The Effects^{of} Mixing & Machine Learning Human Judgment

**COLLABORATION
DOES NOT
NECESSARILY
LEAD TO BETTER
OUTCOMES.**

MICHELLE VACCARO AND JIM WALDO

In 1997 IBM's Deep Blue software beat the World Chess Champion Garry Kasparov in a series of six matches. Since then, other programs have beaten human players in games ranging from Jeopardy to Go. Inspired by his loss, Kasparov decided in 2005 to test the success of Human+AI pairs in an online chess tournament.² He found that the Human+AI team bested the solo human. More surprisingly, he also found that the Human+AI team bested the solo computer, even though the machine outperformed humans.

Researchers explain this phenomenon by emphasizing

that humans and machines excel in different dimensions of intelligence.⁹ Human chess players do well with long-term chess strategies, but they perform poorly at assessing the millions of possible configurations of pieces. The opposite holds for machines. Because of these differences, combining human and machine intelligence produces better outcomes than when each works separately. People also view this form of collaboration between humans and machines as a possible way to mitigate the problems of bias in machine learning, a problem that has taken center stage in recent months.¹²

We decided to investigate this type of collaboration between humans and machines using risk-assessment algorithms as a case study. In particular, we looked at the COMPAS (Correctional Offender Management Profiling for Alternative Sanctions) algorithm, a well-known (perhaps infamous) risk-prediction system, and its effect on human decisions about risk. Many state courts use algorithms such as COMPAS to predict defendants' risk of recidivism, and these results inform bail, sentencing, and parole decisions.

Prior work on risk-assessment algorithms has focused on their accuracy and fairness, but it has not addressed their interactions with human decision makers who serve as the final arbitrators. In one study from 2018, Julia Dressel and Hany Farid compared risk assessments from the COMPAS software and Amazon Mechanical Turk workers, and found that the algorithm and the humans achieved similar levels of accuracy and fairness.⁶ This study signals an important shift in the literature on risk-assessment instruments by incorporating human

subjects to contextualize the accuracy and fairness of the algorithms. Dressel and Farid's study, however, divorces the human decision makers and the algorithm when, in fact, the current model indicates that humans and algorithms would work in tandem.

Our work, consisting of two experiments, therefore first explores the influence of algorithmic risk assessments on human decision-making and finds that providing the algorithm's predictions does not significantly affect human assessments of recidivism. The follow-up experiment, however, demonstrates that algorithmic risk scores act as anchors that induce a cognitive bias: If we change the risk prediction made by the algorithm, participants assimilate their predictions to the algorithm's score.

The results thus highlight potential shortcomings with the existing human-in-the-loop frameworks. On the one hand, when algorithms and humans make sufficiently similar decisions their collaboration does not achieve improved outcomes. On the other hand, when algorithms fail, humans may not be able to compensate for their errors. Even if algorithms do not officially make decisions, they anchor human decisions in serious ways.

EXPERIMENT ONE: HUMAN-ALGORITHM SIMILARITY, NOT COMPLEMENTARITY

The first experiment examines the impact of the COMPAS algorithm on human judgments concerning the risk of recidivism. COMPAS risk scores were used because of the data available on that system, its widespread usage in prior work about algorithmic fairness, and the use of the system in numerous states.

Methods

The experiment entailed a 1 x 3 between-subjects design with the following treatments: *control*, in which participants see only the defendant profiles; *score*, in which participants see the defendant profiles and the defendant COMPAS scores; and *disclaimer*, in which participants see the defendant profiles, the defendant COMPAS scores, and a written advisement about the COMPAS algorithm.

Participants evaluated a sequence of defendant profiles that included data on gender, race, age, criminal charge, and criminal history. These profiles described real people arrested in Broward County, Florida, based on information from the dataset that *ProPublica* used in its analysis of risk-assessment algorithms.¹ While this dataset originally contained 7,214 entries, this study applied the following filters before sampling for 40 profiles that were presented to participants:

- Limit to black and white defendants. Prior work on the accuracy and fairness of the COMPAS algorithm limits their analyses to white and black defendants.^{3,4,6} To compare the results from this experiment with those in prior studies, this study considers only the subset of defendants who identify as either African-American (black) or Caucasian (white).
- Exclude cannabis crimes. Interestingly, the pilot study showed participant confusion about cannabis-related crimes such as possession, purchase, and delivery. In the free-response section of the survey, participants made comments such as “Cannabis is fully legal here.” To avoid confusion about the legality of cannabis in various states, this study excludes defendants charged with crimes

containing the term *cannabis*.

