GOVERNING THE ALGORITHMIC TURN: LYFT, UBER, AND DISPARATE IMPACT

Bao Kham Chau

INTRODUCTION

The arrival of the digital age, with a drastic increase in computational power and new forms of software-based technological innovations, ushered in the beginning of what is now commonly known as the Big Data revolution. Hitherto confined to analyzing small datasets, advances in parallel computing, data science, and Artificial Intelligence/Machine Learning (AI/ML) opened new epistemological doors to knowledge production based on Zettabytes of data. These new insights are perceived to be so significant that the British mathematician Clive Robert Humby referred to data as "the new oil."[1] By taking advantage of these new resources, BigTech companies and Silicon Valley startups have been able to reshape the global economy, leading to significant changes to the daily lives of the average American.[2] Amazon, for example, utilized AI/ML algorithms to personalize advertisements (ads) for its users, improve its delivery logistics, and disrupt the traditional retail model.[3] Similarly, advances in AI/ML also transformed the legal profession by introducing new tools into a lawyer's toolkit. Attorneys can now use ML-powered software to find "relevant" case law,[4] predict outcomes in tax law to facilitate settlement,[5] and use eDiscovery platforms to reduce resources spent on the discovery process.[6]

This algorithmic turn, however, is rife with bias. This paper addresses this issue by investigating the ways in which some private companies' usage of AI/ML implicates antidiscrimination laws. Specifically, this project will examine Transportation Network

---

[1] Charles Arthur, *Tech giants may be huge, but nothing matches big data*, The Guardian (2013), https://www.theguardian.com/technology/2013/aug/23/tech-giants-data (last visited Feb. 18, 2022).
[2] *See*, *e.g.*, Nicholas Carr, *Is Google Making Us Stupider?*, Atlantic (July 2008), https://www.theatlantic.com/magazine/archive/2008/07/is-google-making-us-stupid/306868/ (last visited Jan. 24, 2022).
[3] *See* The Amazon Effect: Impacts on Shipping and Retail, https://www.shorr.com/packaging-news/2015-06/amazon-effect-impacts-shipping-and-retail (last visited Jan. 27, 2022).
[4] *See* Michael A. Livermore et al., *Law Search in the Age of the Algorithm*, 2020 Mich. St. L. Rev. 1183 (2020).
[5] *See* Benjamin Alarie et al., *Using Machine Learning to Predict Outcomes in Tax Law*, 58(3) Can. Bus. L. J. 231 (2016).
[6] *See* The Sedona Conference, *The Sedona Principles*, 19 Sedona Conf. J. 1 (2018).

Companies (TNCs), such as Lyft and Uber, and explore how their dynamic pricing and matchmaking models could reinforce or exacerbate existing inequalities. This paper will thus be divided into three sections. The first section will look at algorithmic bias in general and its possible impact on ML software of TNC firms. The second section will explore the viability of using the disparate impact framework to protect consumers against algorithmic bias in TNC software, specifically, and in private companies, generally. The third section will recommend changes to the software development process to guard against algorithmic bias. Finally, the paper will conclude with potential future research avenues.

I. ALGORITHMIC BIAS IN TNC MOBILE APPLICATIONS

TNCs such as Lyft and Uber connect, on demand, idle drivers with passengers. Although the ridesharing economy helps facilitate economic transactions, it also poses several new regulatory challenges. Specifically, research shows that TNCs increased both pollutants and traffic congestion. In California, for example, TNCs have been shown to "emit approximately 50 percent more [Greenhouse Gas] emissions per [Passenger Miles Traveled] than California passenger vehicles."[7] Research of San Francisco traffic also shows that, on average, Lyft and Uber contributed to more traffic congestions.[8] Additionally, TNCs also compete with and reduce ridership from public transportation. A study "suggest[s] that for each year after Transportation Network Companies (TNCs) enter a market, heavy rail ridership can be expected to decrease by 1.3% and bus ridership can be expected to decrease by 1.7%. This TNC effect builds with each passing year and may be an important driver of recent ridership declines."[9]

While these problems have spawned numerous discussions, the legal ramifications of algorithmic biases in TNCs have yet to be fully studied. To address this gap, this section will define algorithmic bias, explore possible ways algorithmic bias can affect machine learning algorithms, and examine how algorithmic bias in

