

Bias in AI-Based Decision-Making



Adheesh Kadiresan, Yuvraj Baweja, and Obi Ogbanufe

Introduction

For much of human history, human beings have been the decision-makers on matters pertaining to humans (Anyoha, 2017). People made decisions in areas such as hiring, loan eligibility, diagnosis of diseases, retail, manufacturing, entertainment, and more (Colson, 2019). However, in recent decades, artificial intelligence (AI) has been able to perform certain tasks more skillfully and reliably than humans could. For example, 1997 marked the date of the defeat of Gary Kasparov, the highest-ranked chess player, by Deep Blue, a computer chess program created by IBM (Anyoha, 2017). Artificial intelligence is now being used to make decisions in areas such as hiring, loan eligibility, housing, medicine (DeGonia et al., 2016), “technology, banking, marketing, and entertainment” (Anyoha, 2017, para. 9). The implementation of AI in these sectors is due to the ability of AI to perform certain tasks more accurately than humans. For example, in the case of medical sciences, AI was able to decrease false positives of breast cancer by 5.7% on a US data set and by 1.2% on a UK data set (McKinney et al., 2020). The potential benefits AI can bring to society are clear.

When humans were the sole decision-makers before the age of AI, biased decision-making was rampant. One of the most historically significant instances of this is the practice of redlining by the Federal Housing Administration. After the Great Depression, in the 1930s, the Home Owners’ Loan Corporation (HOLC) created maps that were intended to stabilize property values and determine the credit-worthiness of entire neighborhoods. However, these maps were in part influenced by the races of the residents of each neighborhood (Aaronson et al., 2017). This caused discrimination of neighborhoods based on race, which denied investment

A. Kadiresan (✉) · Y. Baweja · O. Ogbanufe
University of North Texas, Denton, TX, USA
e-mail: Obi.Ogbanufe@unt.edu

© The Author(s), under exclusive license to Springer Nature
Switzerland AG 2022

M. V. Albert et al. (eds.), *Bridging Human Intelligence and Artificial Intelligence*, Educational Communications and Technology: Issues and Innovations, https://doi.org/10.1007/978-3-030-84729-6_19

275

opportunities in communities with higher populations of African Americans. This discriminatory practice, termed redlining due to the red color used on maps (Aaronson et al., 2017), created effects that lasted a much longer duration. Redlined areas were “associated with a 5% decrease in 1990 house prices” (Appel & Nickerson, 2016, p. 24). Clearly, biased decision-making can have severe long-lasting effects on society and can contribute to discrimination.

However, using computers to make decisions does not automatically eliminate or even reduce bias. In this chapter, we will explore ways in which decision-making by computers can introduce and exacerbate certain biases. Specifically, this chapter will explore biases that affect human lives. These biases can be prejudiced against “race, gender, age, or any other trait” (DeGonia et al., 2016, p. 16).


Defining Bias

Bias is defined as “an inclination of temperament or outlook” (Merriam-Webster, n.d.-a, para. 1). Thus, favoring one entity more than another would be an instance of bias. In the real world, bias closely relates to the idea of discrimination, which is defined to be a “prejudiced or prejudicial outlook, action, or treatment” (Merriam-Webster, n.d.-b, para. 1). It can be seen that discrimination is a product of bias: a skewed outlook can lead to prejudicial actions toward others. This makes clear the reason why bias is a problem: bias causes discrimination against various groups of people, which has had severe historical consequences, such as in the case of redlining provided above. In this chapter, we explore bias in decisions carried out by artificial intelligence systems.