From this filtered dataset 40 defendants were randomly sampled. A profile was generated containing information about the demographics, alleged crime, criminal history, and algorithmic risk assessment for each of the defendants in the sample. The descriptive paragraph in the control treatment assumed the following format, which built upon that used in Dressel and Farid's study:⁶

The defendant is a [RACE] [SEX] aged [AGE]. They have been charged with: [CRIME CHARGE]. This crime is classified as a [CRIMINAL DEGREE]. They have been convicted of [NON-JUVENILE PRIOR COUNT] prior crimes. They have [JUVENILE-FELONY COUNT] juvenile felony charges and [JUVENILE-MISDEMEANOR COUNT] juvenile misdemeanor charges on their record.

The descriptive paragraph in the score treatment added the following information:

COMPAS is risk-assessment software that uses machine learning to predict whether a defendant will commit a crime within the next two years. The COMPAS risk score for this defendant is [SCORE NUMBER]: [SCORE LEVEL].

Finally, the descriptive paragraph in the disclaimer treatment provided the following information below the COMPAS score, which mirrored the language the Wisconsin Supreme Court recommended in *State v Loomis*:¹⁸

Some studies of COMPAS risk-assessment scores have raised

questions about whether they disproportionately classify minority offenders as having a higher risk of recidivism.

Upon seeing each profile, participants were asked to provide their own risk-assessment scores for the defendant and indicate if they believed the defendant would commit another crime within two years. Using dropdown menus, they answered the questions shown in figure 1.

We deployed the task remotely through the Qualtrics platform and recruited 225 respondents through Amazon Mechanical Turk, 75 for each treatment group. All workers could view the task title, “Predicting Crime”; task description, “Answer a survey about predicting crime”; and the keywords associated with the task, “survey, research, and criminal justice.” Only workers living in the United States could

FIGURE 1: **DEFENDANT PROFILE FROM SCORE TREATMENT**

The defendant is a Caucasian male aged 47. They have been charged with: Grand Theft in the 3rd Degree. This crime is classified as a felony. They have been convicted of 1 prior crime. They have 0 juvenile felony charges and 0 juvenile misdemeanor charges on their record.

COMPAS is risk assessment software that uses machine learning to predict whether a defendant will commit a crime within the next two years. The COMPAS risk score for this defendant is 4: Low-Risk.

Please indicate the likelihood that this person will commit another crime within two years by assigning them a recidivism risk score: 1 to 4 (low risk), 5 to 7 (medium risk), 8 to 10 (high risk).

Do you think this person will commit another crime within two years?

complete the task, and they could do so only once. During the pilot study among an initial test group of five individuals, the survey required an average of 15 minutes to complete. As the length and content of the survey resembled that of Dressel and Farid's,⁶ we adopted their payment scheme, giving workers \$1 for completing the task and a \$2 bonus if the overall accuracy of the respondent's predictions exceeded 65 percent. This payment structure motivated participants to pay close attention and provide their best responses throughout the task.^{6,17}

Results

Figure 2 shows the average accuracy of participants in the control, score, and disclaimer treatments. The error bars represent the 95 percent confidence intervals. The results suggest that the provision of COMPAS scores did not significantly affect the overall accuracy of human predictions of recidivism. In this experiment, the overall accuracy of predictions in the control treatment (54.2 percent) did not significantly vary from those in the score treatment (51.0 percent) [$p = 0.1460$].

The inclusion of a written advisement about the limitations of the COMPAS algorithm did not significantly affect the accuracy of human predictions of recidivism, either. Participants in the disclaimer treatment achieved an average overall accuracy rate of 53.5 percent, whereas those in the score condition achieved 51.0 percent; a two-sided t-test indicated that this difference was not statistically significant [$p = 0.1492$].

Upon the conclusion of the task block in the exit survey, 99 percent of participants responded that they found the

FIGURE 2: **ACCURACY RATE IN TREATMENT GROUPS**

instructions for the task clear, and 99 percent found the task satisfying. In their feedback, participants indicated they had positive experiences with the study, leaving comments such as: “I thoroughly enjoyed this task”; “It was a good length and good payment”; and “Very good task.”