[7] California Air Resources Board, *Clean Miles Standard: 2018 Base-year Emissions Inventory Report*, 42 (2019).
[8] *See* Gregory D. Erhardt et al., *Do transportation network companies decrease or increase congestion?*, 5(5) Sci. Adv. 1 (2019).
[9] Michael Graehler Jr. et al., *Understanding the Recent Transit Ridership Decline in Major US Cities: Service Cuts or Emerging Modes?*, https://trid.trb.org/view/1572517 (last visited Dec. 25, 2021).

machine learning algorithms impact TNC operations such as the determination of a trip's price and the matching of drivers and riders.

A. Artificial Intelligence and Algorithmic Bias

AI/ML software enables a computer to mimic human learning behavior in order to analyze a particular input and uncover patterns, to predict future data, or to act on the derived insights. At the heart of these software is a mathematical model built by algorithms based on sample training and cross validation data.[10] Algorithms can range from simple Bayesian inferences and Logistical Regression methodologies to more complex deep learning models like Convolutional Neural Networks.[11] Most ML algorithms can be used in both supervised and unsupervised learning. The former type takes in both labeled input and output data while the latter does not. In supervised learning, the algorithm "learns" from the training dataset by iteratively making predictions and adjusting for the correct output. In unsupervised learning, the algorithm would try to infer some structure or groups based on heuristics such as proximity of attributes.[12]

Despite the veneer of mathematical objectivity, algorithms can encapsulate bias that produces unfair results.[13] There are generally two types of algorithmic biases — human biases and system biases.[14] A type of human bias is unconscious out-group homogeneity bias, which can affect how a ML model is designed. A relatively homogeneous group of developers, for example, could build an algorithm that captures more nuances of their in-group attributes than their out-group counterparts. This, in turn, would make the model less predictive for members of the out-group than members of the in-group. Systemic bias, on the other hand, refers to data collection errors such as sampling bias. If the data are not collected randomly from the target group, for example, the AI/ML algorithm could mistakenly generalize insights that only belong to a

---

[10] *See* Christopher M. Bishop, *Model-based machine learning*, 371 Phil. Trans. Math. Phys. Eng. Sci. 1 (2013).
[11] *See* Jason Brownlee, *A Tour of Machine Learning Algorithms*, https://machinelearningmastery.com/a-tour-of-machine-learning-algorithms/ (last visited Dec. 27, 2021).
[12] *See* Julianna Delua, *Supervised vs. Unsupervised Learning: What's the Difference?*, https://www.ibm.com/cloud/blog/supervised-vs-unsupervised-learning (last visited Dec. 27, 2021).
[13] *See* Deborah Hellman, *Measuring Algorithmic Fairness*, 106(4) Va. L. Rev. 99 (2020)(exploring two metrics of algorithmic fairness — equal predictive value and error rate balance).
[14] Google Inc., Machine Learning Glossary, https://developers.google.com/machine-learning/glossary (last visited Feb. 28, 2022).

subset of the group to cover the entire dataset. This, in turn, would also make the model less predictive for the target population.

B. Algorithmic Bias in TNC Mobile Applications

The TNC applications (apps) are primarily divided into two types — drivers and passengers. Each type has a "frontend" and a "backend." The frontend component is what the user sees and interacts with. It is similar to the steering wheel of the user's car. Conversely, the backend component refers to things under the hood of the car. It consists of the server, database, and algorithms that process user inputs, produce the appropriate outputs, and send those outputs back to the frontend to be displayed to the user. While cognitive bias also affects the frontend,[15] this section will focus on how algorithmic bias can impact results produced by the backend. Specifically, it will analyze how Lyft pricing models and matchmaking algorithms could be affected by both human and system bias.[16]

a. Bias in TNC Dynamic Pricing Models

TNC companies like Lyft and Uber developed and rely on dynamic pricing models that automatically adjust the price per miles traveled depending on the number of drivers and rides requested in an area.[17] While dynamically changing the price based on the supply of drivers and the demands of passengers in an area is economically sound, researchers have shown the presence of racial discrimination in these dynamic pricing algorithms.[18] Akshat Pandey and Aylin Caliskan observed that, in Chicago, "[t]rips beginning or ending in neighborhoods with more non-white residents, or more residents living below the poverty line, have higher fares, as well as higher trip duration."[19] Although Lyft reiterated that they do not condone

---

[15] *See* Jaime Ventura, *Cognitive Bias & UX Design: the Good, the Bad, and the Ugly*, https://u.group/thinking/cognitive-bias-ux-design-the-good-the-bad-and-the-ugly/ (last visited Dec. 28, 2021).