The Formation of Bias

AI decisions can become biased in numerous ways. One way a decision can become biased is when AI models confuse correlation with causation (DeGonia et al., 2016). Correlation simply refers to an instance when two variables change together. However, causation refers to a relationship between a causing factor and an affected factor (DeGonia et al., 2016). Confusing causation and correlation involve assuming one factor is causing another when in fact the two factors only happen to change in a noncausal relationship. The canonical example here is a potential correlation between ice cream sales and violence. Although higher frequencies of ice cream sales may be correlated with higher violence rates which may occur in warmer months, concluding that ice cream causes violence would be an obvious fallacy of confusing correlation with causation. The true explanation may be that warmer weather is correlated with both higher ice cream sales and also more acts of violence (DeGonia et al., 2016). Figure 1 shows an example of when two correlated factors do not necessarily lead to causation.

Fig. 1 Two correlated factors not necessarily in causation



One situation of mistaking correlation for causation more relevant to artificial intelligence decision-making comes from the usage of zip codes to determine employment. For instance, if a computer model finds that a certain zip code is correlated with the locations of better-performing employees, an incorrect and biased causal relationship may be assumed: that a zip code causes employees to be better (DeGonia et al., 2016). Such an assumption may lead to zip codes being used to determine the employability of hires. The problem here has to do with the historical issue of discrimination in housing, such as the aforementioned practice of redlining. Since redlining disproportionately affected African-American communities (Aaronson et al., 2017), the incorrectly assumed causal relation may end up contributing to racial discrimination in hiring (DeGonia et al., 2016).

Bias can also cause discrimination against certain groups when irrelevant factors are taken into account by a decision-making algorithm. While the addition of

certain parameters relevant to the problem at hand may improve accuracy, having irrelevant parameters can harm accuracy and strengthen “racism, sexism, and other inequalities.” (DeGonia et al., 2016, p. 47). Such irrelevant factors include those that do not have any effect on the end goal. For example, factors that will be referred to as *personal factors*, such as nationality and ethnicity, have no effect on the skills of that person. Providing such information to an algorithm will not improve the accuracy of the algorithm, and the possibility of bias or discrimination means that such factors are harmful in AI decision-making.

A decision-making AI model may find correlations that happen to form across lines of different groups of people, which could exacerbate discrimination of these groups (DeGonia et al., 2016). Thus, it follows that removing personal factors from AI models can decrease discrimination against members of certain groups.

A related source of bias in computerized decision-making comes from skewed data sets that do not accurately represent the target population of an algorithm. As a result, algorithms can make more errors on under-represented demographics in the data set. One example comes from facial recognition algorithms. A groundbreaking paper by Buolamwini and Gebru (2018) describes how all of the observed facial recognition classifiers had error rates from 20.8% (for Microsoft’s classifier) to 34.7% (IBM) for “darker-skinned females” (pp. 9–10). On the other hand, the study found that “the maximum error rate for lighter-skinned males is 0.8%” (Buolamwini & Gebru, 2018, p. 1). This large disparity shows that facial recognition classifiers are biased against certain groups. The paper also analyzes the IARPA (Intelligence Advanced Research Projects Activity) Janus Benchmark A data set (IJB-A data set) which is designed to be “geographically diverse” as well as Adience, which is a “gender and age classification benchmark” (Buolamwini & Gebru, 2018, p. 3). The paper notes that the IJB-A data set consists of 79.6% “lighter-skinned individuals” and Adience consists of 86.2% “lighter subjects” (Buolamwini & Gebru, 2018, p. 7). In this case, we can see that these data sets tend to under-represent people with darker skin.

Another study focuses on how facial recognition algorithms from East Asia tended to perform better on Asian subjects than did algorithms developed in the Western hemisphere (Klare et al., 2012). Similarly, for white subjects, algorithms developed in the western hemisphere performed better. The paper continues suggesting that “this discrepancy was due to the different racial distribution in the training sets for the Western and Asian algorithms” (Klare et al., 2012, p. 3). These examples show how important the samples of data used to develop an algorithm are in providing unbiased results. Data sets must represent all demographics more equally in order to reduce bias in classification.