Participants did not mention the advisement when asked how they took the COMPAS scores into account. Rather, their responses demonstrated that they used the COMPAS scores in different ways: some ignored them, some relied heavily on them, some used them as starting points, and others used them as sources of validation.

Figure 3 has excerpts of participant responses with a summary of answers to the free-response question: “How did you incorporate the COMPAS risk scores into your decisions?”

FIGURE 3: PARTICIPANT RESPONSES TO FREE-RESPONSE QUESTION

	COMPAS	DISCLAIMER
Ignore	“I tried <i>not</i> to look at them after awhile, because I felt some were off (lol) but I still took them into account somewhat. I mostly went with my gut and opinions, though.”	“I thought it was fairly random, so I didn’t invest much faith in it.” “Generally I just ignored it and made my own guess.”
Rely Heavily	“I kept my scores within 2 points of the compass score.”	“I relied on it, it eliminates bias.”
Starting Point	“I used that as a baseline.”	“It was only a starting point. I paid more attention to the criminal charge, prior charges and guessing on whether the defendant would be convicted.” “I took a look at the risk score to ballpark it.”
Validation	“I used to judge my final answer.” “I compared my answer to their answer and adjusted mine slightly based on theirs.”	“I made my own guess on what I thought and then checked the compass to find a score I was happy with.” “I used the Compas core to verify my decision.”
Factor	“I used it in combination with other factors to make my ultimate decision.” “I used them in consideration with the other data provided.”	“I tended to look at them but used the seriousness of crime or amount of past crimes and age as my main deciding factor.” “I took it into consideration along with all the other information given.”

Discussion

When assessing the risk that a defendant will recidivate, the COMPAS algorithm achieves a significantly higher accuracy rate than participants who assess defendant profiles (65.0 percent vs. 54.2 percent). The results from this experiment, however, suggest that merely providing humans with algorithms that outperform them in terms of accuracy does not necessarily lead to better outcomes. When participants incorporated the algorithm's risk score into their decision-making process, the accuracy rate of their predictions did not significantly change. The inclusion of a written advisement providing information about potential biases in the algorithm did not affect participant accuracy, either.

Given research in complementary computing that shows coupling human and machine intelligence improves their performance,^{2,9,11} this finding seems counterintuitive. Yet successful instances of human and machine collaboration occur under circumstances in which humans and machines display different strengths. Dressel and Farid's study demonstrates the striking similarity between recidivism predictions by Mechanical Turk workers and the COMPAS algorithm.⁶ This similarity may preclude the possibility of complementarity. Our study reinforces this similarity, indicating that the combination of human and algorithm is slightly (although not statistically significantly) worse than the algorithm alone and similar to the human alone.

Moreover, this study shows that the accuracy of participant predictions of recidivism does not significantly change when a written advisement about the appropriate

usages of the COMPAS algorithm is included. The Wisconsin Supreme Court mandated the inclusion of an advisement without indicating that its effect on officials' decision-making was tested.¹¹ Psychology research and survey-design literature indicate that people often skim over such disclaimers, so they do not perform their intended purpose.¹⁰ In concurrence with such theories, the results here suggest that written advisements accompanying algorithmic outputs may not affect the accuracy of decisions in a significant way.

EXPERIMENT TWO: ALGORITHMS AS ANCHORS

The first experiment suggested that COMPAS risk scores do not impact human risk assessments, but research in psychology implies that algorithmic predictions may influence humans' decisions through a subtle cognitive bias known as the *anchoring effect*: when individuals assimilate their estimates to a previously considered standard. Amos Tversky and Daniel Kahneman first theorized the anchoring heuristic in 1974 in a comprehensive paper that explains the psychological basis of the anchoring effect and provides evidence of the phenomenon through numerous experiments.¹⁹ In one experiment, for example, participants spun a roulette wheel that was predetermined to stop at either 10 (low anchor) or 65 (high anchor). After spinning the wheel, participants estimated the percentage of African nations in the United Nations. Tversky and Kahneman found that participants who spun a 10 provided an average guess of 25 percent, while those who spun a 65 provided an average guess of 45 percent. They rationalized these results by explaining that people make estimates by

starting from an initial value, and their adjustments from this quantity are typically insufficient.