[16] I chose Lyft because its Engineering blog (https://eng.lyft.com/) provides more transparency into its pricing and matchmaking models.

[17] *See* Uber, Inc., *How surge pricing work*, https://www.uber.com/us/en/drive/driver-app/how-surge-works/ (last visited Dec. 28, 2021); Lyft, Inc., *Personal Power Zones*, https://help.lyft.com/hc/e/articles/115012926807-Personal-Power-Zones (last visited Dec. 28, 2021).

[18] See Kyle Wiggers, *Researchers find racial discrimination in 'dynamic pricing' algorithms used by Uber, Lyft, and others*, https://venturebeat.com/2020/06/12/researchers-find-racial-discrimination-in-dynamic-pricing-algorithms-used-by-uber-lyft-and-others/ (last visited Dec. 28, 2021).

[19] Akshat Pandey and Aylin Caliskan, *Disparate Impact of Artificial Intelligence Bias in Ridehailing Economy's Price Discrimination Algorithms*, AAAI/ACM Conf. Artificial Intelligence 1, 6 (2020).

discrimination on the platform in any form, whether through algorithms or decisions made by their users, it did acknowledge that "technology can unintentionally discriminate."[20] Indeed, careful examination of how Lyft designed its pricing model reveals potential areas where human and system bias could contaminate the process.

Examination of the simulation system that Lyft uses to test out new dynamic pricing models, for example, reveals that it is a place where out-group homogeneity bias could be introduced. In Lyft's simulation system, "[s]imulations start with real historical passenger app sessions — not just rides, but also people who had the app open but decided not to request. For drivers, [Lyft engineers] use historically based geographical distributions to determine starting locations."[21] To create a simulated scenario, data scientists merely have to log on to the simulation platform, specify the ratio of passengers to drivers in a particular area, and run the scenario to experiment on which PrimeTime[22] price model would yield the desired result.

Total drivers in simulation by 10 second time step. Red represents the most drivers and green represents the fewest.

Figure 1: Lyft simulated scenario for the Seattle Metropolitan area[23]

---

[20] Donna Lu, *Uber and Lyft pricing algorithms charge more in non-white areas*, https://www.newscientist.com/article/2246202-uber-and-lyft-pricing-algorithms-charge-more-in-non-white-areas/ (last visited Jan. 10, 2022).
[21] Adam Greenhall, Experimentation in a ridesharing marketplace, https://eng.lyft.com/https-medium-com-adamgreenhall-simulating-a-ridesharing-marketplace-36007a8a31f2 (last visited Dec. 28. 2021).
[22] Lyft PrimeTime is an additional charge that is automatically added to the final total fare. The charge is added when the demand for Lyft rides is greater than the number of Lyft drivers on the road (e.g., during rush hours or after popular events).
[23] *Id.*

While this experimental setup is logically sound, if the data scientist running the experiment is a Seattle native, he or she might be more familiar with the number of passengers and drivers in that area. The experimenter could thus properly calibrate the scenario for the Seattle Metropolitan area while simply setting the default passengers to drivers ratio for other metropolitan areas. This could thus compromise the result of the experiment, introducing bias into the dynamic pricing model that is eventually chosen.

