

Sibyl: Understanding and Addressing the Usability Challenges of Machine Learning In High-Stakes Decision Making

Alexandra Zytek, Dongyu Liu, Rhema Vaithianathan, and Kalyan Veeramachaneni

Abstract—Machine learning (ML) is being applied to a diverse and ever-growing set of domains. In many cases, domain experts — who often have no expertise in ML or data science — are asked to use ML predictions to make high-stakes decisions. Multiple ML usability challenges can appear as result, such as lack of user trust in the model, inability to reconcile human-ML disagreement, and ethical concerns about oversimplification of complex problems to a single algorithm output. In this paper, we investigate the ML usability challenges that present in the domain of child welfare screening through a series of collaborations with child welfare screeners. Following the iterative design process between the ML scientists, visualization researchers, and domain experts (child screeners), we first identified four key ML challenges and honed in on one promising explainable ML technique to address them (local factor contributions). Then we implemented and evaluated our visual analytics tool, SIBYL, to increase the interpretability and interactivity of local factor contributions. The effectiveness of our tool is demonstrated by two formal user studies with 12 non-expert participants and 13 expert participants respectively. Valuable feedback was collected, from which we composed a list of design implications as a useful guideline for researchers who aim to develop an interpretable and interactive visualization tool for ML prediction models deployed for child welfare screeners and other similar domain experts.

Index Terms—Machine learning, XAI, Usability, child welfare, visualization

1 INTRODUCTION

Thanks to innovations in machine learning (ML), computers can now help with many tasks previously performed by humans alone, often improving both speed and precision. However, in many domains, human decision-makers provide essential insights that cannot be replaced by existing algorithms. In such cases, decision-making outcomes are improved when ML output is used to augment human decision-making, rather than replace it.

ML models often output only a single number or classification, such as the risk score seen in the upper right corner of Figure 2. This can make it difficult for human decision-makers to incorporate the model into their decision making. As a result, many ML algorithms lack *usability*, or the attribute of being able to be efficiently used by humans to make better decisions.

The machine learning, data science, and data visualization communities have offered a multitude of algorithms and tools to augment ML predictions and address these usability challenges — we refer to these as **ML augmentation tools**. These tools, when chosen carefully for the domain, have the ability to greatly improve the usability of ML models for decision making. Examples of such tools include data visualizations, global and local explanations [2], cost-benefit analysis [9], performance metrics, and information about historic usage and results of the ML model. However, research aimed at augmenting ML predictions often focuses on an audience of ML/data experts [27] [22] [8] or domain experts in more technical or data-driven fields such as medicine [15] [12]. For example, Zhang et. al. [27] developed a framework for helping data scientists and ML experts interpret and debug ML models, and Lundberg et. al. [15] developed an interface for helping anaesthesiologists prevent hypoxaemia during surgery through detailed data visualization. In contrast, many fields are more qualitative in nature, with decisions following discussion more than data crunching. In this paper, we focus on these more qualitative fields, and the usability challenges they face.

To concretely assess usability challenges, we investigated the need

- Alexandra Zytek, Dongyu Liu, and Kalyan Veeramachaneni are with MIT. E-mail: zyteka, dongyu@mit.edu, kalyan@csail.mit.edu.
- Rhema Vaithianathan is with Auckland University of Technology. E-mail: rhema.vaithianathan@aut.ac.nz.

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org. Digital Object Identifier: xx.xxx/TVCG.201x.xxxxxx

for additional auxiliary information alongside ML predictions through a comprehensive literature review. We first selected a number of papers that had ML and explainability as topics. We selected and reviewed 55 papers covering ML applications and explainability. We identified challenges in human-ML interactions described in these papers. Three of us then began codifying a set of factors that decrease usability in ML models. Table 1 summarizes a set of usability challenges we codified that are relevant when a model is actively used for decision-making.

Some challenges stem from not understanding where a model’s predictions come from, making it difficult for human decision-makers to trust the model (**TR**), and to handle any disagreements between their opinions and the model’s output (**DIS**). Others are caused by a lack of information about the real effects of a decision. A lone model prediction often does not explicitly indicate the expected results of a decision (**CON**), suggest accountability (**ACC**), or provide ethical assurances (**ETH**). Finally, challenges may arise when the output of the model is not a direct suggestion of a decision, but rather auxiliary information. In this case, the output may be confusing (**CT**) or entirely irrelevant (**UT**).

Determining which usability challenges exist, the best tools to address them, and the necessary design choices for these tools depends highly on specific aspects of the domain and the decision-makers involved. Through our literature review and the case study discussed in this paper, we identified a subset of context-dependent factors that should be considered when working to make an ML model more usable in a particular domain. Table 2 lists some examples of these factors.

To investigate the problem of finding and mitigating usability challenges in more qualitative fields, we selected the domain of child welfare screening. In terms of the relevant context factors, child welfare screeners are domain experts without ML/data science expertise, making decisions using an ML model as an auxiliary tool, with about a few minutes per decision, in a high-risk field.

Addressing usability challenges is a non-trivial task that requires collaboration with end-users. In this paper, we engaged in three forms of collaboration: **observations** to understand their existing workflow and its possible usability challenges, **interviews** to gain additional insights into the desires of end-users, and **user studies** with possible ML augmentation techniques to get concrete feedback on design.

The main research questions, and our key findings with regards to these questions, are as follows.

- RQ1** What ML usability challenges exist in the domain of child welfare screening? Through interviews and field observations, our work identifies four main challenges, described in Section 4.

Table 1. List of challenges that could negatively impact the usability of an ML model.

Usability Challenge	Code	Mitigating Tools
Lack of TR ust	TR	Global explanations, local explanations, performance metrics, historical predictions and results
Difficulty Reconciling human-ML Disagreements	DIS	Local explanations
Unclear Consequences of actions	CON	Cost-benefit analysis, historical predictions and results
Lack of AC countability or protections from accountability	ACC	Local explanations, performance metrics
Ethical Concerns (ex. possible bias, concerns about oversimplification)	ETH	Global explanations, local explanations, ML fairness metrics, historical predictions and results
Confusing or unclear prediction T arget (ie. the measure predicted by the model has an unclear meaning or significance)	CT	Cost-benefit analysis, further analysis of prediction target impact
Unhelpful prediction T arget (ie. the measure predicted by the model is not relevant to the required decision)	UT	Retrain model with new prediction target

RQ2 What interfaces can be helpful in mitigating these ML usability challenges? We designed and implemented five augmentation interfaces, making up our SIBYL tool, described in Section 5. Based on our interviews and a formal user study, described in Section 6, we suggest that local factor contributions are most useful for this domain.

RQ3 What design choices must be made when building these tools to optimize them for use by child welfare screeners and other experts in similar domains? Based on our interviews and formal user study, we identified a list of design considerations that should be made when choosing model features and developing ML augmentation visualizations, described in Sections 6 and 7.

2 STUDY CONTEXT: CHILD WELFARE SCREENING

In this section, we introduce the domain of child welfare screening

Child Abuse in the U.S. Child abuse is an active issue affecting the health and well-being of communities. The Centers for Disease Control and Prevention (CDC) estimates at least 1 in 7 children have experienced child abuse and/or neglect in the past year [1]. Child abuse victims can suffer physical and emotional injuries, and may experience trauma resulting in long-term mental health problems [1]. More than one-third of American children are investigated as potential victims of abuse or neglect by age 18 [10]. Still, in 2018, there were 1,770 reported fatalities resulting from child abuse and neglect [1].

