A Case for Humans-in-the-Loop: Decisions in the Presence of Erroneous Algorithmic Scores

Maria De-Arteaga*

Heinz College
Machine Learning Department
Carnegie Mellon University
Pittsburgh, PA, USA
mdearte@andrew.cmu.edu

Riccardo Fogliato*

Department of Statistics and
Data Science
Carnegie Mellon University
Pittsburgh, PA, USA
rfogliat@andrew.cmu.edu

Alexandra Chouldechova

Heinz College
Carnegie Mellon University
Pittsburgh, PA, USA
achould@cmu.edu

ABSTRACT

The increased use of algorithmic predictions in sensitive domains has been accompanied by both enthusiasm and concern. To understand the opportunities and risks of these technologies, it is key to study how experts alter their decisions when using such tools. In this paper, we study the adoption of an algorithmic tool used to assist child maltreatment hotline screening decisions. We focus on the question: Are humans capable of identifying cases in which the machine is wrong, and of overriding those recommendations? We first show that humans do alter their behavior when the tool is deployed. Then, we show that humans are less likely to adhere to the machine's recommendation when the score displayed is an incorrect estimate of risk, even when overriding the recommendation requires supervisory approval. These results highlight the risks of full automation and the importance of designing decision pipelines that provide humans with autonomy.

Author Keywords

Human-in-the-loop; Decision support; Algorithm aversion; Automation bias; Algorithm assisted decision making; Child welfare

CCS Concepts

•**Human-centered computing** → **Human computer interaction (HCI)**; User studies; •**Information systems** → **Decision support systems**; •**Applied computing** → **Computing in government**;

INTRODUCTION

Risk assessment tools are increasingly being incorporated into expert decision-making pipelines across domains such as criminal justice, education, health, and public services [8, 26, 23, 7,

45]. These tools, which range from simple regression models to more complex machine learning models, distill available information on a given case into a risk score reflecting the likelihood of one or more adverse outcomes. Bolstered by decades of research showing that statistical models outperform human experts on prediction tasks, there is widespread optimism that these tools will increase the quality of decisions [32, 11, 20, 1, 25]. This optimism is tempered by evidence that, while humans provided with machine predictions may achieve improved performance, not only do they continue to underperform the machine predictions, but they may also uptake the information in ways that leads to increased disparities in decision outcomes across racial [19] and socioeconomic groups [42]. Such findings raise critical questions about the role of humans in the loop and human-machine complementarity in key societal domains. In this work we focus on one important question in this area: Are humans capable of identifying certain cases where the machine's recommendation is wrong, and appropriately overriding the recommendation in such cases?

We analyze a real world child welfare decision making context where call workers are tasked with deciding whether a call concerning potential child neglect or maltreatment should be screened in for investigation. While in many instances the information communicated in the call may be enough for the call worker to make a determination, in other instances the information may be vague and inconclusive. In an effort to better focus resources on investigating cases where the children are at greatest risk, Allegheny County has deployed a risk assessment tool called the Allegheny Family Screening Tool (AFST) to assist call workers in their screening decisions. The tool uses multi-system administrative data to assess the likelihood that children on the case will experience adverse child welfare events in the near future. More information about the tool and its development can be found on the County's AFST website [38].

Some time after the tool was deployed it was discovered that a technical glitch had resulted in a subset of model inputs being incorrectly calculated in realtime. This in turn resulted in misestimated risk scores being shown in some cases. While, as we elaborate on below, the misestimation was often mild, and the shown score generally provided reasonable risk information, the glitch permits us a rare opportunity to investigate

* Indicates equal contribution.

real world decision making in the presence of misestimated risk.

Before proceeding, we pause to make an important point. These types of technical issues are not uncommon. What is uncommon is for organizations to choose to be transparent about their occurrence. We recognize Allegheny County for their transparency and hope that this approach will become the norm in the deployment of algorithmic systems in sensitive societal domains.

The remainder of this paper is organized as follows. We begin with a discussion of related work in which we provide an overview of phenomena such as algorithm aversion and automation bias that come to bear on human decision making in the presence of algorithmic decision support tools. We then describe the child welfare decision making context, the risk assessment tool deployment setup, and our available data. Our analysis of the data begins by demonstrating that there was a marked change in workers' screening decisions in the post-deployment period. Having established that an overall change in behavior did occur, we then investigate the extent to which call workers deviate from recommendations based on a misestimated risk score. We show that workers are able to appropriately override the tool in many such cases. We also probe questions of potential disparities in adherence to recommendations across racial and socio-economic groups, finding that the deployment of the tool neither significantly mitigates nor exacerbates disparities observed at the given level of analysis. Lastly, we conclude with a discussion of human and system factors that we believe may have contributed to the observed results, and outline opportunities for further research to better understand relevant factors.

BACKGROUND AND RELATED WORK

Prior research has attempted to answer *whether* and *how* the deployment of algorithmic risk assessment tools affects users' decisions. While many have advocated for the adoption of these tools on the basis of their superior predictive accuracy, findings are mixed on whether integrating prediction tools into decision making significantly improves decision quality. Indeed, research in the field suggests that the outcomes of decisions taken by a human aided by a decision support system are often no better than those taken by the human alone.

Recent work has paid special attention to the introduction of risk assessment in the context of pretrial decision making in the criminal justice system. Although the integration of the risk assessment tools was, *ex ante*, expected to lead to a sharp and persistent decrease in incarceration rates, recent findings suggest that there is no impact at all [15] or find there is a decrease but of much smaller magnitude than initially hoped [44]. There is consensus that these lackluster results are due at least in large part to the wide heterogeneity in judges' compliance with the tools' recommendations [9]. Notably, differential compliance has been shown to be a factor driving increased poor-rich [42] and black-white [46, 2] disparities in the post-deployment period. For instance, [2] found that the increased racial gap in incarceration rates post-deployment was due both to inter-variation—judges in whiter counties showing higher compliance—and by intra-variations—overrides of low and

moderate risk being more frequent for black than for white defendants.

