolken Speak at Annenberg Workshop
6. New Publications
    a. The exciting potential and daunting challenge of using GPS human-mobility data for epidemic modeling
    b. This Land is {Your, My} Land: Evaluating Geopolitical Biases in Language Models
    c. Misunderstanding the harms of online misinformation
    d. Quantifying the impact of misinformation and vaccine-skeptical content on Facebook
    e. Causally estimating the effect of YouTube's recommender system using counterfactual bots

PODCASTS:

Brooke and Micah Enter the MAGA-verse. Plus, Liberal YouTubers Fight Back (April 18, 2025)- New York Radio's On the Media Podcast included brief soundbites of Duncan Watts, taken from a separate interview, where he discussed the right-wing media ecosystem.

Turns out common sense isn't all that common (March 6, 2025)- Mark Whiting was a guest on The Colin McEnroe Show, where he talked about his latest work, The common sense project.

How choice can skew the public's understanding of what is real and what is fake (August 22, 2024)- David Rothschild talks about the Media Bias Detector on the *It's All Journalism* podcast.

No, the NFL isn't using facial recognition software to monitor fans (August 5, 2024)- Jenny Allen interviews with Fox 4 Dallas to talk about the repercussions of presenting information out of context.

Misconceptions on misinformation (July 6, 2024)- David Rothschild was featured in the *Behind the Markets* podcasts, where he discusses online misinformation, the role of algorithms, and the Media Bias Detector.

Penn researchers develop AI tool to help evaluate media bias (July 2, 2024)- Amir Tohidi interviewed with KYW Radio to talk about the new Media Bias Detector.

- Reposted by Penn Today

Duncan Watts's radio interview with Sirius XM Business Briefing on the Media Bias Detector (July 1, 2024)- Duncan Watts discussed the Media Bias Detector Dashboard and insights on media coverage on the June 27, 2024 Presidential Debate.

PRESENTATIONS:

Conference Presentations

- **Initiative for PeaceMaking conference: Building Peace by Mitigating Political Polarization** (May 8-9, 2025)
    - "What Works Where: An Integrative Approach to Cross-Partisan Conversation" -presented by James Houghton
- **Program on Negotiation (PON) AI Summit** (March 8-9, 2025)
    - "AI as Explorer: Quantifying Conversations with Natural Language Processing"- Emily Hu
- **Sixth Arab Center for Research and Policy Studies International Winter School: Media in Wartime** (January 11, 2025)
    - "A Genocide vs. a 'War with Hamas': Documenting the Western Media Narrative on Palestine"- presented by Baird Howland
- **Society for Judgment and Decision Making Annual Conference 2024** (November 22-25, 2024)
    - "Making Robust Inferences from Text Data: LLMs and the Natural Language Processing Toolkit"- presented by Emily Hu
- **Philosophy of Science Association Conference** (November 14-17, 2024)
    - "INEM: International Network for Economic Method"- presented by Calvin Isch
- **Behavioral Way Summit II Madrid** (November 14-15, 2024)
    - "Which nudge works?"- presented by Linnea Gandhi
- **The 2024 Conference on Empirical Methods in Natural Language Processing** (November 12-16, 2024)
    - "BordIRlines: A Dataset for Evaluating Cross-lingual Retrieval-Augmented Generation"- presented by Bryan Li, Samar Haider, Chris Callison-Burch
- **FIRE Faculty Network Conference** (October 24-26, 2024)
    - "The Politics of Virtue Signaling"- presented by Calvin Isch
- **2024 Conference on Digital Experimentation @ MIT (CODE@MIT)** (October 18-19, 2024)
    - "Investigating the Impact of Media Bias on News Readers"- presented by Samar Haider, Amir Tohidi, and Duncan Watts