Child Welfare Screening. In the U.S., regional Child Protective Services (CPS) agencies are tasked with handling child abuse and neglect referrals from concerned members of the community, including mandated reporters such as teachers, who are required by law to report any suspicion of abuse or neglect. These referrals are examined by child welfare specialists (“call screeners”), who decide whether to screen in or screen out each case. A screened-in case will be investigated further, while a screened-out case will be recorded but not investigated.

Both false negatives (real abuse cases that are screened out) and false positives (cases with no abuse that are screened in) can have heavy consequences. False negatives lead to prolonged child suffering and, in extreme cases, child fatality. False positives can lead to long-term emotional distress for parents, children, and other family and community members, as well as damaged financial, career and social prospects for parents and other caretakers [18].

In 2018, CPS agencies in the United States received 4.3 million referrals from concerned parties about potential child abuse [1]. 56% of these referrals were screened in and investigated, but only 16.8% of the screened in cases were found to involve abuse or neglect [1].

ML for Child Welfare. One important motivation for computerized assistance in child welfare call screening is repeated cases of missed abuse. Fatal child abuse cases in which children were referred several times but were never screened in are tragic, and although such cases are rare, they are avoidable [7]. An ML solution can quickly scan for red

flags, such as repeated referrals, that busy human call screeners may miss in the overload of data.

In recent years, predictive risk modelling (PRM) has been deployed in child welfare contexts in multiple counties, with the goal of enabling more efficient and consistent decision-making and improving the overall health and safety of county residents [24]. One example of such a model was deployed in Allegheny County, PA by Vaithianathan et. al. in 2016 [24].

Currently, PRM is being introduced to our collaborating county in Colorado by Vaithianathan et. al. [23], through a LASSO regression model trained on 461 features, which include information such as the child and parents’ age, past referrals and their outcome, and past court involvements. The model predicts the likelihood of removal from home in the next two years, translated to a 1 through 20 risk score where the higher the score, the higher the risk [24]. This paper focuses on the usage of this model.

Study Participants. We collaborated over the course of a year with a pool of 19 social workers and supervisors working for the child welfare department in a collaborating county in Colorado. All participants regularly act as screeners in the county’s child welfare screening decision making process. Our collaborations began in December 2019, with two days of in-person field observations (Section 4.1). Following this, we observed a simulated case review session via video conferencing (Section 4.2), and conducted several interviews also via video conferencing (Sections 4.3 and 5). Finally, our collaborating screeners participated in our user study digitally (Section 6).

3 RELATED WORK

In this section, we discuss related work in human-centered and explainable ML, as well as existing ML augmentation tools.

Human-Centered ML. Past literature has advocated for a *human-centered* perspective to ML [4] — one that considers machines and algorithms as part of collaborative systems alongside humans. This perspective considers how humans use, interact with, adapt to, and evaluate ML applications [4]. A truly human-centered ML approach acts end-to-end, beginning with human-in-the-loop training systems and ending with evaluation systems based on the metrics that end users are most interested in [4]. In this paper, we take a deeper look at one step of this extensive pipeline: the use of ML algorithm predictions by humans for real-world decision making.

Explainable ML. A common usability challenge addressed by the literature is the black box nature of most ML algorithms. Humans struggle to use ML predictions because they do not understand where they came from. This usability challenge is addressed through the fields of interpretable or explainable ML. Doshi-Velez and Kim proposed that the need for ML interpretability stems from an “incompleteness in the problem formulation” [2], which prevents the system from being thoroughly evaluated or trusted. This incompleteness can take several forms, including a need for scientific understanding, concerns about

Table 2. Domain context factors that may influence the usability of an ML model. The context factors relevant to child welfare screening are in bold. By *technical expertise*, we refer to the ML or data science expertise of the end-user. This is not meant to be a complete list — there are many other factors that could also be relevant.

Context Factors	Categorizations	Example domain
Degree of Automation	Fully autonomous, humans not involved in decision making	Level 5 self-driving car [17]
	Machine makes decisions, humans monitor	Level 4 self-driving car [17]
	Machine suggests decisions, humans make decisions	Content violation flagging
	Machine provides auxiliary information, humans make decisions	Child welfare screening
Decision time	Immediate	Level 5 self-driving car [17]
	Seconds	Aircraft emergency response
	Minutes	Child welfare screening
	Hours	Non-emergency medical procedures
Technical expertise of humans	Little to none	Child welfare screening
	Experience with data science	Finance
	Machine learning/Data science expert	ML model training and debugging
Domain expertise of humans	Little to none	Autonomous aircraft (to passengers)
	Basic understanding/Intuition	Crowdsourcing
	Domain expert	Child welfare screening
Associated Risk	Low	Camera roll image sorting
	Medium	Mail sorting
	High	Emergency medical procedures

safety or ethics, or mismatched objectives between the model output and the human goal [2].

Doshi-Velez and Kim [2] also define three evaluation approaches for ML interpretability. In *application-grounded approaches*, domain experts work with explanations within a real application. This provides the most realistic quantification of explanation quality, but may require high time commitments from a potentially small pool of domain experts. In *human-grounded approaches*, researchers develop simpler problems for experimentation using non-expert subjects. Finally, in *functionally-grounded evaluation*, a formal definition of interpretability is used as a proxy to evaluate an explanation without using human subjects. This paper focuses on an example of such an application-grounded approach.

Wang et. al. [25] developed a human-driven conceptual framework for building explainable AI systems. They found that decision-makers seek explanations to justify unexpected occurrences, monitor for important events, or facilitate learning. They created a taxonomy of AI techniques based on how they support human reasoning and represent information. Finally, the authors discuss how explainable AI can mitigate cognitive biases. Our work builds on this by finding cognitive biases that can be *caused* by explainable AI, as listed in Section 7.

ML Augmentation Tools. Spinner et. al. [21] developed a conceptual framework for explainable AI, along with a corresponding implementation called EXPLAINER. This framework includes single-model explainers (which are the focus of this paper) as well as multi-model explainers, which can be used to compare and select between models. The authors also distinguish between different model audiences, including novices, model users, and developers — our work focuses on model users.

Krause et. al. [11] developed PROSPECTOR, a comprehensive visualization system for data scientists. This system includes interactive functionality for both global and local feature-focused explanations.

Hohman et. al. [6] developed a visual analytics system called GAMUT to investigate how machine learning practitioners and data scientists interact with machine learning. To develop this tool for use on GAMs, the authors interviewed technical experts to generate a list of common questions asked about predictions. In total, they identified six question types, which they address using three views. GAMUT was tested by having 12 data scientists use the tool while thinking out loud, followed by an interview.

Lundberg et. al. developed PRESCIENCE, an explanatory ML system focused on preventing hypoxaemia during surgeries [15]. This tool predicts the risk of hypoxaemia in the next five minutes using a gradient-boosting algorithm trained on time series. It also includes several visualizations to explain the prediction, including SHAP feature

contribution explanations.