Measurement error bias could also affect the development of Lyft's dynamic pricing models. Data scientists and software engineers building the pricing model rely on accurate historical GPS data. Due to spotty cell phone coverage, however, these GPS data could be corrupted or incomplete. "Agilent Technologies, one of several firms that systematically track weak spots in cell phone coverage, identified cavernous gaps in [San Francisco], Oakland hills, Skyline Boulevard on the Peninsula and sections of the South Bay."[24] When drivers or passengers are in these dead zones, their Lyft apps would thus not be able to send their geo-location data to the backend.[25]

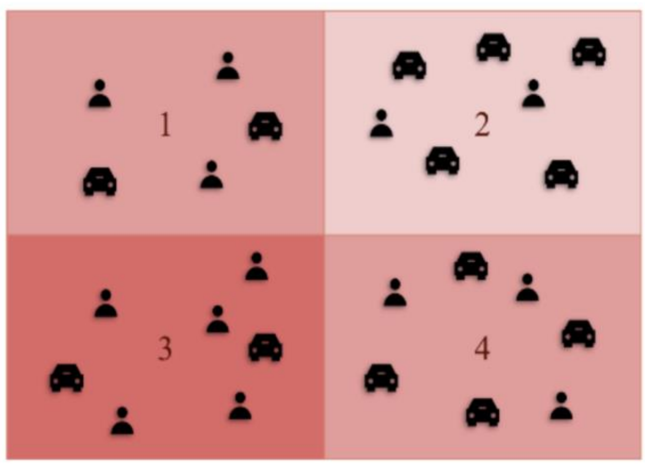

Figure 2: A hypothetical experiment with four geographical regions[26]

---

[24] Todd Wallack, *DEAD ZONES / WHERE CELL PHONES DON'T WORK / Companies' own data show wide gaps in coverage in Bay Area*, https://www.sfgate.com/bayarea/article/DEAD-ZONES-WHERE-CELL-PHONES-DON-T-WORK-2936411.php (last visited Dec. 28, 2021).
[25] *See* Yunjie Zhao et al., Sensor Data in Locations: Wi-Fi, https://eng.lyft.com/sensor-data-in-localization-wi-fi-329ea84db959 (last visited Dec. 28, 2021).
[26] Davide Grapis and Chris Sholley, *Dynamic Pricing to Sustain Marketplace Balance*, https://eng.lyft.com/dynamic-pricing-to-sustain-marketplace-balance-1d23a8d1be90 (last visited Dec. 28, 2021).

If the historical data contained gaps in GPS data or corrupted GPS data, the experiments conducted using these historical data would also be skewed. In the hypothetical experiment in Figure 2, for example, if region 1 is in a cell phone dead zone, then it is possible that the number of drivers or passengers in this zone is underrepresented. Any pricing model built based on this experimental setup would then encapsulate this measurement error bias.

b. Bias in TNC Matchmaking Algorithms

Many newspaper articles have criticized that TNCs, similarly to their taxi predecessors, redlined low-income and underserved communities.[27] Most of the criticisms, however, focused on Lyft and Uber giving drivers the ability to know every ride's destination in advance and to reject ride requests without penalty. There are "well-developed patterns of racial exclusion in the U.S., whether housing, land use policy or transportation … Folks of color, particularly black folks, are not able to get a taxi pickup as often in their neighborhoods."[28] Not many advocates, however, are aware of the ways in which algorithmic bias could influence TNCs' matchmaking algorithms — mathematical models that match drivers to passengers — and make it harder for drivers to be matched to passengers in low-income and underserved neighborhoods.

A closer look at how Lyft matches drivers to passengers reveals an area where both implicit and selection bias could be introduced. According to Lyft, "[w]hen you request a ride, Lyft tries to match you with the driver most suited for your route. To make a dispatch decision, [Lyft] first need[s] to ask: where are the drivers? Lyft uses GPS data from the drivers' phones to answer this question. However, the GPS data that [Lyft] get[s] is often noisy and does not match the road."[29]

---

[27] *See* Jim Motavalli, *Do Uber and Lyft Redline Low-Income Communities?*, https://www.cartalk.com/blogs/do-uber-and-lyft-redline-low-income-communities (last visited Dec. 28, 2021).

[28] Carolyn Said, *Uber's New Policies Could Encourage Discrimination, Advocates Fear*, https://greenlining.org/press/news/2020/ubers-new-policies-could-encourage-discrimination-advocates-fear/ (last visited Dec. 28, 2021).

[29] Marie Douriez, *James Murphy*, *Kerrick Staley, A New Real-Time Map-Matching Algorithm at Lyft*, https://eng.lyft.com/a-new-real-time-map-matching-algorithm-at-lyft-da593ab7b006 (last visited Dec. 28, 2021).