Biases in AI can also stem from biased humans that contribute to a computer algorithm. One example of humans contributing to bias is in the labor market. One study found that resumes with “white names receive 50 percent more callbacks for interviews” than those with “African-American names” (Bertrand & Mullainathan, 2003, p. 2). This experiment made it clear that discrimination from humans exists in the labor market. Since artificial intelligence, computer algorithms, and data collecting all have a human component, it is easy to see how such biases in people can become manifest in computer and AI algorithms and then lead to discrimination.

This kind of bias is not due to AI itself, since the bias originates from humans. In other words, for human-introduced bias, removing artificial intelligence from the task at hand would not necessarily eliminate or even reduce bias.

To conclude, a biased algorithm may derive its bias from an incorrect assumption of causation from correlation, personal factors that are irrelevant to the algorithm's decision, skewed or incomplete data sets that leave out certain demographics, or the humans that created the algorithm in the first place. All of these factors can lead to bias, which has the potential to exacerbate existing discrimination.

Real-World Example of Bias

One example of bias from an AI algorithm comes from [Amazon.com](#), Inc.'s 2014 AI-hiring program. Amazon had to cancel this program after the discovery of biases against women in the hiring algorithm. In this case, the discrimination stemmed from the training data, which included ten years of submitted resumes to Amazon. However, due to the long-running "male dominance across the tech industry" (para. 6) from the gender hiring gap in technology, the algorithm learned to favor resumes from men and "penalized resumes that included the word 'women's'" (Dastin, 2018, para. 7). The source of bias in this case is the data set for the algorithm. However, in this case, the main issue is not with the data set misrepresenting the target audience; instead, the data set reflected the reality of male-skewed hiring. When the algorithm "was trained on historical hiring decisions, which favored men over women, it learned to do the same" (Hao, 2019, para. 5).

Beyond the obvious issue of a biased, and thus ineffective hiring algorithm, the example above is significant for worsening the long history of gender discrimination. From this example, we can see how AI decision-making can both harm potential workers and also contribute to the larger issues of discrimination within society.

In addition to the ethical issues of unfairly rejecting applicants based on gender and contributing to gender discrimination, the example of bias in hiring also faces legal issues. The Federal Equal Opportunity Laws "[prohibit] employment discrimination based on race, color, religion, sex, or national origin" ("Federal Laws Prohibiting Job Discrimination Questions and Answers," n.d., para. 1). Therefore, discriminatory hiring would not be legal, in addition to being unethical. The legal issues will be covered in more detail later in this chapter.

Trust in Artificial Intelligence

AI has the potential to make decisions that can benefit the areas of science, wellbeing, economics, and solutions to environmental issues (Rossi, 2019). The example of breast cancer above indicates that AI can help make decisions that are more accurate and help solve important problems in society.

However, before AI can be widely deployed to solve problems, people must trust it to carry out accurate decisions. Specifically, artificial intelligence must “be aware and aligned to human values” and “explain its reasoning and decision-making” (Rossi, 2019, para. 5). Stefan Jockusch from Siemens presents trust as “justified by statistics” (para. 26) and that trust in facial recognition algorithms, which utilize AI, led to the use of those algorithms in the important task of “recognizing identity” (MIT Technology Review Insights, 2020, para. 25).

The relevant area in which trust must be established in AI is in the avoidance of discrimination. One field where eliminating discrimination is especially important is hiring. While AI has the potential to quickly sift through job applications (Fatemi, 2019), these hiring decisions may be biased against certain groups. Allowing a biased hiring process can erode trust in AI hiring. This is substantiated by the fact that 35% of the US adults that would apply to a position using AI hiring would do so due to their trust that AI can be “fairer, less biased than humans” (Smith, 2017, para. 14). Thus, it is reasonable to infer that if the fairness of hiring with AI were to be compromised, public trust in the abilities of computers to carry out employment decisions would decrease, resulting in reduced usage of a technology that would have had significant advantages.

Efforts to Prevent AI Bias

Sources of Bias

We have learned that AI decisions can have serious biases that contribute to discrimination in society. Specifically, such biases can come from mistaking causation and correlation, relying on factors irrelevant to an algorithm’s decision, skewed data sets, or human contributions. How can such biases in computerized decision-making be resolved? In order to eliminate bias from AI algorithms, all of the above issues must be addressed.