While initial experiments investigating the anchoring effect recruited amateur participants,¹⁹ researchers also observed similar anchoring effects among experts. In their seminal study from 1987, Gregory Northcraft and Margaret Neale recruited real estate agents to visit a home, review a detailed booklet containing information about the property, and then assess the value of the house.¹⁶ The researchers listed a low asking price in the booklet for one group (low anchor) and a high asking price for another group (high anchor). The agents who viewed the high asking price provided valuations 41 percent greater than those who viewed the lower price, and the anchoring index of the listing price was likewise 41 percent. Northcraft and Neale conducted an identical experiment among business school students with no real estate experience and observed similar results: the students in the high anchor treatment answered with valuations that exceeded those in the low anchor treatment by 48 percent, and the anchoring index of the listing price was also 48 percent. Their findings, therefore, suggested that anchors such as listing prices bias the decisions of trained professionals and inexperienced individuals similarly.

More recent research finds evidence of the anchoring effect in the criminal justice system. In 2006 Birte Englich, Thomas Mussweiler, and Fritz Strack conducted a study in which judges threw a pair of dice and then provided a prison sentence for an individual convicted of shoplifting.⁷ The researchers rigged the dice so that they would land on a low number (low anchor) for half of the participants and

a high number (high anchor) for the other half. The judges who rolled a low number provided an average sentence of five months, whereas the judges who rolled a high number provided an average sentence of eight months. The difference in responses was statistically significant, and the anchoring index of the dice roll was 67 percent. In fact, similar studies have shown that sentencing demands,⁷ motions to dismiss,¹³ and damages caps¹⁵ also act as anchors that bias judges' decision-making.

Methods

This second experiment thus sought to investigate if algorithmic risk scores influence human decisions by serving as anchors. The experiment entailed a 1 x 2 between-subjects design where the two treatments were as follows: low-score, in which participants viewed the defendant profile accompanied by a low risk score; and high-score, in which participants viewed the defendant profile accompanied by a high-risk score.

The low-score and high-score treatments assigned risk scores based on the original COMPAS score according to the following formulas:

$$\text{Low-score} = \max(0, \text{COMPAS} - 3)$$

$$\text{High-score} = \min(10, \text{COMPAS} + 3)$$

This new experiment mirrored the previous one: Participants evaluated the same 40 defendants, met the same requirements, and received the same payment. The study also employed the format on the Qualtrics platform.

Results

Figure 4 shows the average scores of participants assigned to defendants versus those provided in the defendant profiles in the low-score and high-score treatments. Error bars represent the 95 percent confidence intervals. The scores that participants assigned defendants highly correlate with those that they viewed in the defendants' profile descriptions. As such, participants in the low-score treatment provided risk scores that were, on average, 42.3 percent lower than participants in the high-score

FIGURE 4: **AVERAGE RISK SCORE BY TREATMENT**

treatment when assessing the same set of defendants. The average risk score from respondents in the low-score treatment was 3.88 [95 percent CI 3.39-4.36], while the average risk score from respondents in the high-score treatment was 5.96 [95 percent CI 5.36-6.56]. A two-sided t-test revealed that this difference was statistically significant ($p < 0.0001$).

At the end of the survey, when participants reflected on the role of the COMPAS algorithm in their decision-making, they indicated common themes, such as using the algorithm's score as a starting point and as a verification of their own decisions. The table in figure 5 summarizes these participant comments by their treatment group and role of the algorithm in their decision-making.

Discussion

The results from this study indicate that algorithmic risk predictions serve as anchors that bias human decision-making. Participants in the low-score treatment provided an average risk score of 3.88, while participants in the high-score treatment assigned an average risk score of 5.96. The average anchoring index across all 40 defendants was 56.71 percent. This anchor measure mirrored that found in prior psychology literature.^{8,14,16} For example, one study investigated the anchoring bias in estimations by asking participants to guess the height of the tallest redwood tree.¹⁴ The researchers provided one group with a low anchor of 180 feet and another group with a high anchor of 1,200 feet, and they observed an anchoring index of 55 percent. Scholars have observed similar values of the anchoring index in contexts such as probability estimates,¹⁹