Kwon et. al. [12] developed RETAINVIS, a visualization tool for explaining recurrent neural networks (RNNs) applied to electronic medical records (EMRs). The tool was developed with active feedback from domain experts (medical practitioners). RETAINVIS includes five different visualizations for looking into RNNs.

Our work is similar to these tools in that it relies on collaboration with end users to develop a tool that provides additional information alongside an ML prediction. However, our users are not expected to have any prior ML or data science expertise, nor are they used to working with data-heavy visualizations. This work is also the first to our knowledge that investigates through a complete case study the usability of machine learning for child welfare screeners.

4 UNDERSTANDING CONTEXT AND END-USER NEEDS

To address **RQ1** (identifying ML usability challenges), we performed a series of field observations and interviews. In this section, we discuss our goals and the findings we made during these steps.

4.1 Understanding Existing Workflows

To better understand the existing child welfare screening workflow, we travelled to our collaborating county in Colorado to observe screeners’ decision-making on referrals without using the ML model. This process led to the following findings:

1. Our collaborating county uses the general procedure for child welfare screening shown in Figure 1. In cases of immediate concern, a referral may be screened in immediately after it is received by CPS (this decision is made by a child welfare supervisor). In most cases, however, the decision as to whether to screen-in or screen-out a referral is made by a team of child welfare experts. It is this team that receives the ML risk score prediction.
2. Five to ten minutes are spent on each case by this team. Most of this time is spent going over the details of the case. The screening decision is made after about one to two minutes of discussion.
3. A large portion of these five-to-ten minutes of screening time is dedicated to determining the factors that are associated with higher and lower likelihood of abuse — referred to as risk and protective factors, respectively — involved in a case, and weighing these against each other. The factors considered will vary based on the details of the case, but may include information such as child’s age (very young children are more vulnerable), criminal

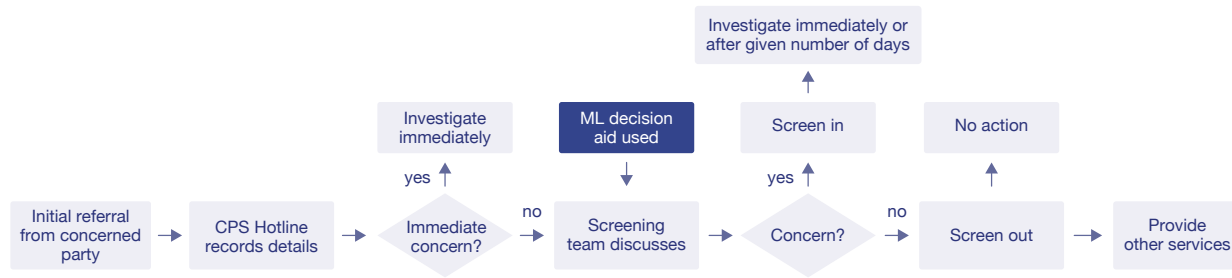


Fig. 1. The general child welfare screening process used by the CPS department of our collaborating county. The referral is first received by the CPS hotline, and then sent to a child welfare supervisor. In a minority of cases, the supervisor will deem the child or children involved to be in serious and immediate risk of danger, and will screen-in the case for immediate further investigation. In most cases, however, the case will be reviewed by a team of child welfare screeners the next day. This team will be given the ML risk score prediction (dark blue box). If this team decides to screen-in the case, it will be investigated further through home visits, interviews, or other means. Otherwise, the case will be recorded but not investigated unless re-referred. In the case of a screen-out, the screeners may elect to provide the family with additional family services.

record of adults involved, whether there are any trusted adults active in the child’s life, and actions and statements made by adults involved (for example, a mother taking actions to separate an aggressive partner from her children).

4.2 Identifying ML Usability Challenges

To identify the particular ML usability challenges relevant to child welfare screening (RQ1), we observed a simulated case review session where social workers used an ML model. In this session, our 19 expert collaborators were split into three teams of typical size, which we will refer to as T1, T2, and T3. They were asked to make decisions about real referrals that the county had handled in the past. For each case, they had access to the information that screeners are typically given during decision-making — including the age of all involved parties and a written description of the potential abuse as given by the referring party — as well as a 1 through 20 risk score provided by the ML model. After receiving this information, the screeners went ahead with their usual process: discussing the case as a team for five to ten minutes and then making a screening decision. The teams were then interviewed and asked to reflect upon how they used each ML score, whether the scores aligned with their expectations, and how the scores impacted their decisions. Each team made decisions on seven to nine cases.

During interview sessions, all three teams expressed reservations about using the ML model for decision-making. Based on their responses to our interview questions, we identified four key usability challenges:

1. **Lack of Trust TR** Screeners expressed a lack of trust when making decisions using the ML model, evidenced by their tendency to not consider the model prediction at all when it disagreed with their intuition. For example, when asked if the score caused them to reconsider their decision, T3 responded

“No, [we were] surprised it is that low.”

2. **Difficulty reconciling human-ML disagreements DIS** Screeners did not have a clear path forward when they disagreed with the model prediction, sometimes electing to ignore it entirely (see previous item) and sometimes trying to justify it based on how they thought the model worked. For example, T2 reported in one case that the score made them

“think a little deeper about why the score is so high [and caused us to] take another look at [the history]”

3. **Unclear prediction target CT** Because the model provides auxiliary information (1-20 risk score based on the likelihood of removal from home in 2 years) rather than a direct decision suggestion, there was some confusion about how to use the model prediction target. For example, when asked how the model affected their decision making process, T3 responded

“[we did not know] enough of what the score means to know how to accurately use it.”

T3 also said they

“Wish we knew how we got to the score.”

4. **Concerns about Ethics ETH** As expected for such a sensitive domain, users were concerned about the ethics of using the ML model score. There was concern that the model may prevent critical thinking. T3 commented

“[The model] could be dangerous for people just looking at the number, need to take everything into account. Makes you stop and think and ask yourself are you critically thinking...”

In addition to determining which usability challenges were relevant to child welfare screening, we also confirmed that some were *not* relevant, and therefore did not require consideration when designing SIBYL. For example, three of the example usability challenges from Table 1 were not relevant. Lack of accountability ACC was not relevant, as accountability is always held by human decision-makers in this domain. Similarly, unclear consequences of actions CON was not relevant, because the domain experts have extensive training and understanding of the potential consequences of screening decisions, and do not intend to offload this understanding to an ML algorithm. Finally, unhelpful prediction target UT was not relevant, as the ML model has been developed specifically to provide relevant information.

4.3 Interviewing Screeners

To begin addressing RQ2 (what tools can be helpful in mitigating usability challenges), we interviewed the 19 screeners about what additional information they would be interested in receiving alongside ML predictions. The format was a semi-structured open-floor session. We began by asking the screeners whether they thought additional information would be useful, and what specific information they would be interested in. We also proposed possible augmentation information (e.g. the relative importance of different factors, answers to what-if questions, and comparisons with past cases) and asked if they might be helpful.

Our findings from this interview included:

1. Screeners were confident that they would want to know why the model made the predictions it made.
2. Screeners believed that understanding how important each factor was to the score prediction would be helpful.
3. Screeners wanted to know what steps the model takes in making predictions.
4. Screeners were interested in getting “what-if” style explanations that give information about what could be changed about a child to reduce his or her risk. Note that the ML model is not trained on causal relationships, so explanations would not be able to provide such information.