The first two issues are closely related: assuming causation from two variables that happen to be correlated contributes to bias if the variables taken in have the potential to discriminate. Such variables are the “personal factors” that must not affect the decision-making of the algorithm. However, eliminating such discriminatory variables is not as easy as it may appear. In fact, data provided to an algorithm can still “include biased human decisions or reflect historical or social inequities, even if sensitive variables such as gender, race, or sexual orientation are removed” (Manyika et al., 2019, para. 4). For example, in the case of Amazon’s gender-biased hiring algorithm, words such as “executed” and “captured” were used to discriminate against women, since resumes from men tended to contain these words more often (Dastin, 2018). Thus, removing explicit personal factors from algorithms is not adequate to prevent discrimination. These factors can manifest themselves in other aspects of the training data.

Another source of bias comes from skewed data sets. Data sets could be skewed due to either real-world biases or data that do not fully represent certain demographics. For example, in the case of Amazon's biased algorithm, the data set was skewed due to real-world inequalities in hiring between women and men (Dastin, 2018). In order to reduce bias in this case, the AI algorithm must make decisions that do not follow the previous patterns in hiring. When data sets are biased from an incomplete representation of all groups of people, data sets must be improved. There are multiple views on this issue. Google's AI department says "Public training data sets will often need to be augmented to better reflect real-world frequencies of people" ("Responsible AI Practices," n.d., para. 9). This view emphasizes how data sets themselves can be biased and need to be altered and improved in order to help reduce bias. Buolamwini and Gebru (2018) created the Pilot Parliaments Benchmark data set that is "gender and skin type balanced" (p. 1) from "male and female parliamentarians from 6 countries" (p. 4). This data set was found to represent "darker female, darker male, lighter female and lighter male subjects" in a more balanced manner than other data sets (Buolamwini & Gebru, 2018, p. 7). IBM AI, on the other hand, through a blog, claims that "machine learning, by its very nature, is always a form of statistical discrimination" and that becomes an issue when "privileged groups [are given a] systematic advantage" and "unprivileged groups [are given a] systematic disadvantage" (Varshney, 2018, para. 2). This point of view emphasizes how machine learning itself aims to discriminate and effort must be applied to prevent discrimination that unjustly harms certain demographics.

The final source of AI bias comes from humans contributing to the field of AI. There are multiple ways to combat this issue as well. For example, the Harvard Business Review recommends "diversifying the AI field itself ... to anticipate, review, and spot bias and engage communities affected" (Manyika et al., 2019, para. 16). However, Joann Stonier, who is the chief data officer of Mastercard, emphasizes "governance and testing methodologies" to combat bias among data scientists (Stonier, 2020, para. 8).

The Issue of Gauging Bias

To approach the issue of bias, it is important to have a method to measure the amount of fairness in an algorithm. Two such fairness measures are group and individual fairness. Group fairness aims for "statistical parity ... for members of different protected groups" whereas individual fairness aims to assign "similar outcomes" to "people who are 'similar' with respect to the classification task" (Binns, 2020, p. 1). Figure 2 provides a simplified hypothetical scenario where these two metrics of fairness yield different classifications. We can see that in group fairness, each group has the same proportion of its members in each outcome. In other words, members of each group have equal probabilities of reaching outcome 1 or outcome 2. Similar qualifications between members of different groups do not necessarily result in similar outcomes. On the other hand, individual fairness means each group does not

SITUATION:

Group A	Group B	
500	1000	Number of members
400	300	Members qualified for outcome 1
100	700	Members qualified for outcome 2

GROUP FAIRNESS:

Group A	Group B	
300	600	Outcome 1
200	400	Outcome 2

INDIVIDUAL FAIRNESS:

Group A	Group B	
400	300	Outcome 1
100	700	Outcome 2

Fig. 2 Group fairness compared to individual fairness

have the same proportion of its members in each outcome. Here, similar qualifications result in similar outcomes regardless of a member’s group.