FIGURE 5: RESPONSES BY TREATMENT GROUP AND ALGORITHM ROLE

ALGORITHM ROLE	LOW TREATMENT	HIGH TREATMENT
Factor	"I took them into consideration but still made my own decisions."	"I took it into account, but did not count on it 100 percent." "I used it as <i>one</i> factor."
Tipping Point	"I used it if I was wavering on a score." "I would look at it if I was close on which way to go."	"For those cases where I felt a 50/50 chance, I sided with the Compas score."
Validation	"Only considered it when it seemed to coincide with my own judgement."	"I looked to see if it was similar to what I thought."
Guideline	"I used it to target my general range of scores, unless I had reason to strongly disagree." "I used the scores to base the start value of my score, read the description of their crime and modified the score."	"I kind of started with the Compas risk score, and then raised or lowered the score based on previous criminal history (or lack of one)." "I used it as a guideline to structure my decisions on."
Deference	"I always considered it as near perfect"	NA
Ignored	"I usually ignored it because it didn't seem like it made much sense to me, but who knows." "It didn't seem very consistent or accurate so I didn't factor it in much, if at all."	"I <i>hate</i> rubrics. You are looking at people, dynamic people, and a computerized rubric or other type of system designed to assess risks is completely ignoring so many other <i>*very important*</i> circumstances that may affect these odds."

purchasing decisions,²⁰ and sales forecasting.⁵
Even though this type of cognitive bias occurs among participants with little training in the criminal justice

system, prior work suggests that the anchoring effect varies little between non-experts and experts in a given field. Northcraft and Neale found that asking prices for homes similarly influenced real estate agents and people with no real estate experience.¹⁶ This study thus suggested that the anchoring effect of algorithmic risk assessments among judges, bail, and parole officers would mirror that of the participants in this experiment. Numerous prior studies demonstrate that these officials are, in fact, susceptible to forms of cognitive bias such as anchoring.^{7,15}

These findings also, importantly, highlight problems with existing frameworks to address machine bias. For example, many researchers advocate for putting a “human in the loop” to act in a supervisory capacity, and they claim that this measure will improve accuracy and, in the context of risk assessments, “ensure a sentence is just and reasonable.”¹² Even when humans make the final decisions, however, the machine-learning models exert influence by anchoring these decisions. An algorithm’s output still shapes the ultimate treatment for defendants.

The subtle influence of algorithms via this type of cognitive bias may extend to other domains such as finance, hiring, and medicine. Future work should, no doubt, focus on the collaborative potential of humans and machines, as well as steps to promote algorithmic fairness. But this work must consider the susceptibility of humans when developing measures to address the shortcomings of machine-learning models.

CONCLUSION

The COMPAS algorithm was used here as a case study

to investigate the role of algorithmic risk assessments in human decision-making. Prior work on the COMPAS algorithm and similar risk-assessment instruments focused on the technical aspects of the tools by presenting methods to improve their accuracy and theorizing frameworks to evaluate the fairness of their predictions. The research has not considered the practical function of the algorithm as a decision-making aid rather than as a decision maker.

Based on the theoretical findings from the existing literature, some policymakers and software engineers contend that algorithmic risk assessments such as the COMPAS software can alleviate the incarceration epidemic and the occurrence of violent crimes by informing and improving decisions about policing, treatment, and sentencing.

The first experiment described here thus explored how the COMPAS algorithm affects accuracy in a controlled environment with human subjects. When predicting the risk that a defendant will recidivate, the COMPAS algorithm achieved a significantly higher accuracy rate than the participants who assessed defendant profiles (65.0 percent vs. 54.2 percent). Yet when participants incorporated the algorithm's risk assessments into their decisions, their accuracy did not improve. The experiment also evaluated the effect of presenting an advisement designed to warn of the potential for disparate impact on minorities. The findings suggest, however, that the advisement did not significantly impact the accuracy of recidivism predictions.

Moreover, researchers have increasingly devoted attention to the fairness of risk-assessment software.

Related articles

➡ The Mythos of Model Interpretability

In machine learning, the concept of interpretability is both important and slippery.

Zachary C. Lipton

<https://queue.acm.org/detail.cfm?id=3241340>

➡ The API Performance Contract

How can the expected interactions between caller and implementation be guaranteed?

Robert F. Sproull and Jim Waldo

<https://queue.acm.org/detail.cfm?id=2576968>

➡ Accountability in Algorithmic Decision-making

A view from computational journalism

Nicholas Diakopoulos, University of Maryland, College Park

<https://queue.acm.org/detail.cfm?id=2886105>

While many people acknowledge the potential for algorithmic bias in these tools, they contend that leaving a human in the loop can ensure fair treatment for defendants. The results from the second experiment, however, indicate that the algorithmic risk scores acted as anchors that induced a cognitive bias: Participants assimilated their predictions to the algorithm's score. Participants who viewed the set of low-risk scores provided risk scores, on average, 42.3 percent lower than participants who viewed the high-risk scores when assessing the same set of defendants. Given this human susceptibility, an inaccurate algorithm may still

result in erroneous decisions.