- **30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining** (August 25-29, 2024)
    - "Quantifying partisanship in ideologically-biased TV news"- presented by Upasana Dutta
- **84th Annual Meeting of the Academy of Management** (August 9-13, 2024)
    - "A Flexible Python-Based Toolkit For Analyzing Team Communication"- presented by Emily Hu
    - "Decoding Teamwork: The Computational Science of Collaboration"- presented by Emily Hu
    - "Recipes for Success: Quantifying the Relative Influence of Teamwork Facets on Performance"- presented by Emily Hu, Abdullah Almaatouq, Mark Whiting, Duncan Watts
    - "What Are We Talking About? Natural Language Processing in Organisations"- presented by Emily Hu
- **International Conference on Computational Social Science 2024** (July 17-20, 2024)
    - "Investigating the Impact of Media Bias on News Readers"- presented by Amir Tohidi
    - "Why is Common Sense Uncommon? A Network Model for Simulating Identity-dependent Shared Beliefs"- presented by Duncan Watts and Mark Whiting
    - "Research Cartography: To navigate our way out of the generalizability crisis in social science, we need a map"- presented by Duncan Watts and Linnea Gandhi
    - "Latent experiments: learning from experiment-like situations in data"- presented by Mark Whiting
    - "Mind the Gap: Trajectory Data Mining in Highly Sparse Location Datasets"- presented by Duncan Watts, Francisco Barreras, and Thomas Li
    - "Hypothetical nudges approximate the direction but not the magnitude of real behavior change"- presented by Anoushka Kiyawat, Duncan Watts, Linnea Gandhi
    - "Measuring partisanship in ideologically-biased TV news shows"- presented by Duncan Watts, Upasana Dutta, Homa Hosseinmardi
    - "Analysis of emotional bias in American news media discourse"- presented by Baird Howland, David Rothschild, Duncan Watts
    - "What differences in people contribute to differences in common sense?"- presented by Duncan Watts, Mark Whiting
    - "Measuring Diversity in Online News"- presented by Chris Callison-Burch, David Rothschild, Duncan Watts, Samar Haider
    - "Uncovering Anti-Democratic Norm Justifications"- presented by Timothy Dorr
    - "This Land is {Your, My} Land: Evaluating Geopolitical Bias in Language Models through Territorial Disputes"- presented by Bryan Li, Chris Callison-Burch, Samar Haider
    - "The Heterogeneous Effects of Peer Punishment Across Cooperative Settings"- presented by Duncan Watts

- **ACM Collective Intelligence 2024** (June 26-29, 2024)
    - "Recipe for Success: Quantifying the Relative Influence of Teamwork Facets on Performance"- presented by Emily Hu
- **2024 Annual Conference of the North American Chapter of the Association for Computational Linguistics** (June 16-21, 2024)
    - "This Land is {Your, My} Land: Evaluating Geopolitical Biases in Language Models Through Territorial Disputes."- presented by Bryan Li, Samar Haider, Chris Callison-Burch
- **Generative AI and Social Science Research, Yale University Institution for Social and Public Policy and Data-Intensive Social Science Center** (April 5-6, 2024)
    - "How Aligned are Humans and Language Models on Common Sense?"- presented by Josh Nguyen

Talks & Oral Panels

- **Social and Political Economy: Polarization in the Age of AI and the Post-Truth Era** (May 2-3, 2025)
    - "Frames Over Falsehoods: Persuasive effect of Framing Outweighs that of Lying"- presented by Amir Tohidi
- **DataX+IDRE 'Research in the Age of AI' Symposium, UCLA** (April 24, 2025)
    - Panel: Addressing the ethical, legal, and social implications of AI- Homa Hosseinmardi
- **NYU Tech & Society Symposium** (April 2025)
    - Homa Hosseinmardi was an invited speaker where she talked about her current PennMAP research
- **Penn Engineering Asset Center- AI-enabled Systems: Safe Explainable, and Trustworthy** (April 23, 2025)
    - "Fake News, Echo Chambers, and Algorithms: A Data Science Perspective"- presented by Duncan Watts
- **CIND 2025 Workshop: The Impact of AI on (Mis)Information** (April 14-15, 2025)-
    - Quantifying the Impact of Misinformation and Vaccine-Skeptical Content on Facebook"- presented by Jennifer Allen
- **Penn Engineering AI Month 2025** (April 1, 2025)
    - "AI Across Disciplines: A Penn Initiative" - Duncan Watts was a panelist on a discussion focused on AI's societal implications and Penn's potential in the future
- USC (March 2025)
    - Homa Hosseinmardi was an invited speaker where she talked about her current PennMAP research
- **Penn Sociology Colloquium Series: Are We Doing Social Science Backwards?** (February 26, 2025)