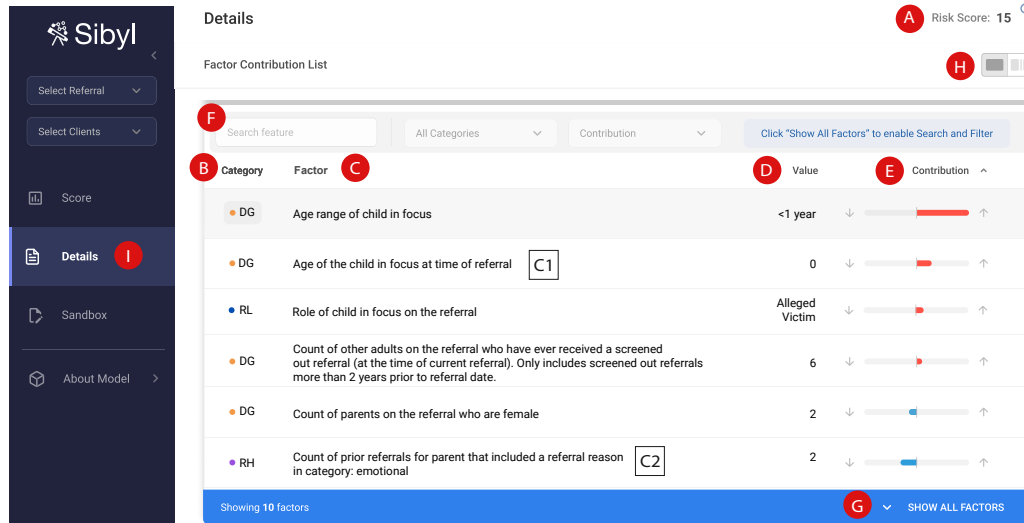


Fig. 2. The final version of the “Case-Specific Details” interface after factoring in user study feedback. This interface shows how particular factors contribute to predictions made by ML models about child welfare. Labeled elements are as follows: (A) The risk score for the case (1-20). (B) Categories for each factor, such as demographics (DG) or referral history. (C) A short description of each factor. (D) The value of numeric or categorical factors. (E) The contribution of each factor (the table can be sorted in ascending or descending order of contribution). (F) UI components for searching by factor name or filtering by category, enabled when the full factor list is shown. (G) A button for switching between a view that shows only the top 10 most contributing factors and one that shows all factors. (H) A button for switching between a single-table view and a side-by-side view, which splits factors that increase and decrease risk. (I) A sidebar for switching between different explanation types, as described Section 5.

- Some screeners were interested in seeing similar cases they dealt with in the past. Others thought this would be too much information to digest in such a short period of time.

5 GETTING FEEDBACK ON POSSIBLE TOOLS

To address **RQ2** (identifying helpful interfaces) and **RQ3** (identifying important design choices), and based on the usability challenges identified (Section 4.2) and responses to our interview (Section 4.3), we engaged in a user-centered iterative design process [16] to develop SIBYL, an ML augmentation tool. We began by designing high-fidelity mock-ups for five augmentation interfaces, each with a separate purpose and goal. Table 3 summarizes the motivation behind each interface in terms of its theorized effect on addressing the usability challenges. The full versions of the original high fidelity mock-ups can be found in Appendix B.

Early in the design process, we learned that the word “factor” is more familiar to screeners than “feature” when referring to pieces of information used when making decisions. For the purposes of consistency, we use the word factor throughout this paper when referring to data inputs used by the model.

5.1 Case-Specific Details: Factor Contributions

The *Case-Specific Details* interface (Figure 2) provides a simple local explanation of where an individual model prediction comes from through factor contributions. The table assigns each factor a contribution (Figure 2E), either positive (red) or negative (blue). The bar length indicates the magnitude of the contribution.

The phrases *positive contribution* and *negative contribution* may cause confusion. In the ML community, a positive contribution indicates that the model prediction will increase in value. In the screener community, however, an increased risk score is a negative occurrence. We decided to avoid using the terms “positive” and “negative” in the app or instructions. Instead, the factors are labelled as “risk factors” or “protective factors” to mirror the screeners’ language. We utilize the bar’s direction along with the two arrows (↑ and ↓) to suggest risk increasing (red bar pointing right) or decreasing (blue bar pointing left).

A local factor contribution explanation may reveal that the young age of a particular child (infant) has caused a significant increase in the risk score (Figure 2-C1), while the low number of past referrals (2) compared to the average referred child resulted in a decrease compared to the average risk (Figure 2-C2).

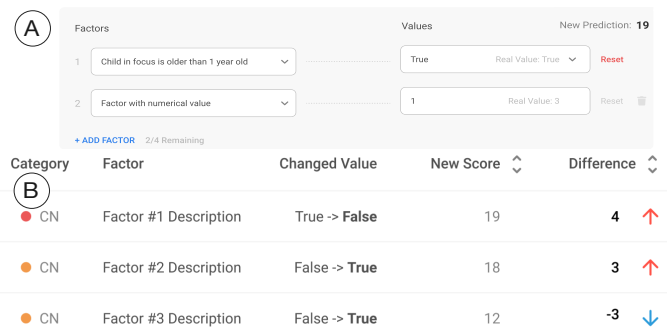


Fig. 3. Sandbox visualizations. (A) Users can change up to four factor values at a time, and the new score will be displayed on the top right corner. (B) The table lists the resulting prediction if each Boolean value is individually reversed.

The local factor contributions were found using the Shapely Additive Explanations (SHAP) algorithm [14]. We chose to use SHAP because it is theoretically grounded in game theory and generates consistent and intuitive explanations for an ML model.

We decided on a local contribution interface as we theorized it may help with all identified usability challenges: the screeners’ lack of trust **TR** by demonstrating that the model relies in part on similar factors as the human screeners in making decisions, difficulty reconciling disagreements **DIS** by highlighting differences in the human and model’s logic, unclear prediction target **CT** by providing a concrete explanation of the scores’ meaning, and concerns about ethics **ETH** by making critical thinking about relevant factors easier.

Feedback. Initial interview feedback suggested that screeners were most interested in the Case-Specific Details interface, so it was kept as the default option. Additionally, screeners said that in their usual workflow, they would list “Risk” and “Protective” factors side by side. To mirror this, the updated version of this interface has a split-view toggle (Figure 2H) that shows negatively and positively contributing factors in two side-by-side tables).

5.2 Sandbox: Investigating “What Ifs”

The *Sandbox* interface allows users to experiment with and see how the model prediction would change if factors differed. It has two parts: (1) The *Experiment with Changes* box (Figure 3A) allows users to

Table 3. The proposed SIBYL interfaces (left column), the challenges they were theorized to address (middle column, using the codes from Table 1), and the reasons we expected these interfaces to address the given challenges (right column). **TR**: Lack of trust in the model. **DIS**: Difficulty reconciling disagreements. **CT**: Confusing prediction target. **ETH**: Ethical concerns.

