

Algorithm Transparency through the Fair Credit Reporting Act (FCRA)

[Karl Schmeckpeper](#)^{1,3*}, [Sonia Roberts](#)^{1,3*+}, [Mathieu Ouellet](#)^{1,3*}, [Matthew Malencia](#)^{1,3*}, [Divya Jain](#)^{1,3*}, [Walker Gosrich](#)^{1,3*}, [Val Bromberg](#)^{2,3*}

¹University of Pennsylvania, School of Engineering and Applied Sciences

²University of Pennsylvania, College of Arts and Sciences

³Penn Science Policy and Diplomacy Group

*Authors contributed equally. Names listed in reverse-alphabetical order.

+The author submitted this article in a personal capacity and was accepted prior to joining the *JSPG* editorial board. They did not contribute to selecting or editing the article as an editorial board member.

<https://doi.org/10.38126/JSPG180415>

Corresponding Author: gosrich@seas.upenn.edu

Keywords: housing discrimination; fair housing; algorithm transparency; algorithm bias; artificial intelligence (AI); Federal Trade Commission (FTC)

Executive Summary: Racial discrimination in housing has long fueled disparities in homeownership and wealth in the United States. Now, automated algorithms play a dominant role in rental and lending decisions. Advocates of these technologies argue that mortgage lending algorithms reduce discrimination. However, “errors in background check reports persist and remain pervasive,” and algorithms are at risk for inheriting prejudices from society and reflect pre-existing patterns of inequality. Additionally, algorithmic discrimination is often challenging to identify and difficult to explain or prosecute in court. While the Federal Trade Commission (FTC) is responsible for prosecuting this type of discrimination under the Fair Credit Reporting Act (FCRA), their enforcement regime “has inadequately regulated industry at the federal and state level and failed to provide consumers access to justice at an individual level,” as evidenced by its mere eighty-seven enforcement actions in the past forty years. In comparison, 4,531 lawsuits have been brought under the FCRA by other groups in 2018 alone. Therefore, the FTC must update its policies to ensure it can identify, prosecute, and facilitate third-party lawsuits against a primary driver of housing discrimination in the 21st century: discrimination within algorithmic decision making. We recommend that the FTC issue a rule requiring companies to publish a data plan with all consumer reporting products. Currently, the FTC recommends that companies make an internal assessment of the components of the proposed data plan to ensure that they are not in violation of the FCRA. Therefore, requiring that these plans be published publicly does not place undue burden on companies and empowers consumers to advocate for themselves and report unfair practices to the FTC. Coupled together, these will reduce the costs of investigation and enforcement by the FTC and decrease the discriminatory impact of automated decision systems on marginalized communities.

I. Introduction and background

The current COVID-19 pandemic has exacerbated entrenched and long-standing racial

discrimination in housing in the United States. Prior to the pandemic, Black homeownership was at a record low of 40.6%. Though that figure

increased to 47% in 2020, data from the US Census Bureau shows that Black Americans still have the lowest rate of homeownership compared to other racial groups, while white Americans have the highest rate of homeownership at 76% (Callis and Kresin 2018). In fact, the gap in homeownership between Black and white Americans is higher today than it was in 1960, when race-based discrimination was legal (Choi et al. 2019). Housing is a significant contributor to intergenerational wealth—in 2019, the median net worth of homeowners was nearly forty times that of those who did not own a property—and the gap in homeownership is a marker of worsening wealth disparity between families of color and white families (Board of Governors of the Federal Reserve System 2019).