    - "An Integrative Approach to Experimental Research"- presented by Duncan Watts
- **Speaker Series on Misinformation, Shorenstein Center for Media, Politics and Public Policy at Harvard Kennedy School and Northeastern University Internet and Democracy Initiative** (February 25, 2025)
    - "Quantifying the Impact of Misinformation and Vaccine-Skeptical Content on Facebook"- presented by Jennifer Allen
- **Penn Presentation** (February 6, 2025)
    - "The Media Bias Detector: Using LLMs to Annotate and Analyze the News at Scale"- presented by Samar Haider
- **Cyber Policy Center, Stanford University Seminar** (January 21, 2025)
    - "Quantifying the Impact of Misinformation and Vaccine-Skeptical Content on Facebook"- presented by Jennifer Allen
- **SciCafe: Digital Divides (American Museum of Natural History)** (December 4, 2024)-
    - Homa Hosseinmardi explored the "how" and "why" of the ways in which society interacts with digital tools, including social media, and how they play a role in rising social and political divisions"
- **Computer Science Department Colloquium Production, Consumption, Absorption, Impact of News at Rutgers** (December 3, 2024)-
    - "LLM and Human-in-the-Loop coding to label news in near-real time"- presented by David Rothschild
- **Knight Science Journalism Program at MIT** (November 7, 2024)
    - "Seminar: Jennifer Allen on the Impact of Misinformation and Vaccine-Skeptical Content on Facebook"
- **Wharton-INSEAD Doctoral Consortium** (October 14-16, 2024)
    - "Perceptions of Fairness in Machine Learning"- presented by Bethany Hsiao
- **AI Horizons: Conversation with David Rothschild, Microsoft (AI at Wharton)** (October 4, 2024)
    - "Large Language Models – Promise or Peril?"- presented by David Rothschild
- **Centre for Digital Health Interventions** (September 27, 2024)
    - "Do nudges change behaviour?"- keynote presentation by Linnea Gandhi
- **National Academy Science Workshop Series** (September 2024)
    - Homa Hosseinmardi was an invited speaker where she talked about her current PennMAP research
- **Invited Oral Presentation in the Research Lab of Professor Matthew Groh, Department of Management & Organizations, Northwestern University** (August 13, 2024)
    - "The Task Space: A Multidimensional Representation of Team Tasks"- presented by Emily Hu
- **Political and Information Networks CIND Workshop** (April 25-26, 2024)
    - Round table talk: "Epistemologies and Data Limits"- presented by Homa Hosseinmardi
    - Round table talk: "Research Priorities in the Era of AI"- presented by Sam Wolken

- **MIT Sloan School of Management Annual Doctoral Research Forum** (May 1, 2024)
    - "Quantifying the Impact of Misinformation and Vaccine-Skeptical Content on Facebook"- presented by Jennifer Allen

Poster Presentations

- **International Conference on Computational Social Science (IC2S2)** (July 17-20, 2024)
    - "Narrative License in the Social Sciences"- presented by Calvin Isch, Neil Fasching, Timothy Dorr
    - "Do Language Models have Common Sense?"- presented by Josh Nguyen, Mark Whiting, and Duncan Watts
    - "A Flexible Python-Based Toolkit for Analyzing Team Communication"- presented by Emily Hu
    - "Using an LLM-based Survey Creation to Track Real-Time News Absorption and Impact"- presented by David Rothschild and Jenny Wang
    - "Atlas: building a map of commensurable knowledge from existing research"- presented by Duncan Watts, Linnea Gandhi, Mark Whiting
    - "Mapping the U.S. News Landscape"- presented by David Rothschild, Duncan Watts, Samar Haider
    - "Political Speech as a Forecasting Tool for Democratic Backsliding"- presented by Regina (Jeanne) Ruane

EVENTS & GATHERINGS:

Multi-lab Collaboration on Reducing Partisan Animosity Planning Meeting - The Deliberation Lab team hosted their first in-person meeting, organized by CSSLab partnerships manager Caroline Rogers, James Houghton, Jeanne Ruane, and Duncan Watts. The meeting was a two day workshop that brought together 25 researchers across the following disciplines: political science, sociology, psychology, public policy, and communications, as well as representatives from non-profit organizations. There were presentations and breakout sessions where researchers discussed potent large-scale experiments on cross-partisan dialogue and outlined next steps.