Interface	Challenges Addressed	How does it address the challenge?
	TR	
Case-Specific Details	TR DIS CT ETH	Reveals relevancy of considered factors Highlights factors that may have been missed or misused Translates score to a concrete factor list Allows for critical thinking about factors and score
Sandbox	DIS ETH	Allows users to test theorized justifications Allows for thinking through what-if scenarios
Similar Cases	TR ETH	Provides information on past performance Provides a deeper look into the nuance of cases
Global Factor Importance	TR CT	Reveals how the model generally makes predictions Translates the score to a concrete factor list
Factor Distributions	TR CT	Shows how well the model performed on past cases Shows the relationship between the risk score and removals from home

change up to four factor values at a time, to investigate specific “what-if” questions. (2) The *Model predictions if each value was changed* box (Figure 3B) shows the resulting prediction if each Boolean factor value was individually reversed (ie. true to false or false to true). Following the consistent design principle, we adopted a similar tabular design and use the red/blue up/down arrows to highlight changes in risk scores.

This interface was added based on feedback from the interviews described in Section 4.3. We theorized it may help with difficulty reconciling disagreements **DIS** by allowing screeners to test their theorized justifications, and with concerns about ethics **ETH** by making more detailed consideration about the model’s output easier.

Feedback. Screeners expressed concern that the Sandbox interface may be misconstrued as suggesting specific actions or reflecting real-world causal structures. However, screeners also said that they saw value in this interface as a supervision tool, used to review model predictions and human decisions rather than being actively used during the decision-making process.

5.3 Similar Cases: Investigating Past Cases

The *Similar Cases* interface shows the complete history of child welfare involvement with past cases that had similar factor values. The similar cases are found using a Nearest Neighbors algorithm. For design purposes, the algorithm used weights all factors equally. This interface includes a timeline for the current case and each similar case, and highlights events such as referrals to child welfare services, investigations, and removals. To facilitate comparative analysis, cases are lined up row by row and share the same timeline.

The interface was added as we theorized it may help with screeners’ lack of trust **TR** by demonstrating past performance, and concerns about ethics **ETH** by providing a deeper look into an individual case.

Feedback. Screeners were concerned that this interface may cause poor decision-making. They explained that basing decisions about a current case on past cases that seem similar is discouraged, as this can lead to biases or self-fulfilling prophecies. Therefore, it was decided that this interface would not be included in a decision-making tool. However, county officials pointed out that it could be used retroactively (outside of decision-making) to investigate unusual predictions made by the model for the purposes of model evaluation. Due to these concerns, we decided not to include this interface in our formal user study.

5.4 Global Factor Importances: Understanding the Model

The *About Model* interfaces offer information about the model’s general logic, outside of the context of a individual prediction.

The first *About Model* interface is the *Global Factor Importance* explanation. This interface shows a global explanation in the form of the general, relative importance of each factor. It also provides a brief description of the model architecture and logic, as well as its performance metrics. A visualization for this interface can be found in the Appendix, Figure 9

The global factor importance rankings were found using the Permutation Importance algorithm [3]. This algorithm computes the change in model performance if each factor is permuted individually. It therefore describes how closely each factor is linked to model performance.

This interface was added as we theorized it may help screeners build trust in the model **TR** by seeing how it generally makes predictions, and because it may clarify the meaning of the prediction target **CT**.

Feedback. Screeners said that the Factor Importance interface seemed intuitive, but may provide too much information to be practical during active decision-making. Instead, they said it may be useful for training and education.

5.5 Factor Distributions: Understanding Past Predictions

The second *About Model* interface is the **Factor Distributions** explanation, which gives a quick retrospective view of how the model performed in the past. This interface shows the distribution of factor values among past cases that were given a particular score (Figure 4A), as well as the percentage of children with that score who were removed from the home (Figure 4B).

Depending on the factor type (Figure 4C), the Factor Distributions explanation uses one of three visualizations to show the value distribution of children who received the selected prediction score. For binary factors, a progress-bar like design is used. For numeric factors, a box-and-whiskers plot shows the global minimum and maximum values for a feature, and the minimum, first quartile, third quartile and maximum for the selected risk score. For categorical factors, segmented bars are used to encode a categorical value distribution; hovering over a segment provides more information about the categories and their corresponding percentages.

This interface was added as we theorized it may help screeners build trust in the model **TR** by seeing how it generally performs, and it may clarify the value of the prediction target **CT** by showing how it relates to a more tangible output of removals from the home.

Feedback. Like the Factor Importance interface, screeners expressed concern that the Factor Distributions interface shows too much information for use during active decision-making. However, they said it may be useful for training and finding gaps in provided services.

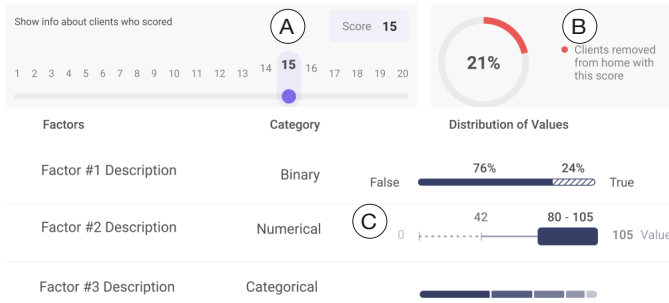


Fig. 4. Factor Distribution visualizations. Users specify a risk score of interest (A) and observe the percentage of children with that score who were removed from the home (B). Three kinds of visualizations (C) are proposed to show the value distribution of children with the selected prediction score.

6 USER STUDIES

To evaluate SIBYL, we ran two formal user studies.

Our first study involved 12 data and/or social scientists. We chose to run a study with non-experts first in order to fix immediate usability problems and iterate on the UI/UX elements of the tool. Although these participants had no prior experience in child welfare screening, it was possible for them to understand the process intuitively enough to use the tool in roughly the same way as experts would.

In our second study, we engaged 13 collaborating child welfare screeners (experts), 2 of whom completed the task while video conferencing and screen-sharing with us. We discuss the results of both user studies in this section.

Data used: For privacy reasons, we used only synthetically generated and deidentified data in this section. Data for different factors was generated using the CTGAN synthetic data generation algorithm [26], in which a generative model is learned from the real data and samples that resemble the real data are drawn from the model.

Case descriptions were real paragraph-form narratives provided by concerned parties during past referrals. Names were changed by a representative of the county for de-identification. For example:

“Caller (teacher) says Abby (age 5) came into school with bruise on arm. Caller says Abby often comes in bruised. Abby told teacher she fell off bike. Teacher asked Abby’s mother about this and mother started acting aggressive...”

6.1 Study Procedure

Participants were first shown a short video explaining how to use SIBYL. Next, they were shown 7 case descriptions, accompanied by model predictions and SIBYL interfaces with simulated data.

Participants were then asked to make a screen-in/screen-out decision, and to answer some reflection questions. These questions included 1) five-point Likert-scale style questions about how much participants trusted the model and how confident they felt in their decisions, 2) multiple choice questions about which SIBYL interfaces were helpful, and 3) free response questions about trust in the model and general feedback.

In total, experts completed 73 individual case analyses, and non-experts completed 75. The procedure for this user study is summarized in Figure 5.

6.2 Study Results

6.2.1 Helpful Interfaces

To address RQ2, we analyzed the self-reported helpfulness of each augmentation interface.