More broadly, there are two competing tendencies that have been observed in studies of human compliance with algorithmic recommendations: *algorithm aversion* and *automation bias*. *Algorithm aversion*—the tendency to ignore tool recommendations after seeing that they can be erroneous—originates from a lack of agency [29, 12] and lack of transparency of the algorithm [49]. Studies have shown that users will knowingly sacrifice accuracy in favor of gaining some control over the algorithm's output [14]. Similarly, [18] reports an experiment in which humans override the machine's predictions when these are highly reliable. Users' reliance on the system is known to vary with the observed [51, 52] and stated accuracy [50] of the system. However, even if the recommendations of more accurate systems are followed more often, agents affected by algorithm aversion may nevertheless prefer human judgment over algorithmic predictions even when evidence known to both the designer and the user clearly indicates that the algorithmic predictions are more accurate than human assessment [13].

Users affected by *automation bias*, on the other hand, will follow tool recommendations despite available (but unnoticed or unconsidered) information that would indicate that the recommendation is wrong. *Automation bias* consists of two classes of errors. Omission errors refer to instances where humans fail to detect problematic cases (or fail to act) because they were not flagged as such by the system. A prominent example is that of pilots in high-tech cockpits, who are prone to relying blindly on automated cues as a heuristic replacement for vigilant information seeking [35]. Commission errors refer to instances where humans take action on the basis of an erroneous algorithmic recommendation, failing to incorporate contradictory external information into the decision process. In the clinical decision support context, commission errors may result in patients being subjected to unnecessary, potentially invasive testing or treatment.

Studies analyzing factors contributing to automation bias have found that complex tasks and time pressure may increase over-reliance on decision support [41, 17]. The users' experience level and their confidence in their own decisions have also been found to be causes of automation bias [31, 33]. Social accountability has been found to reduce automation bias [43], an important result when considering decision support systems used by experts with high public visibility or who are publicly elected, such as judges. Meanwhile, studies focused on the causes of algorithm aversion have found that repeatedly seeing the algorithm make the same mistake leads to decreased reliance of the agent on the system [13], while giving some control over the algorithm can counter this phenomenon [14].

Automation bias and algorithm aversion are opposing phenomena. While automation bias degrades decision quality by driving over-compliance with algorithmic recommendations, algorithm aversion does so by driving under-compliance. There are two characteristics of the decision context that are indicative of which form of bias is likely to dominate: the type of task, and the level of automation. A significant portion of

the literature analyzing automation bias has studied high-tech cockpits [34, 43, 10, 41], while others have looked at health-care diagnosis [31], and automated control systems for fault management in a thermal-hydraulic environment [33]. All these are diagnostic tasks, where it can be assumed that there is a “ground truth” that is in principle knowable to humans. Algorithm aversion has been mostly discussed in tasks of a different nature: prognostic tasks, where predictions pertain what will happen in the future. These tasks are often behavioral prediction tasks, which concern people’s future actions or performance. Examples include predicting students’ performance on an MBA program using admissions data [13], forecasting sales [29, 18], and predicting which defendants will be rearrested or will fail to appear for court if released [12, 15, 44]. In these tasks there is an irreducible degree of uncertainty, and any predictive model, whether human or machine, is bound to make mistakes. The role of the level of automation on over-reliance in decision support systems has been further discussed in [10].

Motivated by these challenges, recent work has explored approaches to characterizing human-machine complementarity in risk assessment contexts and devising approaches to combine the strengths of both. [47] study differences in factors relied upon by human (crowdworker) recidivism predictions and those of COMPAS, a commercial risk assessment tool. The authors identify systematic differences in human and machine predictions, but find that those differences could not be leveraged to improve predictions. [21] suggest that, instead of generating predictions, the algorithm should be trained with a humans-in-the-loop to incorporate the human decision process, and should only report to the human simpler and useful representations of the features. [30] propose a learning to defer model, where the algorithm chooses to make a prediction or defer to the human taking into account both the model’s and the human’s accuracy. Lastly, focusing on a medical diagnosis task, [40] propose a method for triaging cases between full automation and focused human effort in a manner that improves diagnostic accuracy above the human and machine predictions alone.

Recent work has also studied the effect of providing explanations on adherence to algorithmic recommendations. [24] and [37] show that the perceived accuracy of the system depends on the degree to which explanations are easily understandable, with very complicated explanations reducing perceived accuracy. At the same time, explanations can also mislead users, as shown by [27]. In the context of Facebook’s News Feed algorithm, [39] found that while explanations helped users understand how the system works, they did not help them in evaluating the correctness of the output. In the context of Reddit post removal, [22] found that providing explanations helped users better adhere to community guidelines. Explanations may also affect perceptions of justice of algorithmic decisions, as studied by [4]. Such perceptions may not only relate to the output but also to the predictors used by the system [48]. Explanations may also affect perceptions of justice of algorithmic decisions, as studied by [4]. Such perceptions may not only relate to the output but also to the predictors used by the system [48].

In many of these studies adherence is taken to be synonymous with trust. Indeed, while there is no single commonly adopted definition of trust in the HCI literature, the term trust typically refers to a measure of, or the factor influencing, the degree to which the human is willing to delegate decision-making to the machine in absence of complete knowledge of the algorithmic pipeline [28, 51, 3, 37, 50]. We note that in our setting—and in any setting where the objective optimized for by the algorithmic system does not fully capture all relevant costs and payoffs of the decision—trust and adherence are not one and the same. In the child welfare context, the risk assessed by the tool is just one factor relevant to determining whether an investigation is appropriate. Other factors such as resources, recent investigations, information conveyed during the allegation call, are all also relevant. Workers are not expected to adhere to the scores in all cases, and may trust the ability of the tool to assess certain dimensions of risk even if in a given instance that risk is not the most relevant decision factor. Both [8] and [5] provide more indepth discussion of algorithmic tools and their use in the child welfare system. Most notably, [5] take a deep dive into factors that influence affected communities perceptions comfort and trust in the use algorithmic decision support in the child welfare context. The authors find that general distrust in the child welfare system, as well as specifics of how and whether risk information is communicated to case workers and families are important to perceptions of procedural and interpersonal justice.