*The **International Conference on Computational Social Science (IC²S²)** (July 17-20, 2024)- IC2S2 is an annual world-wide conference where researchers from the fields of economics, sociology, political science, psychology, cognitive science, management, computer science, and statistics come together to share their novel findings and address current challenges in the field of computational social science. The conference was held at the University of Pennsylvania and was

organized by the following CSSLab members: Jeanne Ruane, Samar Haider, Sam Wolken, Chris-Callison Burch, Caroline Rogers, Homa Hosseinmardi, and Duncan Watts. The CSSLab was strongly represented at the conference, with 16 accepted papers, 12 oral presentations and 4 posters by lab members.

***The Summer Institutes in Computational Social Sciences** (July 8-20, 2024)- Hosted by the University of Pennsylvania last year, this two-week program brings together scholars interested in computational social science, where participants have the opportunity to develop collaborations and conduct interdisciplinary research. Duncan Watts served on the advisory council, Samar Haider and Sam Wolken were program organizers, and Josh Nguyen served as a teaching assistant.

AWARDS:

Baird Howland (Annenberg PhD student) has been awarded the George Gerbner Postdoctoral Fellowship; this prestigious fellowship will allow Baird to remain at Penn and the Computational Social Science Lab (CSSLab) for an additional two years after he completes his PhD.

Emily Hu (Wharton OID PhD student) won the 2024 IACM Technology Innovation award for her team process mapping project. The award honors a scholar who has made innovative use of technology in research, teaching, or practice in the area of negotiation and conflict management. According to the awards committee: "The toolkit you have been building to quantify conversations is remarkable, and we are delighted to see more work like it."

Other Mentions:

Samar Haider (Computer and Information Science PhD student) was nominated for Best Oral Presentation at IC2S2 2024.

Duncan Watts (CSSLab Director) was appointed to the AI council in the Fall of 2024 which will support and strengthen the AI community at Penn.

STUDENT CONTRIBUTIONS:

- Samar Haider served as Program and Tutorial co-chair for the 10th International Conference on Computational Social Science (IC2S2 2024) and was an organizer for the 2024 Summer Institute for Computational Social Science. He will also contribute as co-chair for The North American Chapter of the Association for Computational Linguistics Student Research Workshop (NAACL SRW) 2025.

- Undergraduate research assistant Thomas Li co-authored a working paper with postdoctoral researcher Francisco (Paco) Barreras: Garden city: A synthetic dataset and sandbox environment for analysis of pre-processing algorithms for GPS human mobility data.
- Undergraduate research assistant Amy Zheng built the Team Communication Toolkit website which was launched in August 2024.
- Research Assistant alum Priya DCosta helped organize a workshop at ICLR 2025 "I Can't Believe It's Not Better: Challenges in Applied Deep Learning."
- Research Assistants Priya DCosta and Evan Rowbotham have contributed to a working paper alongside PhD student Emily Hu, titled: What you say or how you say it? Predicting Conflict Outcomes in Real and LLM-Generated Conversations

CSSLab 3rd Annual Student Mini-Conference (August 2, 2024)- This past Summer, 10 undergraduate and Master's students presented their work and contributions to the CSSLab through six research projects.

Student Presentations:

1. Media Bias Detector- by Emily Duan
2. Human Mobility/COVID- by Thomas Li, Bethany Hsiao
3. Common Sense- by Angie Cao, Dan Kim
4. Deliberation Lab- by Christa Dushime, Evan Ping
5. Team Communication Toolkit- by Evan Rowbotham, Amy Zheng
6. Nudge Cartography- by Ricky Wang, David Feng

POSTDOC & RESEARCH STAFF PLACEMENT:

Jennifer Allen is a postdoctoral researcher at the CSSLab who will be joining NYU as an Assistant Professor in the Technology, Operations, and Statistics group and faculty affiliate of NYU's Center for Social Media and Politics (CSMaP_NYU) in the Summer of 2025.

Mark Whiting (former Senior Computational Social Scientist) started a new position in January 2025 as Chief Technology Officer at Areto, an AI startup based in San Francisco.