AI Bias and the Law

We have explored the formation of biases in AI and the search to mitigate such biases. However, what are the consequences of biased algorithms? Specifically, what laws surround bias in general, what laws specifically target biases in AI right now, and what direction could these laws go in the future?

To make the discussion more focused, we will focus on employment discrimination. This is because the process of hiring can have significant bias, as previously discussed. Furthermore, AI is widely used in hiring. In fact, LinkedIn reported that in 2018, 67% of surveyed recruiters reported that they save time by using AI technology (“*LinkedIn 2018 Report Highlights Top Global Trends in Recruiting*,” 2018).

In order to compare the extent to which AI biases are recognized and acted on by the law, we must first analyze the overall biases in job recruiting and laws surrounding these biases. One study establishes the extent to which hiring can be biased by revealing the discrimination toward those with “African-American names” compared to those with “White names” (Bertrand & Mullainathan, 2003, p. 2). In this case, the Equal Employment Opportunity Commission (EEOC) enforces the Civil Rights Act’s Title VII, “which makes it illegal to discriminate against a person on the basis of race, color, religion, sex, or national origin” (*What Laws Does EEOC Enforce?*, n.d., para. 2).

Due to the EEOC laws, discriminating against these categories would not be legal. However, from the example of Amazon's biased algorithm, it is evident that algorithms can learn to discriminate between groups, such as gender, through separate, seemingly unrelated words. Thus, in order to comply with antidiscrimination laws, AI decision-making must not only prevent bias from explicitly labeled demographics, but also from demographic information that can be inferred through other means.

Another issue here is the possibility of a more indirect but still harmful form of discrimination. For example, if an algorithm makes decisions based on the location of each employee, a zip code may come to determine employment for citizens of certain places (DeGonia et al., 2016). Then, due to long-existing racial discrimination in housing (Aaronson et al., 2017), zip-code-based decision-making could lead to racial discrimination (DeGonia et al., 2016). One example of laws being enforced to fight this sort of discrimination from location comes from a lawsuit against Abercrombie & Fitch. After the discriminatory hiring practices were revealed, Abercrombie & Fitch was "barred from utilizing its previous recruitment strategies, such as targeting particular predominately white fraternities or sororities" (*Case: Abercrombie & Fitch Employment Discrimination*, 2006, para. 5). The action taken against Abercrombie & Fitch here could serve as a model for dealing with more indirect but still very harmful forms of AI discrimination.

It is clear that the laws surrounding AI have holes as discrimination can still surface through unexpected factors. Therefore, one can conclude that progress must be made in law to prevent discrimination through any factors fed into a hiring algorithm. One example of a recent house bill is the "Algorithmic Accountability Act of 2019." In this case, "covered entities" (p. 4) must "conduct automated decision system impact assessments of ... high-risk automated decision systems" including analysis of the algorithm's purpose, benefits, risks, privacy, and risk minimization procedures (Algorithmic Accountability Act of 2019, 2019, p. 9). This is an example of a law that may ensure more fair decision-making algorithms.

Conclusions

While AI algorithms can increase efficiency in hiring, entertainment, and other industries, these algorithms can also contribute to bias and discrimination. This bias can come from mistaking the correlation of two variables for causation, depending on discriminatory personal factors, using skewed or incomplete data sets, or human sources. The effort to combat AI bias is an ongoing one, and there are no simple solutions in this area. AI bias can also encounter issues with the law, as bias can be introduced in subtle ways that still discriminate against certain demographics. As society starts to apply AI to a wider variety of decisions, the issue of bias must not be overlooked.