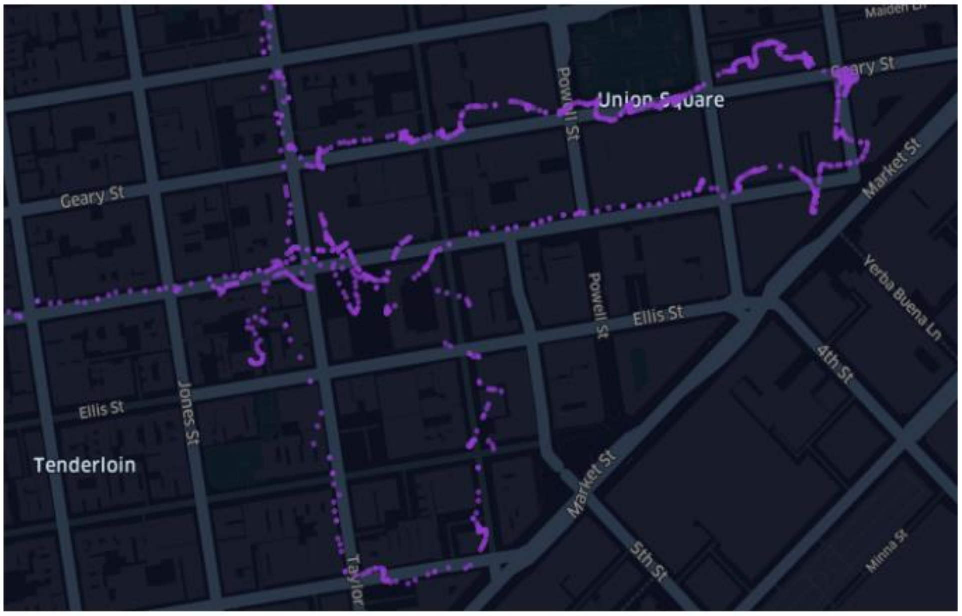

Location data is often noisy and does not match the road.

Figure 3: A hypothetical driver's self-reported location data, as sent by the mobile app[30]

While trying to locate drivers with those noisy GPS data, the Lyft matchmaking algorithm also simultaneously tries to determine potential matches between passenger routes. In order to reduce computational complexities of these two tasks, the Lyft engineers used a bucketing system called geohashing. "Geohashing is a technique for bucketing latitude and longitude coordinates. Using historical data from past Lyft rides, [Lyft] could record the average speed of [Lyft] rides from one geohash to another and store that in a simple hash-table lookup."[31]

[30] *Id.*
[31] Timothy Brown, *Matchmaking in Lyft Line — Part 2*, https://eng.lyft.com/matchmaking-in-lyft-line-691a1a32a008 (last visited Dec. 28, 2021).