One of the main drivers of the homeownership and thus, wealth gap is discriminatory lending practices. Most banks use Automated Underwriting Systems (AUS) to evaluate applicants in a timely fashion. These systems appeared in 1995 and were quickly becoming ubiquitous six years later in 2002 (Gates, Perry, and Zorn 2002). The role of AUSs has switched from being a valuable helper in the underwriting process to being a fully automated online lending tool used by Fintech, a fast-growing market that has quadrupled in the last four years (Fuster et al. 2019). Other industries have also capitalized on fully automated algorithmic decisions. In 2019, approximately 45% of 2,000 large mortgage lenders relied on online or app-based interfaces to originate mortgages (Choi, Kaul, and Goodman 2019). Unfortunately, these new automated approaches have not removed racial disparities from lending practices. From 2009 to 2015, lenders using automated algorithms rejected approximately 1.3 million applications from credit-worthy applicants of color (Bartlett et al. 2019). A recent study by Lending Tree found that the home loan denial rate, as mediated by algorithms, for Black Americans was 17.4%, more than double the 7.9% for non-Hispanic whites. Lenders charged people of color an extra \$765M per year in interest (Bartlett et al. 2019).

Discriminatory practices not only affect homeownership, but also housing stability of renters in communities of color. Prior to the pandemic, studies found that neighborhoods with

more Black and Latinx renters experienced higher eviction rates. In the COVID-19 crisis, a greater proportion of Black and Latinx adults experienced job loss in comparison to their non-Hispanic white counterparts. Within the five industries most vulnerable to COVID-19 (food & accommodation, construction, entertainment, retail, and other services) (Housing Finance Policy Center 2020), 45.5% of Latinx workers and 31.8% of Black workers are renters, leaving them more vulnerable to eviction. While the Coronavirus Aid, Relief, and Economic Security (CARES) Act of March 2020 implemented an eviction moratorium preventing landlords from initiating eviction proceedings due to nonpayment of rent, it did not prevent the accrual of housing related debt. This puts people at risk of eviction when the federal moratorium ends on October 3, 2021. In fact, researchers at JP Morgan Chase Institute found that renters need a greater financial safety net than homeowners during the pandemic (Greig, Zhao and Lefevre 2021). The accumulation of housing-related debt and looming end of eviction moratoriums coupled with higher rates of COVID-19 cases, hospitalizations, and deaths in these communities has only exacerbated existing inequity in housing stability.

i. Promise of AI in eliminating biases

Decisions in the rental and lender markets are increasingly being made by algorithmic decision-making tools. They speed up decisions and have the potential to reduce face-to-face discrimination in markets prone to implicit and explicit biases, such as housing. It is estimated that 90% of landlords use automatically generated reports for tenant screening (Kirchner and Goldstein 2020). Despite this popularity, automated tenant screening can be prone to errors and mismatches. A review of federal lawsuits filed against screening companies found that these reports make hasty matches, wrongly labelling potential tenants as criminals or sex offenders (Nelson 2019; Kirchner and Goldstein 2020). In March 2019, CoreLogic, a tenant screening company that deploys automated algorithms, was held liable for discrimination claims because the algorithm analyzed arrest records, disability, race, and national origin, which is illegal under the Fair Housing Act (Fair Housing Act 1968).

In a recent study, researchers at Berkeley found that in the mortgage market, both algorithms and face-to-face lenders discriminate against borrowers of color. However, they found that algorithms discriminate 40% less than face-to-face lenders (Bartlett 2019). This is an indicator of the potential for algorithms to reduce inequity, but also a warning: although the algorithms may be less biased, their biases are *consistent* and are applied on a massive scale. Bias in algorithms can be caused by biases present in the data used to develop the algorithm. For example, underlying data may under-represent or fail to generalize to high-minority populations. Also, longitudinal data may capture historic racial inequities. Other underlying data may not explicitly capture race, but features of the data may include proxies for race (Lee and Singh 2021). In lending and tenant screening, existing housing data reflects years of segregation, Jim Crow laws, redlining, white flight, and other biased practices and regulations. This use of biased data risks perpetuating inequities at scale, cementing racial bias into code that is seldom questioned and difficult to interpret (Lee and Floridi 2021).

ii. Opportunity for intervention

Currently, under the FCRA, the FTC opens investigations into and prosecutes companies providing consumer reports that are biased against protected classes. With additional resources, the FTC can more effectively protect against discrimination in housing-related lending practices. However, without proper regulatory intervention, inequity in homeownership and housing stability will persist, widening existing wealth inequities.