Homa Hosseinmardi (former Research Scientist) accepted a position as an associate professor at the UCLA Department of Communication in the Summer of 2024.

(Both Whiting and Hosseinmardi are still affiliated with the CSSLab as collaborative researchers.)

COVERAGE:

Our lab's research and achievements have received extensive media coverage, both within Penn and from external outlets.

A research agenda for encouraging prosocial behaviour on social media

Penn:

- Can Social Media Be Less Toxic? (March 27, 2025)- coverage by Annenberg, featured in Penn Population Studies Newsletter

Divergence between predicted and actual perception of climate information

External:

- News on climate change is more persuasive than expected (April 22, 2025)- Phys.org

Short-term exposure to filter-bubble recommendation systems has limited polarization effects: Naturalistic experiments on YouTube

External:

- Research Record: YouTube's Algorithm and its Effect on Political Polarization (March 18, 2025)- Princeton School of Public and International Affairs

Media bias in portrayals of mortality risks: Comparison of newspaper coverage to death rates

Penn:

- Annenberg study explores how media coverage underemphasizes chronic illnesses as cause of death (March 6, 2025)- Coverage by the Daily Pennsylvanian
- How News Coverage Distorts America's Leading Causes of Death (February 13, 2025)- Annenberg repost of CSSlab blog; also reposted by SEAS and Penn Today

External:

- HMN 2025: How information protection distorts America's main causes of loss of life (February 25, 2025)- Health Medicine Network
- How news coverage distorts America's leading causes of death (February 24, 2025)- Medical Xpress

Trends of Violence in Movies During the Past Half Century

External:

- The words in movies are shifting, and it's not good news (2/27/2025)- MSN
- Violence and death getting mentioned more often in movie scripts, study finds (January 2, 2025)- ABC Action News
- Now Streaming: More Murders in the Movies (January 6, 2025)- The Doctor Will See You Now
- Do movies really contain more violence than before? (January 6, 2025)- ETX Studio
- Movies are using more violent language including 'murderous verbs' (January 2, 2025)- Earth.com
- Murder, she spoke: Violence and death getting mentioned more often in movie scripts, study finds (January 2, 2025)- KVIA
- A Rising Trend of 'Murderous Verbs' in Movies Over 50 Years (January 1, 2025)- Newsgram
- Ohio State University study reveals surge in 'Murder' verbs in movies since 1970 (December 31, 2025)- The Jerusalem Post
- Has Violence In Movies Increased Over The Years? (December 31, 2024)- RTTNews
- Rise of 'murderous verbs': 160,000 film study exposes growing on-screen violence (December 31, 2024)- Interesting Engineering
- Movies talk more about murder and killing (December 31, 2024)- Bluewin
- 50-Year Rise in Murderous Verbs in Film (December 31, 2024)- Mirage
- Movie Violence Has Increased Across All Genres Over Past 50 Years, Massive Study Finds (December 30, 2024)- Ohio State University
- Violence in films has soared since 1980s, say scientists (December 30, 2024)- MSN
- Rise in talk about killing in films raises health concerns, researchers say (December 30, 2024)- The Guardian
- A rising trend of 'murderous verbs' in movies over 50 years (December 30, 2024)- Ohio State University
- Analysis of 160,000 films shows rise in 'murderous verbs' since 1970 (December 30, 2024)- Phys.org
- News Release: Trends of violence in movies during the past half century (December 30, 2024)- EurekAlert!
- Violence in films has soared since 1980s, say scientists (December 30, 2024)- The Telegraph

The Common Sense Project

Penn:

- Penn researchers develop tool to study 'commonsensicality' with public online survey (November 6, 2024)- The Daily Pennsylvanian

External:

- I took this new viral personality test – and realised something shocking about myself (October 2, 2024)- The Independent

- I took a common sense test – and my result appalled me (September 30, 2024)- The Guardian
- The fascinating truth about why common sense isn't really that common (September 24, 2024)- NewScientist
- Quantifying common sense: New research suggests it's not so common (April 23, 2024)- Northeastern University Science Magazine

Team Communication Toolkit

Penn:

- The mechanics of collaboration (August 30, 2024) - Penn Today

Media Bias Detector

Penn

- Mack Institute and CSS Lab Team Up On Media Bias Detection (April 29, 2025)- Mack Institute News
- This Tool Uses AI to Detect Bias in News Articles (November 1, 2024)- Knowledge at Wharton
- Detecting Media Bias (August 26, 2024)- Penn Gazette
    - Reposted by the Penn Population Studies Center
- Duncan Watts and CSSLab's New Media Bias Detector (June 28, 2024)- Penn Today
    - Reposted by Annenberg, Penn Engineering, and Knowledge at Wharton; also featured in Penn Today's A look back at 2024 through Penn Today stories
- AI-Powered Bias Detector Transforms News Analysis- Annenberg
    - Reposted by Penn Population Studies Center
- Mapping Media Bias: How AI Powers the Computational Social Science Lab's Media Bias Detector (June 25, 2024)- Penn Engineering
    - Reposted by Annneberg and EurekAlert!
- Computational Social Science Lab Bias Detection Video- Annenberg

External

- As colleges grapple with AI's pitfalls, U. of Delaware uses technology to transform faculty lectures into interactive study aides (August 7, 2024)- The Philadelphia Inquirer
- Media Bias: Understanding and Avoiding Skewed News (July 28, 2024)- Ted Jordan Consulting Blog
- Philly daily roundup: AWS donates $1M to CHOP; 3D modeling at a boxing gym; Penn prof's Media Bias Detector (June 28, 2024)- Technical.ly
- Media Bias Detector: New AI Tool Provides Insights Into How News Outlets Report on Various Topics (June 26, 2024)- Tech Times
- Mapping media bias: How AI powers a new media bias detector (June 25, 2024)- Phys.org
- This tool detects media bias using OpenAI's GPT-4- HNR Studio

News Consumption Dashboard

Penn:

- Mapping How People Get Their (Political) News (June 10, 2024)- Annenberg
    - Reposted by Penn Engineering

Misunderstanding the harms of online misinformation

Penn

- What Public Discourse Gets Wrong About Misinformation Online (June 28, 2024)- Annenberg
    - Reposted by Knowledge at Wharton and Fair Observer

External

- Trump's Victory Has Opened the Disinformation Floodgates (November 25, 2024)- Foreign Policy
- Can democracy work without journalism? With the US election upon us, we may be about to find out (November 2, 2024)- The Guardian
- The Regime Thinks Free Speech is "Fanaticism" (October 30, 2024)- Mises Institute
- America's political polarization is driving its hunger for fake news (October 31, 2024)- The Korea Times
- America's political polarization is driving its hunger for fake news (October 26, 2024)- Bloomberg
- Are We Misinformed About Misinformation? (October 24, 2024)- Undark
- Misinformation is a threat to society – let's not pretend otherwise (October 8, 2024)- London School of Economics and Political Science
- You might discover a conspiracy theory on social media — but you're more likely to believe it if you hear it from a friend (October 1, 2024)- Nieman Lab
- Book Review: A Misinformation Researcher's Guide to the 'Carnival of Mirrors' (September 3, 2024)- Just Security
- Even Disinformation Experts Don't Know How to Stop It (July 14, 2024)- The New York Times
- America's virulent anti-vaccine lies (July 3, 2024)- The New York Times
- Supreme Court Cases on Content Moderation and Mifepristone Access (June 27, 2024)- Kaiser Family Foundation
- The secret digital behaviors of Gen Z (June 25, 2024)- Business Insider
- There's No Dodge Button for Disinformation (June 18, 2024)- Foreign Policy
- Truth Sandwiches: Disinfo on Disinfo (June 10, 2024)- Daily Kos
- New study: Few see disinformation online - and the debate around it lacks support from science (June 10, 2024)- Warp News

- Cutting off the advertising revenue of websites that spread hoaxes is the best way to end misinformation (June 8, 2024)- El País
- Ctrl-Alt-Speech: The Internet Is (Still) For Porn, With Yoel Roth (June 7, 2024)- Techdirt
- After Jan. 6, Twitter banned 70,000 right-wing accounts. Lies plummeted. (June 6, 2024)- The Washington Post
- Strangulating the advertising of websites that spread hoaxes is the best way to end misinformation (June 5, 2024)- Pledge Times

Quantifying the impact of misinformation and vaccine-skeptical content on Facebook