Our work contributes to the nascent literature on human-machine complementarity in risk assessment by investigating whether humans are able to correct for misestimated risk scores in a real world decision making context. Our study relies on a retrospective analysis of observed call worker decisions before and after a risk assessment tool was integrated into the decision making pipeline. A similar context but in a different setting is analysed by [6], which study the effects of inconsistencies in sentencing recommendations due to human errors in the judicial setting in Maryland. The authors find that judges are more likely to go along with mistakenly lesser sentences for violent offenses, but tend to discount recommendations that are mistakenly too high.

DEPLOYMENT SETUP AND DATA

Call workers at the child welfare hotline are tasked with deciding whether a call alleging potential child maltreatment or neglect should be screened in for investigation. In making their decisions, call workers in Allegheny County have access to the information communicated in the referral call, along with multi-system administrative data on demographics, child welfare involvement, criminal history, and other information related to the children and adults associated to a referral. The administrative data consists of hundreds of data elements. It is therefore challenging for workers to make systematic and effective use of the administrative data in each case. In order to help workers make better use of this data, Allegheny County introduced a risk assessment tool that distills the information contained therein into a single risk score reflecting the likelihood that the children on the referral will experience adverse child welfare related outcomes in the months following the referral. The intended use of the tool is to help workers

identify high-risk cases in instances where the information communicated in the call may be insufficient, inconclusive, or otherwise incomplete in reflecting the immediate or long-arc risk of the children. As we further discuss in Section 3.2, specific guidelines were created to strongly encourage screen-ins (investigations) for the highest scoring cases.

Predictive model

Throughout the paper, we refer to the case associated to a call as a *referral*, each of which may have several *referral records* associated to it, one for each children involved in the call. The deployed tool was trained with all referral records collected by Allegheny County between April 2010 and July 2014. Two distinct predictive models estimate the probability of *out-of-home placement* and of *future referral* for each child based on features that include demographics, past welfare interaction, public welfare, adult and juvenile criminal justice involvement, and behavioral health information available on all persons associated with each referral. *Out-of-home placement* refers to whether the child is placed out of the home following an investigation, and a *future referral* refers to a future call involving the child coming in to the hotline. The predicted probabilities are then converted into an integer score in the range from 1 to 20, corresponding to the ventiles. The score to be shown to the workers is calculated as the maximum score over both models, over all children involved in a referral. We denote this aggregated score as $S \in \{1, \dots, 20\}$. A more detailed description of the model can be found on the County’s website [38].

Deployment

By design, call workers are only shown risk scores for cases with sufficient information. During the analyzed deployment period, 92.5% of referrals had an associated score shown. In addition to displaying a risk score, the tool assigns a label of “mandatory screen-in” to certain referrals, which means that the supervisor’s approval is required in order to screen out the referral. These correspond to the cases whose calculated score for out-of-home placement is greater or equal to 18. That is, these are cases where at least one child associated to the referral has a *placement* score of at least 18. Figure 1 presents a graphical representation of the decision-making pipeline.

Score misestimation

During deployment, a glitch in the system led to certain model inputs not being calculated correctly in realtime. Figure 3a shows a histogram of the fraction of values for each feature whose realtime values did not match retrospectively recalculated values. There are two reasons for this mismatch. The first is due to an issue, since resolved, wherein the realtime database queries were erroneously returning counts and indicators of 0. Figure 3b compares the fraction of referral records where the feature was calculated as 0 in real time vs. retrospectively for the top six most frequently mismatched binary indicators. The second issue is that as cases evolve the roles of different adults associated to the case and information about them may evolve. For instance, the individual identified as a perpetrator may change between the initial run and the

Figure 1: Decision pipeline for child welfare services in Allegheny County. Re-referral risk score and out-of-home placement (ooh) risk score are calculated for all children associated to a referral. The maximum over these scores is shown to the call worker, who also has access to historical records and information conveyed in the call. The call worker decides whether to screen in the call for investigation. Screening out calls with an ooh risk score ≥ 18 requires supervisor’s approval.

retrospective analysis. These more naturally occurring mismatches are unlikely to have as significant an impact on the calculated score.

Figure 2: Heatmap of the density of shown score $\tilde{S}$ conditional on assessed score S . A cell at row r and column c shows the fraction of the time that a referral assessed at a score of $S = c$ was shown to have a score of $\tilde{S} = r$.

As a result of the glitch in how model inputs were being calculated in realtime, the score displayed (‘shown’) to workers during deployment did not always correspond to the score that should have been shown. Figure 2 shows the distribution of the scores. The concentration of cases around the diagonal indicates that the shown scores were most often equal or very close to the assessed score that should have been displayed. These circumstances of deployment allow us to study the behavior of call workers when the shown score is inaccurate. In

Figure 3: Analysis of the glitch. (a) shows the normalized histogram of the fraction of referral records affected by the glitch for each feature. For the six (binary) features with highest fractions, (b) shows their distribution with and without the glitch.

particular, we are able to analyze what happens when differences between the assessed and the shown score result in a different “mandatory screen-in” policy being applied.

The terminology and notation we use throughout the paper are detailed in Table 1.

Data

The risk assessment tool was deployed in August 2016. We use data from January 2015 to July 2016 to analyze the behavior of call workers before the deployment of the tool, and data from August 2016 until December 2017 to study their behavior after adoption.