II. Argument and evidence

Although the FTC assumes a prominent role in the enforcement of the FCRA, it has only prosecuted eighty-seven furnishers and credit reporting agencies in the past forty years (Sickler 2015). Private consumer lawsuits for FCRA violations, such as class-action litigation, far exceed FTC enforcement (Feder and Muttreja 2016). Currently, the process within the FTC to discover and prosecute violations may take years. For example, a tenant screening company, RealPage, did not properly audit or scrutinize the accuracy of the criminal record information it was providing to

landlords from January 2012 to September 2017 (FTC v. RealPage Inc 2018). RealPage spent over five years incorrectly reporting criminal records for innocent applicants before the FTC was able to intervene. To expedite the process of identifying and investigating potential violations of the FCRA, the FTC should require companies to publish a data plan. Legislation at the state level has achieved significant success in data regulation and protection. However, implementing a federal rule will allow for stronger mitigation of algorithmic discrimination.

The California Consumer Privacy Act (CCPA), one of the strictest privacy laws in the US, was implemented on January 1, 2020 with bipartisan support. This law requires large companies such as Google to divulge the types of data they collect, the sources they use to collect that data, why that information must be collected or sold, and the third parties with which that information is shared. The act also allows consumers to consent to having their data sold (CCPA 2021). Companies that implement algorithms in their products also must abide by these rules, forcing them to reevaluate or face consequences for any discriminatory practices underlying the algorithms. Under this legislation, companies may be required to reveal what data is used to assist algorithmic decisions, such as browsing history, making it easier for the public and regulatory agencies to detect and prevent algorithmic bias. Accordingly, over fifty lawsuits have already invoked this act following its recent enactment (Kaye 2021).

Outside of the United States, the General Data Protection Regulation (GDPR) covers data collection and privacy in the European Union. This law gives citizens the right to access data that has been collected on them (Article 15), to know how the data is being used (Article 15), to rectify any inaccuracies (Article 16), and to not be subjected to decisions made entirely by automated systems (Article 22, Council regulation 2016, 679). The GDPR helps to equip individuals to find and stop biased or inaccurate algorithmic decision systems.

III. Policy options

We propose three policy options to improve the FTC's ability to prosecute unlawful discrimination in consumer reporting products that use complex

and opaque algorithms to automate decision making.

i. Option 1: Provide funding to form a team of experts to investigate and support consumer protection issues that arise from algorithms used for automated decision making

This team could be formed into a new Office housed within the Bureau of Consumer Protection (BCP) or could become part of the Office of Technology Research and Investigation. The proposed team would focus on issues related to automated decision systems used by businesses, providing the specialized knowledge required to analyze the impact of automated algorithms on consumer privacy, data security, credit information, and financial technology with respect to housing and other issues under the FTC's enforcement. We propose that **\$7.5M in new funding be allocated to establish the team, representing a budget increase of 2%**. This funding increase is equivalent to about twenty-five new full-time employees, according to the FTC's 2020 numbers. As the FTC has averaged a return of \$30.70 for every \$1 spent on consumer protection over the last four years, it is expected that the proposed team will fund itself (FTC 2020). This team would seamlessly integrate with other Offices providing technical expertise, such as the Office of Technology Research and Investigation (FTC 2015), which has helped the FTC win settlements for consumers in several high-impact cases involving advanced technology (FTC 2017A; FTC 2017B). A dedicated team working on algorithm use in housing would better equip the FTC to enforce its rules as they are applied to new automated technology that can obfuscate harmful bias.

Benefits

The inclusion of content experts in algorithmic bias into the FTC workflow will improve the enforcement of laws even when complex algorithms obscure violations. This option offers a significant return on investment and does not require new legislation.