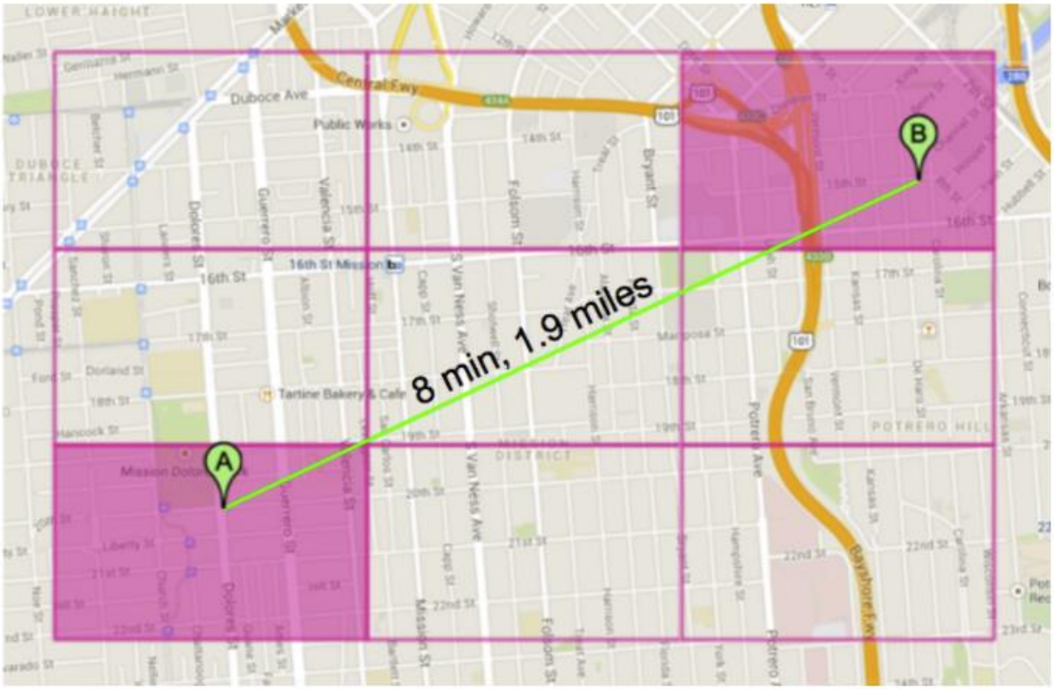

Geohash Level 6

Figure 4: A hypothetical geohashing scheme for a part of Downtown San Francisco[32]

While geohashing did drastically reduce the computational time, it also opened the matchmaking algorithm up to both implicit and selection bias. As seen in Figure 4, geohashing schemes are designed to improve algorithmic performance and do not necessarily need to take into consideration the socioeconomic conditions on the ground. As long as a geohash produces more efficient results, it will be used in the matchmaking algorithm. Since low-income neighborhoods have been historically underserved by both the taxi and TNC industries, geohashing could replicate and reinforce this dynamic. Thus, this purely technological solution could have introduced implicit and selection bias into the matching algorithm.

II. ALGORITHMIC GOVERNANCE AND THE DISPARATE IMPACT FRAMEWORK

As seen above, the algorithmic turn raises several challenges to the existing regulatory framework. As AI/ML powered software are being integrated into the administrative state, there are increasing demands for more algorithmic accountability and transparency.[33] Instead of advocating for new legislations, this paper

---

[32] *Id.*

[33] *See, e.g.*, Hannah Bloch-Wehba, *Access to Algorithms*, 88 Fordham L. Rev. 1265 (2020); Noah Bunnell, *Remedying Public-Sector Algorithmic Harms: The*

suggests using the current disparate impact framework to regulate algorithmic bias.

The disparate impact framework enables plaintiff to prove discrimination without needing to prove the intent to discriminate. Under *Griggs v. Duke Power Co.*, a company could be sued for violating antidiscrimination laws if the company's practices produced statistically significant discriminatory results.[34] Here, discrimination could be proven via statistics. The Equal Employment Opportunity Commission (EEOC) often relies on the "four-fifths rule", which presumes discrimination if the protected group's selection rate is less than or equal to four-fifths, or eighty percent, of the selection rate of the group with the highest selection rate.[35] If the company does not have a legitimate business purpose that could only be achieved via the practices that produced the discriminatory result, then the plaintiff would most likely prevail even if those practices are facially neutral.[36]

As algorithmic bias in ridesharing platforms could produce discriminatory results that are statistically significant, passengers could thus potentially rely on the disparate impact framework to sue. It is, however, unlikely that passengers could sue under Title VII of the 1964 Civil Rights Act (Title VII) or the 1967 Age Discrimination in Employment Act (ADEA). Indeed, both Title VII and ADEA only forbid employment discrimination "because of such individual's race, color, religion, sex, or national origin"[37] or "because of such individual's age."[38] Discrimination against customers, however, are not actionable under the two statutes.

To initiate disparate impact legal actions against algorithmic discrimination on TNC platforms, passengers should thus rely on Title II of the 1964 Civil Rights Act (Title II) and on the 1990 Americans with Disabilities Act (ADA). Title II outlawed discrimination or segregation on the ground of race, color, religion, or national origin in public accommodations engaged in interstate commerce,[39] while the ADA prohibits discrimination on the basis of disability in "the full and equal enjoyment of the goods, services,

---

Case for Local and State Regulation via Independent Agency, 54 Colum. J. L. & Soc. Probs. 261 (2021).

[34] *Griggs v. Duke Power Co.*, 401 U.S. 424 (1971).