We constrain our analysis to the 92.5% of referrals that had an associated score shown. Moreover, existing regulation dictates that certain calls must be investigated. Referrals that fall under this category include those in which the call concerns bodily injury and sexual abuse. Since the call worker had no discretion on deciding whether or not to screen in these calls, we exclude from our analysis the 19% of referrals that fall under this legislation. Finally, we also exclude the referrals that are associated to open investigations, since the call worker does not make any determination in these cases. This excludes 19% of the remaining observations. Once we have constrained the data to those cases that had a score shown and for which the call worker made a determination, we are left with 27,575 referral records, corresponding to 11,802 referrals (recall that multiple children may be associated to a call, leading to multiple referral records within each referral).

ANALYSIS

Change in call workers’ behavior

The main question we seek to answer in this section is whether call workers updated their screening behavior after the tool was deployed. In other words, did algorithm aversion lead to the tool being ignored altogether?

We first analyze the behavior of call workers with respect to the assessed score, S , before and after the deployment of the tool. While ideally we would also display similar results for the shown score, $\hat{S}$, the glitch that produced $\hat{S}$ is not retrospectively reproducible. We are unable to calculate what $\hat{S}$ would have been in the pre-deployment period. However, since pre- and post-deployment decisions would remain the same with respect to both S and $\hat{S}$ if workers did not alter their decisions, it is nevertheless insightful to study variation with S .

As we now discuss, the results indicate that workers did update their behavior. First, we note that the overall screen-in rate did not vary before and after deployment, remaining around 45%. This stability over time is likely explained by the fact that the agency investigates as many cases as their resources allow. However, we do observe a change on *which* cases were investigated. Figure 4 shows the screen-in rates across values of the assessed risk score S . The steeper slope of the post-deployment curve, particularly for very high and very low risk cases, indicates that post-deployment screen-ins are better aligned with the score. We see a pronounced increase in the screen-in rate of highest risk cases, and a pronounced decrease in the screen in rates for low and moderate risk cases.

Recall that any referral where the placement score was 18 or greater was flagged as a “mandatory screen-in”. Shown mandates required supervisory approval to be screened out. While not all assessed mandates ($M = 1$) where shown mandates ($\hat{M} = 1$) and vice versa, it is nevertheless instructive to look at how screen-in rates varied pre- and post- deployment with respect to M . Figure 5 shows a timeline of screen-in rate of cases corresponding to $M = 1$. A sharp increase in screen-in rates for assessed mandatory cases can be seen the month of deployment. Overall, in the period before deployment, cases that would have fallen in the category $M = 1$ had a screen-in rate of 58%, while after deployment this rate increased to 71%. As there appears to be no change in the jurisdiction’s policies

Table 1: Terminology and notation used throughout the paper.

Term	Notation	Description
Assessed score	$S \in \{1, \dots, 20\}$	Score assessed by the predictive model. Maximum over assessed risk of out-of-home placement and assessed risk of re-referral.
Shown score	$\tilde{S} \in \{1, \dots, 20\}$	Score shown to the call worker. It should correspond to the assessed score, but is subject to a glitch in the system.
Assessed mandatory screen-in	$M \in \{0, 1\}$	$M = 1$ if assessed risk of out-of-home placement greater or equal to 18. Supervisor’s approval required for screen out.
Shown mandatory screen-in	$\tilde{M} \in \{0, 1\}$	Mandatory screen-in label shown to the call worker. It should correspond to assessed mandatory screen-in, but is subject to a glitch in the system.

Figure 4: Screen-in rates by assessed score S before and after deployment. After deployment screen-in decisions are better aligned with the score.

or types of calls received around the time of deployment, we attribute the sudden increase in screen-in rate to the deployment of the model. This increase is to be expected given that there is a barrier to overriding cases $\tilde{M} = 1$, which is strongly correlated with M . Notably, the increase is smaller than what would be expected if automation bias were to occur. Even though there is a cost-barrier to override the machine (required supervisory approval), only 66% of cases shown as mandatory screen-in, $\tilde{M} = 1$, are screened in.

Overrides of erroneous scores

In this section we investigate whether call workers detect and correct mistakenly computed algorithmic scores. We study both types of errors common to the phenomenon of automation bias: omission errors and commission errors. That is, we investigate whether call workers commit errors of omission—screening out shown low-risk cases that are assessed as high(er) risk—and errors of commission—screening

Figure 5: Timeline of screen-in rates for cases with mandatory screen-in policy, $M = 1$. Dotted line marks date of deployment. We find a sharp and persistent increase in screen in rates for assessed mandatory cases. Some of this increase is attributable to assessed mandatory cases that were also correctly shown as mandatory screen-ins. Figure 6: Alignment of screen-in rates with respect to assessed score S and shown score $\tilde{S}$. We find that post-deployment decisions are better aligned with the assessed score than the score that was shown, indicating that workers were successful in correcting for the score miscalculation in certain cases.

Figure 7: Analysis of human decisions. Fraction of screen-in's for: (a) overestimation ($\tilde{S} > S$), equality ($\tilde{S} = S$), underestimation ($\tilde{S} < S$) of the assessed score for binned shown scores; (b) binned shown score for cases in which the screen-in should have been mandatory, $M = 1$. Error bars indicate 95% confidence intervals.

in shown high-risk cases that are assessed as lower risk. For this analysis, we use the fact that the shown score did not always correspond to the predicted score during the analyzed post-deployment period. Therefore, in some cases the score shown corresponded to an underestimation of risk, while in others it corresponded to an overestimation.