Drawbacks

This option requires initial investment from Congress and does not make law violations by algorithms any easier to enforce.

ii. Option 2: The FTC should issue a rule requiring companies to publish a data plan with all consumer reporting products

The proposed data plan includes three components:

- 1) **The types of data that the company may use, and either the method of obtaining this data or the source from which it is obtained.** The company is not required to specify which data types are used in their proprietary algorithms, but they may not use data not listed in their data plan. This would allow the company some agility in technological development and some protection of intellectual property by giving the option to slightly obfuscate specific data types in use, while still respecting the FTC's requirement that certain data types cannot be used to make decisions about consumer reports.
- 2) **The level of explanatory detail that the consumer can expect if the product is used to evaluate them.** The FCRA requires that a consumer receive a detailed and specific explanation when denied something of value. When applying this precedent to other automated decision systems, the appropriate level of explanatory detail may vary slightly. However, it should always provide information to the consumer that allows them to improve their score on the automated system in the future. The exact types of explanations that a consumer can receive should be specified in the data plan.
- 3) **An auditing plan to ensure that the product is not biased.** The company may either provide a detailed description of the internal tests conducted and how the results may be publicly accessed or submit to third-party auditing at the company's expense. The FTC already prohibits using protected classes of data, not providing sufficient explanatory detail, and producing biased results. In order to avoid prosecution, the FTC already recommends that companies consider each component of the proposed data plan. Publishing a data plan should, therefore, not create any undue hardship for a company. The data

plan should also significantly accelerate the investigation of complaints, thus reducing money spent on investigations.

Benefits

Public release of data makes enforcement easier for the FTC, empowers consumers to identify and report violations to the FTC, and is not an undue burden on companies.

Drawbacks

This option requires an FTC rule change and the development of standards related to the data plan. It may require additional investment from companies.

iii. Option 3: The FTC should issue a rule requiring third-party auditing of all high-risk commercial automated decision systems

While companies using low-risk automated decision systems may choose to develop their own evaluations for bias, companies using high-risk automated decision systems will need to pay for their automated decision systems to undergo a third-party audit every year. We define “high-risk” systems as those whose reports could affect 1) quality of life, 2) access to housing, or 3) access to employment for the consumer.

Third-party auditing would require additional infrastructure. The FTC would work with the National Institute of Standards and Technology to develop benchmarking datasets and tests on a yearly basis which can be used to review companies that provide third-party auditing services and provide accreditation for the next year. The datasets and tests would be devised to separately test whether the auditing company can detect bias in a dataset and bias in an algorithm.

Third-party auditing can be used to provide an appropriate financial incentive to the companies developing automated decision systems. If the company does not publish an appropriate data plan or scores poorly on its audit, it must pay for the full cost of the audit plus a fine and it may not legally use the automated decision system until it passes the audit. Third-party auditing may be indicated for a low-risk company if the company’s data plan is insufficiently detailed or it admits to using illegal data types, or if the company is already being

investigated for wrongdoing by the FTC. If the low-risk company is found to be in the wrong during this investigation or its data plan is the cause for the investigation or audit, the low-risk company should be required to pay for the audit.

Paying for audits may be difficult for small companies and startups. To mitigate these effects, we propose that even high-risk companies whose algorithms are used on fewer than 1,000 people will not need to pass an audit. The FTC should also offer easily attainable grants for companies to pay for their first audit.

Benefits

This option strengthens incentive structures so that companies are more likely to avoid bias in their systems, makes enforcement significantly easier for the FTC by requiring companies to specify the data types in use and report the explanatory details provided to consumers about decisions. It also standardizes enforcement so that companies are on even ground.

Drawbacks

This option requires an FTC rule change, the development of standards related to auditing in coordination with other federal agencies, and infrastructure and regulation around 3rd party auditors. It could be more expensive for companies and would require additional investment in the FTC.

VI. Policy recommendation

We recommend **Option 2: The FTC should issue a rule requiring companies to publish a data plan with all consumer reporting products.** This policy option best targets unfair business practices prohibited by the FCRA, which are often difficult to prosecute because they are obscured by complex automated decision systems. This option places the responsibility of transparency on companies that use these systems, which would not place undue burden on companies that are following current laws or on consumers to find and report unfair practices. This allows third-party auditors, government regulators, and consumers to better detect algorithmic biases, improving the accountability of companies that use automation in making housing decisions. This option enables the FTC to better enforce its rules against racial

discrimination in housing by improving the investigation and enforcement processes in cases with automated system.