Considered in tandem, these findings indicate that collaboration between humans and machines does not necessarily lead to better outcomes, and human supervision does not sufficiently address problems when algorithms err or demonstrate concerning biases. If machines are to improve outcomes in the criminal justice system and beyond, future research must further investigate their practical role: an input to human decision makers.

References

1. Angwin, J., Larson, J. 2016. Machine bias. *ProPublica* (May 23).
2. Case, N. 2018. How to become a centaur. *Journal of Design and Science* (January).
3. Chouldechova, A. 2017. Fair prediction with disparate impact: a study of bias in recidivism prediction instruments. *Big Data* 5[2], 153–163.
4. Corbett-Davies, S., Pierson, E., Feller, A., Goel, S., Huq, A. 2017. Algorithmic decision making and the cost of fairness. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM Press, 797–806.
5. Critcher, C. R., Gilovich, T. 2008. Incidental environmental anchors. *Journal of Behavioral Decision Making* 21[3], 241–251.
6. Dressel, J., Farid, H. 2018. The accuracy, fairness, and limits of predicting recidivism. *Science Advances* 4[1], eaao5580.
7. Englich, B., Mussweiler, T., Strack, F. 2006. Playing dice with criminal sentences: the influence of irrelevant anchors on experts' judicial decision making. *Personality and Social Psychology Bulletin* 32[2], 188–200.
8. Furnham, A., Boo, H. C. 2011. A literature review of the anchoring effect. *The Journal of Socio-Economics* 40[1], 35–42.
9. Goldstein, I. M., Lawrence, J., Miner, A. S. 2017. Human-machine collaboration in cancer and beyond: the Centaur Care Model. *JAMA Oncology* 3[10], 1303.
10. Green, K. C., Armstrong, J. S. 2012. Evidence on the effects of mandatory disclaimers in advertising.

- Journal of Public Policy & Marketing* 31(2), 293–304.
11. Horvitz, E., Paek, T. 2007. Complementary computing: policies for transferring callers from dialog systems to human receptionists. *User Modeling and User-Adapted Interaction*, 17(1-2), 159–182.
 12. Johnson, R. C. 2018. Overcoming AI bias with AI fairness. *Communications of the ACM* [December 6].
 13. Jukier, R. 2014. Inside the judicial mind: exploring judicial methodology in the mixed legal system of Quebec. *European Journal of Comparative Law and Governance* (February).
 14. Kahneman, D. 2011. *Thinking, Fast and Slow*. Farrar, Straus and Giroux.
 15. Mussweiler, T., Strack, F. 2000. Numeric judgments under uncertainty: the role of knowledge in anchoring. *Journal of Experimental Social Psychology* 36(5), 495–518.
 16. Northcraft, G. B., Neale, M.A. 1987. Experts, amateurs, and real estate: an anchoring-and-adjustment perspective on property pricing decisions. *Organizational Behavior and Human Decision Processes* 39(1), 84–97.
 17. Shaw, A. D., Horton, J. J., Chen, D. L. 2011. Designing incentives for inexperienced human raters. In *Proceedings of the ACM Conference on Computer-supported Cooperative Work*. ACM Press, 275-284.
 18. *State v Loomis*, 2016.
 19. Tversky, A., Kahneman, D. 1974. Judgment under uncertainty: heuristics and biases. *Science* 185(4157), 1124–1131.

20. Wansink, B., Kent, R. J., Hoch, S. J. 1998. An anchoring and adjustment model of purchase quantity decisions. *Journal of Marketing Research* 35(1), 71.

Michelle Vaccaro received a bachelor's degree in computer science in 2019 from Harvard College. She is particularly interested in the social implications of new technologies, and she hopes to pursue further research opportunities in that area.

Jim Waldo is a Gordon McKay Professor of the practice of computer science at Harvard University, where he is also a professor of technology policy at the Harvard Kennedy School. His interests include distributed systems, the intersection of technology, policy, and ethics, and privacy-preserving mechanisms. Prior to joining Harvard, he spent more than 30 years in the industry, much of that at Sun Microsystems.

Copyright © 2019 held by owner/author. Publication rights licensed to ACM.