It is remarkable to note that the human decisions are better calibrated with respect to the assessed score S than with respect to the shown score $\tilde{S}$, as seen in Figure 6. The relatively high screen-in rates for cases with very low shown scores suggest that the call workers are effectively using information at their disposal to avoid many errors of omission, and choosing to screen in around 30% of cases that are shown to have the lowest risk scores. Meanwhile, only a very small fraction of cases with an assessed score in the lowest buckets were screened in. This is consistent with what we observe at the high risk end of the scales. Again, call workers seem to be making use of available information to choose not to screen in cases that the machine marks as high-risk, thereby avoiding errors of commission. This result is particularly surprising given that overriding many of these cases would require a supervisor's approval. When focusing on those that required approval to be screened out, $\tilde{M} = 1$, call workers screen in 66% of the cases. This rate goes down to 57% when only considering those that should not have been in this bucket (cases where $\tilde{M} = 1$ and $M = 0$). This indicates that humans are more likely to override a shown mandatory screen in when it is not an assessed mandatory. It is not surprising that screen in rates are high for shown mandatories overall, as high shown scores correlated strongly with high assessed scores.

In order to explore this behavior further, we look at decisions for under- and over-estimated scores in Figure 7. Figure 7a shows the percentage of screen-in's for binned $\tilde{S}$. *Correct*

indicates that the assessed score is equal to the shown score, *underestimation* means that the score shown was *lower* than the assessed score, and *overestimation* means that the score shown was *higher* than the assessed score. If the call workers were to blindly follow the risk tool, we would observe that within each score bucket the screen-in rates are the same across all three classes. We observe something very different. It is clearly evident from the analysis that, among cases with similar shown scores, cases for which the score was underestimated are screened in at *much* higher rates than the others. What this means is that humans are seemingly able to identify that the risk is being underestimated for these cases. For example, for the cases with a shown score between 11 and 15, those for which this score was an underestimation of the assessed score were screened in at a rate of almost 60%, while the other cases with a shown score between 11 and 15 were screened in at rates around 30%. This suggests that the call workers make use of other pieces of information—either from the call or from the administrative data system directly—and respond appropriately by screening-in with higher probability. Somewhat surprisingly, we do not observe such a stark pattern for overestimated cases except in the highest risk bucket, where call workers are less likely to screen-in cases that should have had lower scores assigned.

While Figure 7a allows us to observe certain patterns, it does not differentiate between over and underestimations of different magnitude. For that reason, we zoom-in on the decisions made for assessed mandatories ($M = 1$). Figure 7b shows the proportion of cases $M = 1$ that were screened in for each bucket of shown score, $\tilde{S}$. This rate is approximately constant across all buckets, suggesting that the underestimation of the assessed score had no effect on the call workers' actions and they were able to make use of other information to identify these high risk cases. Even when the shown placement score

Figure 8: Rate of screened-in cases accepted for service across buckets of score shown, grouped according to the relationship between the score shown and the assessed score.

was more than 12 points lower than the assessed placement score, the screen-in was still as likely as for a shown mandatory case.

The first part of our analysis established that workers did update their behavior when the tool was deployed. Meanwhile, the findings summarized in Figure 6 and Figure 7 indicate that they do not follow the score blindly. Instead, workers successfully make use of other sources of information and are more likely to override the machine when the shown score is significantly miscalculated.

What can we say about the quality of these decisions, beyond their relationship with the assessed predicted risk? When a call is screened in, a social worker visits the family and is tasked with deciding whether to open a child welfare case, which is termed as “accept for service”. If call workers’ performance had been perfect and there were no resource constraints, we could expect all referrals they screen in to have an case opened. Before deployment of the tool, 18% of investigated referrals were accepted for services, 19% were connected to an existing open case involving the family, and 63% were not accepted for service. Post-deployment, these rates were 21%, 23%, and 56%, respectively. This indicates a higher precision in the post-deployment period: more of the screened-in referrals were being provided with services.

Figure 8 shows the rates of cases accepted for service across buckets of scores, grouped according to the relationship between the score shown and the assessed score—the same grouping used in Figure 7a. For example, for the cases with a shown score between 11 and 15, more than half of those for which this score was an underestimation of the assessed score were accepted for services upon screen-in. Meanwhile, this statistic drops to 15% for all other cases with a score shown in this bucket. The similarity between Figure 8 and Figure 7a is striking.

Figure 9: Screen-in rates by race pre- and post-deployment. (a) overall; (b) for subset of cases $M = 1$.

ing. In Figure 7a we observed that humans are very good at identifying cases that had an underestimated score, being more likely to screen these in. In Figure 8 we observe that these are indeed more likely to be accepted for service following an investigation.

Disparities in decision-making

One of the main concerns surrounding the implementation of risk assessment tools is that this may exacerbate disparities, harming already marginalized groups. In the scenario of child welfare, disparities across race and income level are of particular concern [16]. While it is important to ensure that children are being protected, it is also important to avoid overburdening a group with interventions that may have harmful unintended consequences.

Figure 9 shows the screen-in rates for Black and White children before and after the deployment of the tool. We observe a slight reduction on screen-in rates for Blacks and a slight increase for Whites. Prior work has identified differences in overrides across groups as a source of disparities. We therefore also analyze if there is a difference in alignment with the assessed mandatory screen-in policy across groups. Figure 9b shows the screen-in rates for the subset of cases for which $M = 1$. Here, we observe that screen-in rates for both races increase, although the increase is slightly sharper for White children. These results indicate that, unlike what has been observed in other domains, there does not appear to be a difference in willingness to adhere to the recommendation that would compound previous racial injustices.

We also consider disparities across socioeconomic factors. Because the features available for risk assessment rely in part in previous interactions with the public services, one concern that arises is whether differential reliance on the risk assessment tool will disadvantage those who have sought out the government’s support in the past. Moreover, and most relevant to the focus of this work, it would be concerning if willingness to override the machine’s prediction varied depending

on how wealthy the children’s family is. If this was the case, we should observe that screen-in rates for cases labeled as $M = 1$ decrease as the wealth increased. Figure 10 shows the screen-in rates by poverty rates, defined as % of households below the poverty level in the persons’ neighbourhood. Here, we see that, among assessed mandatory cases, screen in rates do not appear to correlate with neighbourhood poverty levels.