References

- Barocas, Solon, and Andrew D Selbst. 2016. *Big data's disparate impact*. Calif. L. Rev. 104 (2016): 671. <http://dx.doi.org/10.15779/Z38BG31>
- Bartlett, Robert, Adair Morse, Richard Stanton and Nancy Wallace. 2019. *Consumer-lending discrimination in the FinTech era (No. w25943)*. National Bureau of Economic Research, 2019.
- Board of Governors of the Federal Reserve System. 2019. *Survey of Consumer Finances (SCF)*. Retrieved from <https://www.federalreserve.gov/econres/scfi/ndex.htm>
- California Consumer Privacy 2018. *TITLE 1.81.5. California Consumer Privacy Act of 2018*, 1798.100 - 1798.199.100, Retrieved from <https://oag.ca.gov/privacy/ccpa>.
- Callis, R, and M Kresin. 2018. *Quarterly residential vacancies and homeownership, first quarter 2018* (Report No. CB18-57). Retrieved from the United States Census Bureau. <https://www.census.gov/housing/hvs/files/currenthvspress.pdf>
- Board of Governors of the Federal Reserve System. 2019. *Survey of Consumer Finances (SCF)*. Retrieved from <https://www.federalreserve.gov/econres/scfi/ndex.htm>
- California Consumer Privacy 2018. *TITLE 1.81.5. California Consumer Privacy Act of 2018*, 1798.100 - 1798.199.100, Retrieved from <https://oag.ca.gov/privacy/ccpa>.
- Callis, R, and M Kresin. 2018. *Quarterly residential vacancies and homeownership, first quarter 2018* (Report No. CB18-57). Retrieved from the United States Census Bureau, <https://www.census.gov/housing/hvs/files/currenthvspress.pdf>.
- Choi, Jung, Karan Kaul and Laurie Goodman. 2019. *FinTech Innovation in the HomePurchase and Financing Market*. The Urban Institute, Washington, DC. Retrieved from https://www.urban.org/sites/default/files/publication/100533/fintech_innovation_in_the_home_purchase_and_financing_market_0.pdf
- Choi, Jung Hyun, Alanna McCargo, Michael Neal, Laurie Goodman and Caitlin Young. 2019. *Explaining the Black-White homeownership gap*. The Urban Institute. Washington, DC. Retrieved from <https://www.urban.org/research/publication/explaining-black-white-homeownership-gap-closer-look-disparities-across-local-markets>
- Fair Housing Act, Title VIII of the Civil Rights Act of 1968, U.S. Code 7 C.F.R. § 1901.203* (1968)
- Feder, Meir, and Rajeev Muttreja. 2016. *Understanding the FCRA*. Practical Law, Litigation. Mai/April: 48-57. <https://www.jonesday.com/-/media/files/publications/2016/04/understanding-the-fair-credit-reporting-act-ipract/files/understanding-the-fcra/fileattachment/understanding-the-fcra.pdf>
- Federal Trade Commission. 2011. *40 Years of Experience with the Fair Credit Reporting Act: An FTC Staff Report with Summary of Interpretations*. P024808. Washington, DC, United States.
- Federal Trade Commission. 2015. *BCP's Office of Technology Research and Investigation: The next generation in consumer protection*. Retrieved from <https://www.ftc.gov/news-events/blogs/business-blog/2015/03/bcps-office-technology-research-investigation-next>
- Federal Trade Commission. 2017A. *FTC Approves Final Consent Order with Online Company Charged with Deceptively Tracking Consumers Online and Through Mobile Devices*. Retrieved from <https://www.ftc.gov/news-events/press-releases/2017/04/ftc-approves-final-consent-order-online-company-charged>
- Federal Trade Commission. 2017B. *Mobile App Settles FTC Allegations That it Failed to Deliver Promised Cash Rewards for Meeting Exercise and Diet Goals*. Retrieved from <https://www.ftc.gov/news-events/press-releases/2017/09/mobile-app-settles-ftc-allegations-it-failed-deliver-promised>
- Federal Trade Commission v. RealPage, Inc. 2018. 3:18-cv-02737-N 20. United States District Court for Northern District of Texas Dallas Division.
- Federal Trade Commission. 2020. *Fiscal Year 2021 Congressional Budget Justification*. Retrieved from <https://www.ftc.gov/reports/fy-2021-congressional-budget-justification>