References

- Aaronson, D., Hartley, D., & Mazumder, B. (2017). The effects of the 1930s HOLC “redlining” maps (No. 2017–12). *EconStor*. Working Paper. <https://www.econstor.eu/handle/10419/200568>
- Algorithmic Accountability Act of 2019, H. R. 2231, House of Representatives, 116th Congress (2019). <https://www.congress.gov/bill/116th-congress/house-bill/2231/text>
- Anyoha, R.. (2017, August 28). *The history of artificial intelligence*. <http://sitn.hms.harvard.edu/flash/2017/history-artificial-intelligence/>
- Appel, I., & Nickerson, J. (2016). *Pockets of Poverty: The Long-Term Effects of Redlining*. SSRN. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2852856
- Bertrand, M., & Mullainathan, S. (2003). Are Emily and Greg More Employable than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination (No. w9873). *National Bureau of Economic Research*. <https://doi.org/10.3386/w9873>
- Binns, R. (2020). On the apparent conflict between individual and group fairness. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (pp. 514–524).
- Buolamwini, J., & Gebru, T. (2018). Gender shades: Intersectional accuracy disparities in commercial gender classification. In S. A. Friedler & C. Wilson (Eds.), *Proceedings of the 1st Conference on Fairness, Accountability and Transparency* (Vol. 81, pp. 77–91). PMLR.
- Case: Abercrombie & Fitch employment discrimination. (2006, March 17). <https://www.naacpldf.org/case-issue/abercrombie-fitch-employment-discrimination/>
- Colson, E. (2019, July 8). *What AI-driven decision making looks like*. Harvard Business Review. <https://hbr.org/2019/07/what-ai-driven-decision-making-looks-like>
- Dastin, J. (2018, October 10). Amazon scraps secret AI recruiting tool that showed bias against women. *Reuters*. <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight-idUSKCN1MK08G>
- DeGonia, A., Kerper, A., Baxter, A., Miller, A., Shah, A., Kik, B., Kowalczyk, B., Williams, C., Grapperhaus, G., Bosnich, G. A., Goscenski, H., Buss, J., Hwang, J., Ng, J., Adamski, J., Chokshi, K., Radulovic, K., Kalesinskas, L., Scarlati, L., ... Kegel, J. (2016). *Automated discrimination: The power and peril of big data: Big data revolutionizes the power of prediction, but can also perpetuate discrimination far beyond the constraints of human ability*. CreateSpace Independent Publishing Platform.
- Fatemi, F. (2019, October 31). How AI is uprooting recruiting. *Forbes Magazine*. <https://www.forbes.com/sites/falonfatemi/2019/10/31/how-ai-is-uprooting-recruiting/>
- Federal Laws prohibiting job discrimination questions and answers*. (n.d.). Retrieved December 19, 2020, from <https://www.eeoc.gov/fact-sheet/federal-laws-prohibiting-job-discrimination-questions-and-answers>
- Hao, K. (2019, February 4). This is how AI bias really happens—And why it’s so hard to fix. *Technology Review*. Retrieved December 19, 2020, from <https://www.technologyreview.com/2019/02/04/137602/this-is-how-ai-bias-really-happensand-why-its-so-hard-to-fix/>
- Klare, B. F., Burge, M. J., Klontz, J. C., Vorder Bruegge, R. W., & Jain, A. K. (2012). Face recognition performance: Role of demographic information. *IEEE Transactions on Information Forensics and Security*, 7(6), 1789–1801.
- LinkedIn 2018 Report Highlights Top Global Trends in Recruiting*. (2018, January 10). Retrieved December 19, 2020, from <https://news.linkedin.com/2018/1/global-recruiting-trends-2018>
- Manyika, J., Silberg, J., & Presten, B. (2019, October 25). What do we do about the biases in AI? *Harvard Business Review*. <https://hbr.org/2019/10/what-do-we-do-about-the-biases-in-ai>
- McKinney, S. M., Sieniek, M., Godbole, V., Godwin, J., Antropova, N., Ashrafian, H., Back, T., Chesus, M., Corrado, G. S., Darzi, A., Etemadi, M., Garcia-Vicente, F., Gilbert, F. J., Halling-Brown, M., Hassabis, D., Jansen, S., Karthikesalingam, A., Kelly, C. J., King, D., ... Shetty, S. (2020). International evaluation of an AI system for breast cancer screening. *Nature*, 577(7788), 89–94.
- Merriam-Webster. (n.d.-a). Bias. In *Merriam-Webster.com dictionary*. Retrieved December 19, 2020, from <https://www.merriam-webster.com/dictionary/bias>