Figure 10: Screen-in rates by fraction of households below the poverty level in the persons’ neighbourhood, pre- and post-deployment, for subset of cases with predicted mandatory screen-in, $M = 1$.

DISCUSSION

Our analysis studied the adoption of a risk assessment tool that assists call workers tasked with deciding which calls concerning potential child neglect or abuse should be screened in for further investigation. We focused our analysis on investigating whether humans successfully override or ignore erroneous algorithmic scores. We have found that call workers did change their behavior when the tool was deployed, showing partial adherence to the machine’s recommendation. We also found that call workers were less likely to adhere to the recommendations in cases where a technical glitch resulted in an incorrect score being displayed.

When considering how humans make use of recommendations provided by an algorithmic system, algorithm aversion and automation bias can be seen as two ends of a broad spectrum, both of which are undesirable. On one end, algorithm aversion leads the human to completely disregard the machine, even when the recommendation may be providing useful information that would improve the quality of decisions. On the other end, automation bias results in humans blindly following the machine’s recommendations, failing to make use of other sources of information and their own judgement to disregard the recommendation when evidence suggests that the machine may be mistaken. The call workers in our study were found to not be at either of these extremes. Call workers changed their decision-making following the deployment of the tool, but they did not adhere to the tool’s recommendations in every instance.

A clear limitation of our study is that it is retrospective in nature. However, this retrospective data also offers a rare opportunity to study a phenomenon that would not be ethically feasible to investigate as a randomized field trial. By virtue of being a study of a real deployed system, our study has the advantage of inherent field validity. While crowd worker studies and randomized field experiments allow for controlled variation of different factors, even the best conceived randomized trials can fail to have field validity in social policy settings [36]. Moreover, given the high-stakes nature of the task, behaviors displayed by call workers when given hypothetical calls in a lab setting may differ from their behavior when faced with a real allegation of child abuse. Our investigation of the technical glitch demonstrates that certain phenomena are in fact observed to occur when trained experts are making high-stakes decisions in a real world setting.

What contributed to these positive results? The retrospective nature of our study means that we were unable to perturb different elements of the decision making framework in order to assess their impact on decision outcomes. However, through our analysis and discussions with jurisdiction staff, we are able to identify certain elements of the deployment setup that we believe likely influenced the observed behavior, which we elaborate on below. An important direction for future research is to further investigate these and other factors in more controlled settings to gain a better understanding of their influence on commission and omission errors in algorithm assisted decision making.

A key contributing factor is that throughout the process call workers continued to have access to not only the referral calls but also the administrative data system. This provided a different view of the case than what was being pulled into the risk score calculation. In particular, even when inputs related to past child welfare history were being miscalculated in real time, workers would still have access to the correct information in the data system. In addition to the role of having access to the raw features, and having the time to inspect these, it is particularly important to highlight that call workers had been previously trained to make decisions without the aid of a risk assessment score. Therefore, they had experience in parsing and interpreting the raw data. A question that arises is whether this previous experience had an important role, and whether the same behavior can be expected from call workers who start working after the deployment of the tool. The answer to this question could inform the need for training that teaches experts how to make use of other sources of data appropriately, in order to avoid an over-reliance on algorithmic recommendations.

Secondly, the risk tool provides workers only with a score, and does not ‘explain’ its predictions, nor does it display values of any of the features involved in the score calculation. If this additional information had been provided, it is possible that the glitch would have been detected by workers. However, it is also possible that this distillation of the data would have been trusted by workers, who would in turn have been dissuaded from examining the original data. In the latter case, the

workers would likely have been more susceptible to omission errors.

With regard to design recommendations, the results highlight the fact that humans-in-the-loop can help guard against harmful effects resulting from erroneous algorithmic recommendations. Evidence in other domains has shown that humans-in-the-loop do not always improve the quality of decisions, leading some to call for complete automation of tasks. However, providing humans with autonomy to contradict the machine mitigated the effects of miscalculated scores in the child maltreatment call screening context. Given that technical glitches as the one studied in this paper are always a risk, and that any statistical model will make mistakes, design should focus on augmenting the human's ability to identify and correct mistakes. Future research in controlled settings that evaluates the effect of individual elements of the decision-making pipeline described in this paper could identify specific design practices that effectively strengthen the human's role.

More research is also needed to understand how to develop machine learning decision support tools that provide information helpful to the humans in assessing whether the machine's predictions effectively account for all factors relevant to a case. This is an especially challenging task in forecasting settings where research shows that human predictions tend to underperform those of statistical models. Whereas humans (or a consensus of humans) may be expected to identify cases where an image classifier incorrectly tags a dog as a cat, humans are not as reliable in assessing which students will succeed, which children are at risk, or which defendants will appear for court. We hope that the findings presented in this paper will motivate further work on developing systems that make it easier for humans to assess whether the machine may be making a mistake in the forecasting setting. In particular, many of the decision-pipeline elements likely contributed to workers selectively trusting the machine by questioning its recommendations. Further research is needed to better understand the role that trust and other factors play in the overall success of human-in-the-loop algorithmic systems.

ACKNOWLEDGMENTS

We are grateful to the Hillman Foundation for funding this research, and to our collaborators at the Allegheny County Department of Human Services. Thanks also to the anonymous reviewers for their numerous comments and suggestions that helped to improve the manuscript.