- Fuster, Andreas, Matthew Plosser, Philipp Schnabl and James Vickery. 2019. *The role of technology in mortgage lending*. The Review of Financial Studies 32, no. 5 : 1854–1899. <https://doi.org/10.1093/rfs/hhz018>
- Gates, Susan Wharton, Vanessa Gail Perry and Peter M Zorn. 2002. *Automated underwriting in mortgage lending: Good news for the underserved?* Housing Policy Debate 13, no. 2 : 369–391. <https://doi.org/10.1080/10511482.2002.9521447>
- Greig, Fiona, Chen Zhao and Alexandra Lefevre. 2021. *Renters vs. Homeowners: Income and Liquid Asset Trends during COVID-19*. JPMorgan Chase Institute. Retrieved from <https://www.jpmorganchase.com/institute/research/household-debt/renters-homeowners-income-and-liquid-asset-trends-during-covid-19> <http://dx.doi.org/10.2139/ssrn.3807563>.
- Housing Finance Policy Center. 2020. *Lenders, Loan Channels & Industry/Occupation by Income & Race/Ethnicity*. The Urban Institute, Washington, DC. Retrieved from https://www.urban.org/sites/default/files/2020/05/08/lenders_loan_channels_industry_occupation_by_income_and_race_ethnicity_0.pdf
- Kaye, Kate. "People File Lawsuits to Test Boundaries of California's Privacy Law." Digiday, February 17, 2021. <https://digiday.com/media/people-file-lawsuits-to-test-boundaries-of-californias-privacy-law/>.
- Kirchner, Lauren and Goldstein, Matthew. 2020. *How Automated Background Checks Freeze Out Renters*. New York Times, 28 May, p.BU1. Academic OneFile, Retrieved from <https://www.nytimes.com/2020/05/28/business/renters-background-checks.html>.
- Lee, Michelle Seng Ah, and Jatinder Singh. 2021. *Risk identification questionnaire for unintended bias in machine learning development life cycle*. SSRN. <http://dx.doi.org/10.2139/ssrn.3777093>
- Lee, Michelle Seng Ah, and Luciano Floridi. 2021. "Algorithmic fairness in mortgage lending: from absolute conditions to relational trade-offs." *Minds and Machines* 31: 165-191. <https://doi.org/10.1007/s11023-020-09529-4>
- Maurer, Roy. 2019. *FCRA Litigation on Track to Reach a New High*. SHRM. Retrieved from <https://www.shrm.org/resourcesandtools/hr-topics/talent-acquisition/pages/fcra-litigation-track-reach-new-high.aspx>
- National Fair Housing Alliance. 2020. *Fair Housing in Jeopardy*. Retrieved from <https://oag.ca.gov/privacy/ccpa>.
- Nelson, Ariel. 2019. *Broken records redux: how errors by criminal background check companies continue to harm consumers seeking jobs and housing*. Boston, MA, National Consumer Law Center.
- Regulation 2016/679, *On the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC*, European Parliament and of the Council of 27 April 2016, OJ L 119, 4.5.2016, p. 1–88.
- Sickler, Alexandra P Everhart. 2015. *The (Un) fair Credit Reporting Act*. Loy. Consumer L. Rev. 28 : 23

Karl Schmeckeper is a Ph.D. student in Computer Science at the University of Pennsylvania. He works with Prof. Kostas Daniilidis in the Daniilidis Group, a computer vision and robotics research group in the General Robotics, Automation, Sensing, and Perception (GRASP) Lab. His Ph.D. work focuses on enabling robots to learn by encoding priors into machine learning models, by improving exploration, and by leveraging human demonstrations. Karl has a bachelor's degree in computer science from the University of Massachusetts Amherst and a master's degree in robotics from the University of Pennsylvania.