The Case-Specific Details interface was by a large margin considered the most helpful interface, by both experts and non-experts. It was labelled as being helpful by experts in 91.8% of case analyses, and by non-experts in 90.7% of case analyses. This was significantly higher than Sandbox (experts: 16.4%, non-experts: 22.6%) and Factor Distributions (experts: 20.5%, non-experts: 8.0%). Factor Importance was never listed as helpful by either group.

6.2.2 Reliance on Sibyl

Unsurprisingly, non-experts were more likely to report listening to the model without considering the added information in SIBYL. One non-expert participant commented,

“No idea what is going on in this case description – so completely defer to the model here.”

Another non-expert participant commented

“I found the score useful - and used it as a justification for screening out without exploring in detail all the factors.”

Additionally, non-experts reported that they used the model “a lot” or “a great deal” on 46.4% of cases, while experts chose these options in 15.6% of cases.

6.2.3 Impact on Trust

The SIBYL interfaces were reported to both increase and decrease expert users’ trust in the model for different reasons. To analyse why this might occur, we thematically analyzed the responses to the open question “What made you trust the model more or less?” We gathered 56 responses to this question and divided the answers into two categories, based on the corresponding answer to the 5-point Likert scale question “How much did you trust the model’s prediction for this case?” 43 free-responses corresponded to trusting the model “a great deal,” “a lot,” or “a moderate amount,” while 13 corresponded to trusting the model “a little” or “not at all.” For all but three of these answers, the response corresponded with the degree of trust listed (ie., participants who reported trusting the model provided reasons why they trusted the model more). We do not include the three exceptions in our analysis.

To conduct the thematic analysis, we sorted according to which (if any) specific SIBYL interfaces were referenced, and further identified what specific elements of the interfaces were referenced. For answers that did not reference a specific interface, we identified themes in the responses, such as agreement with the model, references to general information provided by the model, or inaccuracies in the model. With these two coding mechanisms, we sorted the responses into 7 themes that corresponded with increasing trust, and 5 that corresponded with decreasing trust.

Table 4 lists these themes. We see that agreeing with the model’s score increases trust of the model the most. Beyond this, the Case-Specific Details page was frequently cited as increasing model trust, either due to specific factors listed or more general elements of the page, such as the number of factors. Trust was reduced when there was confusion or inconsistencies in the presented information, or when the model did not consider important factors that participants knew about.

6.2.4 Information Presentation

To address RQ3, we categorize and summarize the comments made by users regarding SIBYL design choices, as well as the steps we took to address them.

1. **Too many factors shown** The model was originally trained on over 400 factors, all of which were presented in the SIBYL interface, but many of these factors have zero or near-zero weight. For example, one participant commented:

“Too many factors listed. I only want to see the material risk and protective factors.”

Our updated version of SIBYL only shows 10 factors by default, with an option to show more.

2. **Confusion caused by correlated factors** The model uses some engineered factors, resulting in factors that have deterministic relationships. For example, there is a numeric factor called AGE OF CHILD, and then a set of binary factors referring to each age group: i.e. CHILD IS LESS THAN 1 YEAR OLD, CHILD IS BETWEEN THE AGES OF 1 AND 3, etc. These factors may cause confusion when shown directly to users. In addition to increasing the cognitive load on users without providing additional useful information, explanations using these factors may reveal seemingly contradictory or unusual relationships. For example,

an age category may contribute greatly, while the numeric age factor does not.

Additionally, having these correlated factors causes confusion on the sandbox page, as it is possible to change one factor without changing all of the other deterministically-correlated factors in its set. One participant commented,

“I’m not sure, in the sandbox, if I change one feature, other features will be changed automatically.”

To solve these problems, we combined the correlated factors in the SIBYL interface, forming categorical factors out of binary one-hot encoded factors, and summing the additive contributions.

- 3. Confusion caused by Boolean terminology** One source of confusion was the method of displaying Boolean factors. In our original design, we displayed the description of the factor, with a value of `True` or `False`. This is the most accurate way of representing the model’s logic, but it is not the most intuitive way for our end-users. One participant said,

“The ‘true’ and ‘false’ is hard to interpret... Would rather have a positive statement (e.g., no perpetrator named)”

Therefore, our final version of SIBYL instead states only true statements about the child — including by negating descriptions of false factors. For example, the factor `CHILD HAS SIBLINGS` with a value of `False` will be displayed as `CHILD DOES NOT HAVE SIBLINGS`.

7 DISCUSSION AND LIMITATIONS

In this section, we summarize general lessons learned regarding how augmentation tools can improve the usability of ML models in the child welfare domain. We speculate that these lessons may generalize to other domains with similar context factors — i.e. those where users with high domain expertise and low technical expertise make high-impact/high-risk decisions with the help of ML models.

7.1 Helpful Augmentation Tools

In Section 6.2.1, we noted that our user study participants did not frequently list our About Model or Sandbox interfaces as helpful, despite our hypothesis that these interfaces could address some usability challenges as noted in Table 3. While further user studies will be required to identify exactly why more interfaces were not used, a few potential reasons come to mind. First, screeners are used to making decisions fairly quickly, and may not have enough time to parse and consider case details while investigating more than one interface. Second, the About Model interfaces may not be very relevant to screeners in the midst of a specific case, and may be more helpful if presented to screeners prior to screening. Third, the experimentation functionality of the Sandbox page is only useful if screeners have a specific what-if question in mind, while the “flipped feature” aspect may be somewhat redundant — it shows similar information to the Case-Specific Details page (the relative importance of each factor), but in a way that appears to be less intuitive to users.

7.2 Accuracy versus Fidelity

Robnik-Sikonja and Bohanec [19] define the *accuracy* of an explanation as how well it generalizes to other unseen examples (i.e., how accurately these rules predict what happens in the real world), and *fidelity* as how well an explanation describes the model itself.

Our users were mostly interested in getting *accurate* explanations that provided information about *the case at hand*. As evidenced by all three findings about the interface design (Section 6.2.4), users wanted to receive information about the model in a language and format that mirrored their own, not the format used by the model itself. This is also evidenced by design requests like using the terms *risk and protective factors*, rather than the more ML-centric terms *negative and positive features*.

7.3 The Importance of Interpretable Factors

Simple models, such as regression, are often cited as being inherently interpretable [20]. However, this case study suggested that even simple models may cause confusion in users, and lead to challenges when attempting to explain model predictions for decision making.

Instead, our work found that, for the purposes of making models usable for end-users, the interpretability of the model factors may be most important. In our study, the screeners were often confused when explanations used factors that did not have clear implications on risk. For example, in our user study, one participant said

“... 2 parents have missing date-of-birth is shown as a significant blue bar which I can’t imagine is protective.”

Additionally, as discussed in Section 6.2.4, one-hot-encoded factors were not interpretable, and many of the reasons screeners trusted the model more or less (Table 4) related to the specific factors.

7.4 Non-Applicable Usability Challenges

One interesting finding of our work was the usability challenges we did not see evidence of in this case.

For example, one possible use of explanations is to give humans the ability to actively correct errors in a model’s logic. We did not see evidence of this behavior from our users, however. There are several possible reasons for this. First, our users are making decisions in a very limited time, and do not have additional time to review the model’s quality. Second, our users are thoroughly analyzing every case on their own and were only using the model as an extra flag. Finally, the users already have some discomfort about the model, likely due to the high associated risk. As a result, our users tended to discount the model altogether if they did not believe it was correct about a particular case.