REFERENCES

- [1] Stefánía Ægisdóttir, Michael J White, Paul M Spengler, Alan S Maugherman, Linda A Anderson, Robert S Cook, Cassandra N Nichols, Georgios K Lampropoulos, Blain S Walker, Genna Cohen, and others. 2006. The meta-analysis of clinical judgment project: Fifty-six years of accumulated research on clinical versus statistical prediction. *The Counseling Psychologist* 34, 3 (2006), 341–382.
- [2] Alex Albright. 2019. *If You Give a Judge a Risk Score: Evidence from Kentucky Bail Decisions*. Technical Report. Working paper.
- [3] Gagan Bansal, Besmira Nushi, Ece Kamar, Daniel S Weld, Walter S Lasecki, and Eric Horvitz. 2019. Updates in human-ai teams: Understanding and addressing the performance/compatibility tradeoff. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 2429–2437.
- [4] Reuben Binns, Max Van Kleek, Michael Veale, Ulrik Lyngs, Jun Zhao, and Nigel Shadbolt. 2018. 'It's Reducing a Human Being to a Percentage': Perceptions of Justice in Algorithmic Decisions. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, 377.
- [5] Anna Brown, Alexandra Chouldechova, Emily Putnam-Hornstein, Andrew Tobin, and Rhema Vaithianathan. 2019. Toward Algorithmic Accountability in Public Services: A Qualitative Study of Affected Community Perspectives on Algorithmic Decision-making in Child Welfare Services. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, 41.
- [6] Shawn D Bushway, Emily G Owens, and Anne Morrison Piehl. 2012. Sentencing guidelines and judicial discretion: Quasi-experimental evidence from human calculation errors. *Journal of Empirical Legal Studies* 9, 2 (2012), 291–319.
- [7] Rich Caruana, Yin Lou, Johannes Gehrke, Paul Koch, Marc Sturm, and Noemie Elhadad. 2015. Intelligible models for healthcare: Predicting pneumonia risk and hospital 30-day readmission. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 1721–1730.
- [8] Alexandra Chouldechova, Diana Benavides-Prado, Oleksandr Fialko, and Rhema Vaithianathan. 2018. A case study of algorithm-assisted decision making in child maltreatment hotline screening decisions. In *Conference on Fairness, Accountability and Transparency*. 134–148.
- [9] Alma Cohen and Crystal S Yang. 2019. Judicial politics and sentencing decisions. *American Economic Journal: Economic Policy* 11, 1 (2019), 160–91.
- [10] Mary Cummings. 2004. Automation bias in intelligent time critical decision support systems. In *AIAA 1st Intelligent Systems Technical Conference*. 6313.
- [11] Robyn M Dawes, David Faust, and Paul E Meehl. 1989. Clinical versus actuarial judgment. *Science* 243, 4899 (1989), 1668–1674.
- [12] Matthew DeMichele, Peter Baumgartner, Kelle Barrick, Megan Comfort, Samuel Scaggs, and Shilpi Misra. 2018. What do criminal justice professionals think about risk assessment at pretrial? *Available at SSRN 3168490* (2018).
- [13] Berkeley J Dietvorst, Joseph P Simmons, and Cade Massey. 2015. Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General* 144, 1 (2015), 114.

- [14] Berkeley J Dietvorst, Joseph P Simmons, and Cade Massey. 2016. Overcoming algorithm aversion: People will use imperfect algorithms if they can (even slightly) modify them. *Management Science* 64, 3 (2016), 1155–1170.
- [15] Jennif Doleac and Megan Stevenson. 2018. *The roadblock to reform*. Technical Report. American Constitution Society Research Report.
- [16] Virginia Eubanks. 2018. *Automating inequality: How high-tech tools profile, police, and punish the poor*. St. Martin's Press.
- [17] Kate Goddard, Abdul Roudsari, and Jeremy C Wyatt. 2011. Automation bias: a systematic review of frequency, effect mediators, and mitigators. *Journal of the American Medical Informatics Association* 19, 1 (2011), 121–127.
- [18] Paul Goodwin and Robert Fildes. 1999. Judgmental forecasts of time series affected by special events: Does providing a statistical forecast improve accuracy? *Journal of Behavioral Decision Making* 12, 1 (1999), 37–53.
- [19] Ben Green and Yiling Chen. 2019. Disparate interactions: An algorithm-in-the-loop analysis of fairness in risk assessments. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*. ACM, 90–99.
- [20] William M Grove, David H Zald, Boyd S Lebow, Beth E Snitz, and Chad Nelson. 2000. Clinical versus mechanical prediction: a meta-analysis. *Psychological assessment* 12, 1 (2000), 19.
- [21] Sophie Hilgard, Nir Rosenfeld, Mahzarin R Banaji, Jack Cao, and David C Parkes. 2019. Learning Representations by Humans, for Humans. *arXiv preprint arXiv:1905.12686* (2019).
- [22] Shagun Jhaver, Amy Bruckman, and Eric Gilbert. 2019. Does transparency in moderation really matter? User behavior after content removal explanations on reddit. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–27.
- [23] Danielle Leah Kehl and Samuel Ari Kessler. 2017. Algorithms in the criminal justice system: Assessing the use of risk assessments in sentencing. (2017).
- [24] René F Kizilcec. 2016. How much information?: Effects of transparency on trust in an algorithmic interface. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, 2390–2395.
- [25] Jon Kleinberg, Himabindu Lakkaraju, Jure Leskovec, Jens Ludwig, and Sendhil Mullainathan. 2017. Human decisions and machine predictions. *The quarterly journal of economics* 133, 1 (2017), 237–293.
- [26] Amanda Kube, Sanmay Das, and Patrick J Fowler. 2019. Allocating interventions based on predicted outcomes: A case study on homelessness services. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- [27] Himabindu Lakkaraju and Osbert Bastani. 2019. "How do I fool you?": Manipulating User Trust via Misleading Black Box Explanations. *arXiv preprint arXiv:1911.06473* (2019).
- [28] John D Lee and Katrina A See. 2004. Trust in automation: Designing for appropriate reliance. *Human factors* 46, 1 (2004), 50–80.
- [29] Joa Sang Lim and Marcus O'Connor. 1995. Judgemental adjustment of initial forecasts: Its effectiveness and biases. *Journal of Behavioral Decision Making* 8, 3 (1995), 149–168.
- [30] David Madras, Toni Pitassi, and Richard Zemel. 2018. Predict responsibly: improving fairness and accuracy by learning to defer. In *Advances in Neural Information Processing Systems*. 6147–6157.
- [31] Katharina Marten, Tobias Seyfarth, Florian Auer, Edzard Wiener, Andreas Grillhösl, Silvia Obenauer, Ernst J Rummeny, and Christoph Engelke. 2004. Computer-assisted detection of pulmonary nodules: performance evaluation of an expert knowledge-based detection system in consensus reading with experienced and inexperienced chest radiologists. *European radiology* 14, 10 (2004), 1930–1938.
- [32] Paul E Meehl. 1954. Clinical versus statistical prediction: A theoretical analysis and a review of the evidence. In *In Proceedings of the 1955 Invitational Conference on Testing Problems*. University of Minnesota Press, 136–141.
- [33] Neville Moray, Toshiyuki Inagaki, and Makoto Itoh. 2000. Adaptive automation, trust, and self-confidence in fault management of time-critical tasks. *Journal of experimental psychology: Applied* 6, 1 (2000), 44.
- [34] Kathleen L Mosier, Melisa Dunbar, Lori McDonnell, Linda J Skitka, Mark Burdick, and Bonnie Rosenblatt. 1998a. Automation bias and errors: Are teams better than individuals?. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, Vol. 42. SAGE Publications Sage CA: Los Angeles, CA, 201–205.
- [35] Kathleen L. Mosier, Linda J. Skitka, Susan Heers, and Mark Burdick. 1998b. Automation Bias: Decision Making and Performance in High-Tech Cockpits. *The International Journal of Aviation Psychology* 8, 1 (1998), 47–63. DOI : http://dx.doi.org/10.1207/s15327108ijap0801_3
- [36] Daniel S. Nagin and Robert J. Sampson. 2019. The Real Gold Standard: Measuring Counterfactual Worlds That Matter Most to Social Science and Policy. *Annual Review of Criminology* 2, 1 (2019), 123–145. DOI : <http://dx.doi.org/10.1146/annurev-criminol-011518-024838>
- [37] Mahsan Nourani, Samia Kabir, Sina Mohseni, and Eric D Ragan. 2019. The Effects of Meaningful and Meaningless Explanations on Trust and Perceived System Accuracy in Intelligent Systems. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, Vol. 7. 97–105.