[35] *Questions and Answers to Clarify and Provide a Common Interpretation of the Uniform Guidelines on Employee Selection Procedures*, 44 Fed. Reg. (March 2, 1979).

[36] See *Albemarle Paper Co. v. Moody*, 422 U.S. 405, 425, 95 S. Ct. 2362, 2375 (1975).

[37] Civil Rights Act of 1964 § 7, 42 U.S. Code § 2000e.

[38] The Age Discrimination in Employment Act of 1967 § 623, 29 U.S.C. §§ 621–634.

[39] Civil Rights Act of 1964 § 2, 42 U.S. Code § 2000e.

facilities, privileges, advantages, or accommodations of any place of public accommodation."[40] As consumers have successfully sued TNCs for failing to meet the ADA Title III accessibility requirements,[41] it is plausible that the ADA does furnish private parties with a suitable disparate impact litigation tool to challenge algorithmic discrimination in TNCs. Similarly, while there has not been a Title II disparate impact case against TNCs, courts have at least assumed *arguendo* that Title II does allow for disparate impact claims.[42]

Recently, a U.S. District Court in Western Texas also held that a disparate impact claim against algorithmic discrimination could be cognizable. In *DeHoyos v. Allstate Corp.*, the court approved a settlement agreement in a class action lawsuit alleging that "the insurers used a credit scoring algorithm as a form of unfair racial discrimination which caused them to be charged higher rates for automobile and homeowners' insurance policies than similarly situated Caucasians."[43] These precedents thus support the argument that passengers could bring suit against algorithmic biases in TNC platforms under the disparate impact framework.

III. RECOMMENDATIONS FOR ALGORITHMIC FAIRNESS

In addition to legal algorithmic governance, I propose that the engineering community borrow a page from the American Bar Association (ABA) and self-regulate by creating and enforcing some ethical standards for software development. Specifically, I suggest having engineering consortiums like the Institute of Electrical and Electronics Engineers (IEEE) address algorithmic bias by supplementing their current proposed standards for combatting said bias with proposals to improve algorithmic fairness and transparency in the subprocesses of the agile software development process – project planning, requirement analysis, technical specifications, implementation of the tech-specs, quality assurance, and acceptance testing.[44]

In the project planning and requirement analysis phase, for example, product managers usually meet with various stakeholders to gather information and solicit feedback for a proposed feature.

---

[40] Americans with Disabilities Act of 1990 § 3, 42 U.S. Code § 12182.
[41] *See Lowell v. Lyft, Inc.*, 352 F. Supp. 3d 248, 251 (S.D.N.Y. 2018).
[42] *See, e.g.*, *Robinson v. Power Pizza, Inc.*, 993 F. Supp. 1462, 1464-65 (M.D. Florida 1997); *Arguello v. Conoco, Inc.*, 207 F.3d 803, 813 (5th Cir. 2000).
[43] *DeHoyos v. Allstate Corp.*, 240 F.R.D. 269, 275 (W.D. Tex. 2007).
[44] *See* Ismail Abiodun Sulaimon, Ahmed Ghoneim, Mubarak Alrashoud,, *A New Reinforcement Learning-Based Framework for Unbiased Autonomous Software Systems*, 8th International Conference on Modeling Simulation and Applied Optimization (ICMSAO), 1 (2019) (discussing current IEEE efforts to come up with and promulgate ethical standards to combat algorithmic bias).

The resulting product specifications are, consequently, heavily dependent on the demographics composition of the participants. The end users of these algorithms, however, rarely have a seat at the table. To remediate this problem, the IEEE standards could broaden the composition of relevant stakeholders who could participate in both the project planning and requirement analysis phase. This would both reduce bias and increase fairness and transparency in the software development process.

IV. CONCLUSION

This paper has explored how algorithmic bias could affect TNC software and how consumers could potentially utilize the disparate impact framework to protect their rights. It has also recommended some private STEM standards to reduce algorithmic bias and improve algorithmic transparency. Although some of these recommendations are specific to TNCs, they could be expanded to cover any entities that use AI/ML software in their decision-making process. Irrespective of their feasibility, however, it is important to continue exploring new forms of algorithmic governance, especially since the algorithmic turn is here to stay.