Additionally, users expressed almost no interest in learning about the model itself through explanations. A common explanation need addressed by the literature [13] is to understand how the model works (model transparency), possibly for debugging. In our case study, however, only once (see Section 4.3, Item 3) did any screener express interest in understanding the details of how the model worked under the hood — and even then, they were mostly looking for a broad overview.

7.5 Cognitive Biases

Wang et. al. [25] introduced a list of the cognitive biases explanations can help address. Our experience with child welfare screeners addition-

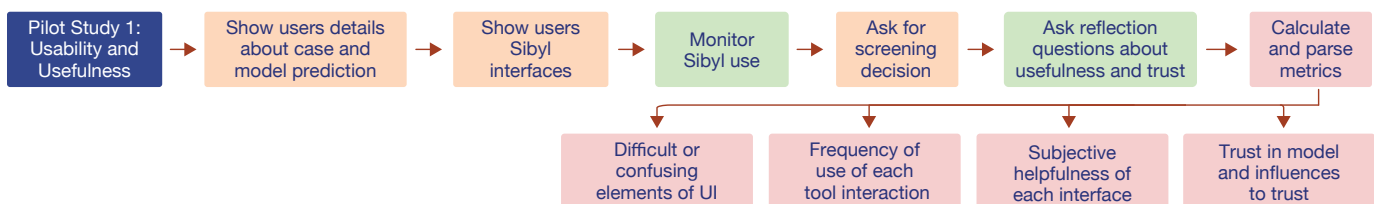


Fig. 5. The procedure for our formal user study. Our participants were first shown the description of potential abuse from a child welfare referral, as well as the corresponding ML prediction risk score. Next, they were given the opportunity to interact with the SIBYL interfaces. Once they were ready, they were asked to make a screen-in or screen-out decision, and then asked a series of reflection questions.

Table 4. Summary of answers to the question “what made you trust the model more/less”. The top section lists reasons for having “a great deal”, “a lot” or “a moderate amount” of trust in the model. The bottom section lists reasons for having “a little” or “not at all” trust in the model. The first column lists the general themes we found in the answers. The second column lists the number of answers that fell within each theme. The final column lists selected answers for each theme.

Factors that increased trust		
Category	# of answers	Sample answers
General comments about being shown protective and risk factors (Details page)	8	“The details and the risk and protective factors and the contribution they have” “Info in the details and risk factors”
Specific factors listed as risk or protective (Details page)	8	“The past number of child welfare involvements (listed in the features listed)”, “The risk factors involved, especially prior placements, benefits, and current CPS involvement”
The number of factors listed (Details page)	5	“Very few risk factors”, “The lack of protective factors”
The score agreed with screener intuition	10	“...not residing with the alleged perpetrators which I would assume would reduce the risk score”, “Model prediction makes sense”
The explanation agreed with screener intuition	3	“Risk factors made sense for the model prediction number”
General comments about sandbox page	1	“Details under sandbox of why the risk level was so high”
General comments about the explanation providing more information or understanding	6	“Allowed for more understanding”, “History clarification”
Factors that decreased trust		
Category	# of answers	Sample answers
A specific, key factor was not considered by the model	3	“This is a case for law enforcement, not CPS”, “...it may have been handled during the open case”
The importance weighting of factors was off (Details page)	2	“Young, vulnerable children being left alone is still cause for concern, despite past involvement”
The score disagreed with screener intuition	3	“Seems high, no health history, no real proof of any drug use, no proof child is at any risk of abuse”
There was some confusion about information presented	2	“There were discrepancies between the info in the referral and the info provided by the tool”
The screener wanted more information	2	“I would want to see other referrals for the family”

ally suggested some of these cognitive biases could be *encouraged* by the explanations and other forms of further information. For example:

1. **Representativeness Bias [25]:** Case-based explanations that offer similar examples to the case at hand (such as our Similar Cases page) risk encouraging users to make decisions based on similarities to another case.
2. **Causation vs Correlation:** Counterfactual-based explanations, which consider how the model prediction would change under different circumstances, made participants more likely to interpret the explanations as containing information about the causal structure of the world.
3. **Availability Bias [25]:** A factor-contribution explanation that is sorted in ascending order (and therefore lists negative contributions first) may result in different decisions than one that is sorted in descending order (and therefore lists positive contributions first) due to availability bias, which causes humans to put too much importance on recent or memorable events or information.

Further work and user studies may better reveal the extent to which these biases are caused or exacerbated by ML augmentation tools.

7.6 Limitations

Our work has some limitations that may be addressed in future work.

Our analysis focused on the qualitative comments and self-reported confidence measures provided by participants, which we used to measure the usability and perceived usefulness of the different explanation interfaces. Because the data used in our formal user study was synthetic, and therefore not associated with real-world outcomes, we did not measure the efficacy of the screening decisions made by these users. Arguably, it would be possible to quantitatively measure certain effects of SIBYL – such as changes in the quality of decisions made, or differences in results depending on how the platform was deployed – or even

to use SIBYL to identify biases in human decision-making. However, this was beyond the scope of this paper, which sought to explicitly augment human decision-making, rather than potentially cross the line into partially automating it.

The literature on ML usability would benefit from a complete, methodological investigation into the ML usability challenges present across domains, to extend our sample subset introduced in Section 1.

8 CONCLUSION AND FUTURE WORK

In this work, we identified the ML usability challenges associated with the domain of child welfare screening. We found one promising tool (factor contributions, on the Case-Specific Details interface) for mitigating many of these challenges, and pinpointed important design decisions that must be made to maximize the usefulness of this tool. Future work should empirically investigate the effect this tool has on decision-making, quantitatively measure how well it mitigates existing usability challenges, and explore innovative visualization designs to better solve the remaining challenges.

ACKNOWLEDGMENTS

We would like to thank Rhema Vaithianathan, Diana Benavides Prado, Megh Mayur, Athena Ning, and Larissa Lorimer for their guidance throughout the process of applying machine learning explainability to child welfare, for connecting us to the child welfare domain experts, and for providing us with their trained model and dataset. We thank the child welfare experts at Larimer County Department of Human Services for their invaluable insights and feedback on the tool. We thank Iulia Ionescu, Sergiu Ojoc, Ionut Radu, and Ionut Margarint for their work on developing the Sibyl application. We thank Arash Akhgari for his work on our diagrams and visualizations. We thank Michaela Henry for support in project management and insights. We thank the participants of both user studies for their time and feedback. Finally, we would like to thank our anonymous reviewers for their comments and suggestions. This work is supported in part by NSF Award 1761812.