- Merriam-Webster. (n.d.-b). Discrimination. In *Merriam-Webster.com dictionary*. Retrieved December 19, 2020, from <https://www.merriam-webster.com/dictionary/discrimination>
- MIT Technology Review Insights. (2020, October 28). With trust in ai, manufacturers can build better. *Technology Review*. Retrieved December 19, 2020, from <https://www.technologyreview.com/2020/10/28/1011266/with-trust-in-ai-manufacturers-can-build-better/>
- Responsible AI Practices*. (n.d.). Retrieved December 18, 2020, from <https://ai.google/responsibilities/responsible-ai-practices/?category=fairness>
- Rossi, F. (2019, February 6). *Building trust in artificial intelligence*. <https://jia.sipa.columbia.edu/building-trust-artificial-intelligence>
- Smith, A. (2017, October 4). *Americans' views toward hiring algorithms*. <https://www.pewresearch.org/internet/2017/10/04/americans-attitudes-toward-hiring-algorithms/>
- Stonier, J. (2020, March 19). Fighting AI bias—Digital rights are human rights. *Forbes Magazine*. <https://www.forbes.com/sites/insights-ibmai/2020/03/19/fighting-ai-bias-digital-rights-are-human-rights/>
- Varshney, K. R. (2018, September 19). *Introducing AI Fairness 360*. <https://www.ibm.com/blogs/research/2018/09/ai-fairness-360/>
- What laws does EEOC enforce?* (n.d.). Retrieved December 19, 2020, from <https://www.eeoc.gov/youth/what-laws-does-eeoc-enforce>

Adheesh Kadiresan is a student at the Texas Academy of Mathematics and Science (TAMS), a residential high-school program run by the University of North Texas (UNT). He is a research assistant at Dr. Mark Albert's Biomedical-AI Lab studying the use of Autoencoders in an EKG Vest. He was also an early summer researcher under Dr. Nicoladie Tam, studying the effect of emotion on decision-making. Under the guidance of Dr. Tam, Adheesh created his poster titled "Effect of Emotion on Decision-Making," which was presented in the 2021 UNT Scholars Day. Additionally, Adheesh is a recipient of the 2020 TAMS Summer Research Scholarship and the 2021 TAMS Summer Research Scholarship. He will be working in the UNT AI Summer Research Program during the summer of 2021.

Yuvraj Baweja is an incoming freshman in the University of Texas at Austin's Computer Science and Business Honors program after completing two years at the Texas Academy of Mathematics & Science (TAMS), an accelerated high school program at the University of North Texas (UNT Denton). He is working on a project in the AI Laboratory at UNT run by Mark V. Albert dubbed Stock2Vec, using Word2Vec embeddings to run predictions on companies given financial data. His research interests include the intersection between Artificial Intelligence and Business, specifically Natural Language Processing, Machine Learning, and Economics. Aside from research, he has given two TEDx talks, one concerning computer science in the core school curriculum and one about the benefits of residential programs. Furthermore, he has competed in numerous computer science and cyber-security competitions, achieving rankings such as the top 25 out of 5000 high school teams in picoCTF.

Dr. Obi Ogbanufe is an Assistant Professor of Information Technology and Decision Sciences at the University of North Texas. She is a recipient of the NSF CyberCorps Scholarship for Service award. She has published in Decision Support Systems, Information Systems Journal, and the International Journal of Human-Computer Interaction. Her research interests include information security, cybercrime, human-AI interaction, and risk management.