- [38] Allegheny County Department of Health and Human Services. n.d. Allegheny Family Screening Tool. (n.d.). <https://www.alleghenycounty.us/Human-Services/News-Events/Accomplishments/Allegheny-Family-Screening-Tool.aspx>
- [39] Emilee Rader, Kelley Cotter, and Janghee Cho. 2018. Explanations as mechanisms for supporting algorithmic transparency. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, 103.
- [40] Maithra Raghu, Katy Blumer, Greg Corrado, Jon Kleinberg, Ziad Obermeyer, and Sendhil Mullainathan. 2019. The algorithmic automation problem: Prediction, triage, and human effort. *arXiv preprint arXiv:1903.12220* (2019).
- [41] Nadine B Sarter and Beth Schroeder. 2001. Supporting decision making and action selection under time pressure and uncertainty: The case of in-flight icing. *Human factors* 43, 4 (2001), 573–583.
- [42] Jennifer L Skeem, Nicholas Scurich, and John Monahan. 2019. Impact of Risk Assessment on Judges’s Fairness in Sentencing Relatively Poor Defendants. *Virginia Public Law and Legal Theory Research Paper* 2019-02 (2019).
- [43] Linda J Skitka, Kathleen L Mosier, Mark Burdick, and Bonnie Rosenblatt. 2000. Automation bias and errors: are crews better than individuals? *The International journal of aviation psychology* 10, 1 (2000), 85–97.
- [44] CarlyWill Sloan, George Naufal, and Heather Caspers. 2018. The Effect of Risk Assessment Scores on Judicial Behavior and Defendant Outcomes. (2018).
- [45] Vernon C Smith, Adam Lange, and Daniel R Huston. 2012. Predictive modeling to forecast student outcomes and drive effective interventions in online community college courses. *Journal of Asynchronous Learning Networks* 16, 3 (2012), 51–61.
- [46] Megan Stevenson. 2018. Assessing risk assessment in action. *Minn. L. Rev.* 103 (2018), 303.
- [47] Sarah Tan, Julius Adebayo, Kori Inkpen, and Ece Kamar. 2018. Investigating Human+ Machine Complementarity for Recidivism Predictions. *arXiv preprint arXiv:1808.09123* (2018).
- [48] Niels van Berkel, Jorge Goncalves, Danula Hettiachchi, Senuri Wijenayake, Ryan M Kelly, and Vassilis Kostakos. 2019. Crowdsourcing Perceptions of Fair Predictors for Machine Learning: A Recidivism Case Study. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–21.
- [49] Michael Yeomans, Anuj Shah, Sendhil Mullainathan, and Jon Kleinberg. 2017. Making sense of recommendations. *Journal of Behavioral Decision Making* (2017).
- [50] Ming Yin, Jennifer Wortman Vaughan, and Hanna Wallach. 2019. Understanding the Effect of Accuracy on Trust in Machine Learning Models. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, 279.
- [51] Kun Yu, Shlomo Berkovsky, Dan Conway, Ronnie Taib, Jianlong Zhou, and Fang Chen. 2016. Trust and reliance based on system accuracy. In *Proceedings of the 2016 Conference on User Modeling Adaptation and Personalization*. ACM, 223–227.
- [52] Kun Yu, Shlomo Berkovsky, Ronnie Taib, Dan Conway, Jianlong Zhou, and Fang Chen. 2017. User trust dynamics: An investigation driven by differences in system performance. In *Proceedings of the 22nd International Conference on Intelligent User Interfaces*. ACM, 307–317.