REFERENCES

- [1] Children's Bureau. Child Maltreatment 2018: Summary of Key Findings. 2020.
- [2] F. Doshi-Velez and B. Kim. Towards A Rigorous Science of Interpretable Machine Learning. *arXiv:1702.08608 [cs, stat]*, Mar. 2017.
- [3] A. Fisher, C. Rudin, and F. Dominici. All Models are Wrong, but Many are Useful: Learning a Variable's Importance by Studying an Entire Class of Prediction Models Simultaneously. *Journal of Machine Learning Research*, 20:1–81, Dec. 2019.
- [4] M. Gillies, R. Fiebrink, A. Tanaka, J. Garcia, F. Bevilacqua, A. Heloir, F. Nunnari, W. Mackay, S. Amershi, B. Lee, N. d'Alessandro, J. Tilmanne, T. Kulesza, and B. Caramiaux. Human-Centred Machine Learning. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems*, pp. 3558–3565. ACM, San Jose California USA, May 2016. doi: 10.1145/2851581.2856492
- [5] Google. Machine Learning Glossary: Fairness. *Google Developers*, 2020.
- [6] F. Hohman, A. Head, R. Caruana, R. DeLine, and S. M. Drucker. Gamut: A Design Probe to Understand How Data Scientists Understand Machine Learning Models. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems - CHI '19*, pp. 1–13. ACM Press, Glasgow, Scotland Uk, 2019. doi: 10.1145/3290605.3300809
- [7] D. Hurley. Can an Algorithm Tell When Kids Are in Danger? *New York Times*, p. 30, Jan. 2018.
- [8] M. Kahng, P. Y. Andrews, A. Kalro, and D. H. Chau. ActiVis: Visual Exploration of Industry-Scale Deep Neural Network Models. In *IEEE Conference on Visual Analytics Science and Technology (VAST)*, vol. 24, p. 10. IEEE, Phoenix, Arizona, USA, Apr. 2017.
- [9] W. Kenton. How Cost-Benefit Analysis Process Is Performed. *Investopedia*, 2021.
- [10] H. Kim, C. Wildeman, M. Jonson-Reid, and B. Drake. Lifetime Prevalence of Investigating Child Maltreatment Among US Children. *American Journal of Public Health*, 107(2):274–280, Feb. 2017. doi: 10.2105/AJPH.2016.303545
- [11] J. Krause, A. Perer, and K. Ng. Interacting with Predictions: Visual Inspection of Black-box Machine Learning Models. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, pp. 5686–5697. ACM, San Jose California USA, May 2016. doi: 10.1145/2858036.2858529
- [12] B. C. Kwon, M.-J. Choi, J. T. Kim, E. Choi, Y. B. Kim, S. Kwon, J. Sun, and J. Choo. RetainVis: Visual Analytics with Interpretable and Interactive Recurrent Neural Networks on Electronic Medical Records. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):299–309, Jan. 2019. doi: 10.1109/TVCG.2018.2865027
- [13] Z. C. Lipton. The Mythos of Model Interpretability. In *2016 ICML Workshop on Human Interpretability in Machine Learning*, vol. 16, pp. 31–57. IMLS, New York, NY, June 2016.
- [14] S. M. Lundberg and S.-I. Lee. A Unified Approach to Interpreting Model Predictions. In *Advances in Neural Information Processing Systems*, vol. 31, p. 10. Curran Associates Inc., Long Beach, California, 2017.
- [15] S. M. Lundberg, B. Nair, M. S. Vavilala, M. Horibe, M. J. Eisses, T. Adams, D. E. Liston, D. K.-W. Low, S.-F. Newman, J. Kim, and S.-I. Lee. Explainable machine-learning predictions for the prevention of hypoxaemia during surgery. *Nature Biomedical Engineering*, 2(10):749–760, Oct. 2018. doi: 10.1038/s41551-018-0304-0
- [16] T. Munzner. A Nested Model for Visualization Design and Validation. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):921–928, Nov. 2009. doi: 10.1109/TVCG.2009.111
- [17] National Highway Traffic Safety Administration (NHTSA). Automated Vehicles for Safety. *NHTSA*, Sept. 2017.
- [18] D. Richardson. The Effects of a False Allegation of Child Sexual Abuse on an Intact Middle Class Family. *IPT*, 2:226–238, 1990.
- [19] M. Robnik-Šikonja and M. Bohanec. Perturbation-Based Explanations of Prediction Models. In J. Zhou and F. Chen, eds., *Human and Machine Learning: Visible, Explainable, Trustworthy and Transparent*, Human-Computer Interaction Series, pp. 159–175. Springer International Publishing, Cham, 2018. doi: 10.1007/978-3-319-90403-0_9
- [20] C. Rudin. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5):206–215, May 2019. doi: 10.1038/s42256-019-0048-x
- [21] T. Spinner, U. Schlegel, H. Schäfer, and M. El-Assady. explAIner: A Visual Analytics Framework for Interactive and Explainable Machine Learning. *IEEE Transactions on Visualization and Computer Graphics*, 26(1):1064–1074, Jan. 2020. doi: 10.1109/TVCG.2019.2934629
- [22] H. Strobelt, S. Gehrmann, H. Pfister, and A. M. Rush. LSTMVis: A Tool for Visual Analysis of Hidden State Dynamics in Recurrent Neural Networks. In *InfoVis*, vol. 24, pp. 667–676. IEEE, Phoenix, Arizona, USA, Oct. 2017.
- [23] R. Vaithianathan, H. Dinh, A. Kalisher, C. Kithulgodla, E. Kulick, M. Mayur, A. Ning, D. Benavides Prado, and E. Putnam-Hornstein. Implementing a Child Welfare Decision Aide in Douglas County. Dec. 2019.
- [24] R. Vaithianathan, E. Putnam-Hornstein, N. Jiang, P. Nand, and T. Maloney. Developing Predictive Models to Support Child Maltreatment Hotline Screening Decisions: Allegheny County Methodology and Implementation. p. 60, 2017.
- [25] D. Wang, Q. Yang, A. Abdul, and B. Y. Lim. Designing Theory-Driven User-Centric Explainable AI. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems - CHI '19*, pp. 1–15. ACM Press, Glasgow, Scotland Uk, 2019. doi: 10.1145/3290605.3300831
- [26] L. Xu, M. Skoularidou, A. Cuesta-Infante, and K. Veeramachaneni. Modeling tabular data using conditional GAN. In *Advances in Neural Information Processing Systems*, 2019.
- [27] J. Zhang, Y. Wang, P. Molino, L. Li, and D. S. Ebert. Manifold: A Model-Agnostic Framework for Interpretation and Diagnosis of Machine Learning Models. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):364–373, Jan. 2019. doi: 10.1109/TVCG.2018.2864499

A APPENDIX: MITIGATING TOOL DEFINITIONS

Here, we define the types of mitigating tools discussed in this paper and Table 1.

Global Explanation: An explanation of a model's general logic, achieved through methods such as quantifying the overall importance of features or visualizing the model boundary [2].

Local Explanation: An explanation as to why a model made an individual prediction, achieved through methods such as quantifying how much each feature contributed to this particular prediction [2].

Cost-Benefit Analysis: A measurement of the expected total reward from taking an action, defined by the expected benefits minus the costs [9]. In the case of machine learning, this would involve providing information about the expected results of a prediction alongside the prediction itself.

ML Fairness Metrics: Mathematical approaches to measuring the level of bias present in models [5]

B APPENDIX: DESIGN MOCKUP REVIEWS

Figures 6 - 10 show the original design mockups that were presented to child welfare screeners, and describe the feedback and changes that were made as a result of this interview.

C APPENDIX: USER STUDY QUESTIONS ASKED

Table 5 contains the complete list of questions we asked during our user study.