Sonia Roberts is a Ph.D. student in Electrical and Systems Engineering at the University of Pennsylvania. She works with Prof. Dan Koditschek in Kod*Lab, a legged robotics research group in the General Robotics, Automation, Sensing, and Perception (GRASP) Lab. Her Ph.D. work focuses on locomotion on soft, fragile substrates like sand, and how robots can respond to these difficult environments without needing complex internal models of them. Sonia's undergraduate degree is from Vassar College where she studied Cognitive Science, and she spent two years working as a research technician at the Howard Hughes Medical Institute's Janelia Farm Research Campus before coming to Penn.

Mathieu Ouellet is a Ph.D. student in Electrical and Systems Engineering at the University of Pennsylvania. He works with Prof. Danielle Bassett in the Complex Systems Lab UPenn, a multidisciplinary laboratory

interested in the many facets of Complex Systems. He also works with Prof. Lee Bassett in the Quantum Engineering Lab @Upenn, a research group focused on understanding and controlling nanoscale quantum properties in semiconductor materials. His Ph.D. work focuses on understanding the properties of systems of multiple interacting components. Mathieu has a bachelor's degree in computer science and physics from the University of Quebec. He also completed a master's degree in applied mathematics, where he studied supersymmetry and the power spectral parameter representation in the context of quantum mechanics.

Matthew Malencia is a robotics researcher, an AI educator, and a science policy advocate. He is pursuing a Ph.D. in Electrical and Systems Engineering at the University of Pennsylvania with advisors Dr. Vijay Kumar and Dr. George Pappas, and he is a visiting researcher at the University of Cambridge with Dr. Amanda Prorok. His research focuses on fairness and diversity on robot teams. Matthew is the co-director of AI4ALL@GRASP, a summer program that teaches Philadelphia area high school students the fundamentals of artificial intelligence and robots. Lastly, Matthew works with the Science Policy & Diplomacy Group at the University of Pennsylvania on AI ethics and policy. He has written multiple AI memos, engaged with policymakers on AI topics, and delivered an intervention on the floor of the United Nations during a science and technology forum.

Divya Jain is a Ph.D. student at the Center for Injury Research and Prevention at the Children's Hospital of Philadelphia and the University of Pennsylvania and is co-advised by Dr. Kristy Arbogast and Dr. Catherine McDonald. Her doctoral work centers on objective methods for diagnosing and monitoring recovery from concussion in adolescents, with a focus on determining when adolescents should return to driving motor vehicles after their injury. Divya holds a bachelor's degree in bioengineering from the University of Maryland, College Park and is an HHMI-NIBIB Interfaces Scholar and a Dwight D. Eisenhower Transportation Fellow.

Walker Gosrich is a Ph.D. Candidate studying robotics in the GRASP Lab at the University of Pennsylvania. He is an NSF Graduate Fellow working under Dr. Mark Yim, and researches planning and control for teams of robots using tools such as machine learning and graph theory. He is the Policy Chair of the Penn Science Policy and Diplomacy Group. His policy interests center on how policy affects the development and implementation of new robotics and AI technologies. Walker received his Bachelors in Mechanical Engineering from the University at Buffalo in 2018.

Val Bromberg is a second-year undergraduate student majoring in Biology, with a concentration in Mechanisms of Disease, and minoring in Chemistry at the University of Pennsylvania's College of Arts and Sciences. She works with Dr. Harvey Friedman in the Friedman lab, a lab which aims to provide protection from Herpes simplex viruses through the creation of a vaccine utilizing mRNA technology. Additionally, she works with a research group led by Dr. Brendan Kelly and Dr. Michael David to assess outcomes in COVID-19 patients who were given empiric antibiotics and/or developed bacterial superinfections throughout their hospital admissions.