Penn:

Re-examining Misinformation: How Unflagged Factual Content Drives Vaccine Hesitancy (May 30, 2024)-

- Reposted by Annenberg, Penn Today, and Knowledge at Wharton

External:

- The big idea: do we worry too much about misinformation? (March 17, 2025)- The Guardian
- Here's How Quickly Could Polio Return to the U.S. without Vaccines (January 30, 2025)- Today Headline
- Polio Vaccines Saved America from the Disease. What Happens If We Lose Them? (January 30, 2025)- Scientific American
- Fake news is driving us apart amid disaster – but slanted news is slowly drowning our democracy (January 12, 2025)- Salon
- Trump's Black voter bump misses the big picture (October 30, 2024)- The Washington Post
- Combating Misinformation Runs Deeper Than Swatting Away 'Fake News' (September 30, 2024)- Scientific American
- Childhood vaccinations seen as "not important," Gallup poll finds (August 7, 2024)- Salon
- Study shows impact of misleading headlines from mainstream news (July 9, 2024)- MIT Sloan School of Management
- How should we address stories of death after vaccination? (June 13, 2024)- Your Local Epidemiologist
- How a blunder by a respected medical journal is fueling an anti-vaccine lie (June 11, 2024)- Los Angeles Times
- Reviewing New Science on Social Media, Misinformation, and Partisan Assortment (June 9, 2024)- Tech Policy Press
- TechCrunch Minute: How misinformation gets amplified on social media (June 4, 2024)- TechCrunch
- Unflagged vaccine-skeptical content on Facebook causes significant vaccine hesitancy, study finds (June 3, 2024)- News-Medical.Net

- Facebook's Fact Checkers Failed to Flag Gray Vaccine Information (June 3, 2024)- Vax-Before-Travel
- Auroras Are on the Horizon, and Bird Flu Is on the Menu (June 3, 2024)- Scientific American
- Clickbait and Misinformation (June 3, 2024)- Neurologica Blog
- Misinformation works and a handful of social 'supersharers' sent 80% of it in 2020 (June 2, 2024)- beSpacific
- Study: Truthful yet misleading Facebook posts drove COVID vaccine reluctance much more than outright lies did (May 31, 2024)- Center for Infectious Disease Research and Policy (CIDRAP) at the University of Minnesota
- Mainstream COVID-19 Misinformation Caused More Facebook Harm Than Fake News (May 31, 2024)- Mirage News
- One Action Has Made a Significant Impact on How Misinformation Spreads Online (May 31, 2024)- ScienceAlert
- Studied on social media: The real-world impacts of factual but misleading content & the characteristics of "supersharers" (May 30, 2024)- EurekAlert!
- Unflagged stories drive vaccine hesitancy more than debunked misinformation (May 30, 2024)- News-Medical.Net
- Misleading COVID-19 headlines from mainstream sources did more harm on Facebook than fake news (May 30, 2024)- MIT Sloan School of Management

Other Coverage:

- AI Across Disciplines event highlights the power of 'breadth and connectivity' at universities (April 2, 2025)
- Best Of: How TV, Twitter and TikTok Remade Our Politics (December 13, 2024)- The Ezra Klein Show, The New York Times references a lab paper from 2022 titled: Quantifying partisan news diets in Web and TV audiences
- Grumpy Voters Want Better Stories—Not Statistics (November 20, 2024)- Scientific American
- University of Pennsylvania Gets $10 Million Funding, Launches New Center for Media, Technology and Democracy (September 14, 2024)- Philadelphia Today
- University of Pennsylvania Launches Penn Center on Media, Technology, and Democracy (September 5, 2024)- Wharton News
- Debunking Disinformation Myths, Part 3: The Prevalence and Impact of Fake News (August 24, 2024)- Conspicuous Cognition

The Computational Social Science Lab was created in March 2021 as a joint venture of the School of Engineering and Applied Science, the Annenberg School for Communication, and the Wharton School. We seek novel, replicable insights into societally relevant problems by applying computational methods to large-scale data. Through our research infrastructure, industry partnerships, and network of collaborators, we also aim to facilitate progress in computational social science more generally.

CSSLab

Computational Social Science Lab
CSSLab@seas.upenn.edu