Learning Optimal Predictive Checklists

Anonymous Author(s)
Affiliation
Address
email

Abstract

Checklists are widely-used clinical decision aids that improve safety and reliability. In this paper, we present a new method to learn predictive checklists for classification tasks. We represent checklists as discrete linear classifiers with unit weights and binary features, and fit checklists that optimize accuracy by solving an integer programming (IP) problem. Our approach includes functionality for adaptive binarization, allows practitioners to fit models that obey custom constraints on model form and performance, and provides an optimality gap that can inform model development and customization. We introduce specialized techniques for initialization and training that reduce the time for our method to fit a checklist that performs well, and to pair it with a small optimality gap. We benchmark the performance of our method on eight clinical prediction problems, and illustrate its potential for customization by training a short-form checklist for post-traumatic stress disorder screening.

1 Introduction

Checklists are widely used to promote safety and reliability in human decision-making, particularly in healthcare [11, 16, 20, 34, 44, 51, 59, 60, 61, 62, 63]. In surgical safety, checklists have led to large reductions in mortality and morbidity [see e.g., 34, 59], and checklists are the diagnostic standard in mental health, e.g., for PTSD [43] and schizophrenia [4]. Checklists are easy to use, easy to understand [33, 55], and they allow clinicians to scrutinize models and to make an informed decision regarding adoption. Unlike other kinds of predictive models, checklists can be integrated into existing clinical workflows without extensive training or technology, e.g., as a printed sheet [53]. Medical checklists are typically developed by panels of experts [30, 40], but there are no standardized approaches for checklist design, development or validation.

A key challenge in checklist development is the lack of reliable methods to learn them from clinical datasets [33]. Given the simplicity of our model class, it is entirely possible that some datasets may never admit a checklist that is sufficiently accurate to deploy in clinical settings. Finally, even checklists with high accuracy may perform poorly in a minority population – models for clinical decision-support often exhibit disparities in prediction and performance across subgroups [56, 67].

In this paper, we introduce a checklist learning algorithm that addresses these challenges. Our method identifies the most accurate checklist that obeys all requirements needed for adoption; as well as identifying datasets that cannot admit a checklist with sufficient accuracy for the clinical classification problem. Our method incorporates a mixed integer programming (MIP) problem that: (1) directly measures and constrains the quantities of interest to optimize accuracy; (2) automatically binarizes features to further enhance performance. We then fit checklists by solving this problem as well as we can. While other work has targeted interpretable machine learning more generally [2, 27, 42], this is generally via fitting simple rule-based models by solving harder optimization problems. In contrast to other work, our program is considerably easier to solve from a computational standpoint because it

Submitted to 35th Conference on Neural Information Processing Systems (NeurIPS 2021). Do not distribute.

restricts coefficients to binary values and does not include ℓ_0-regularization. Here, we make use of this "slack" in computation to add functionality that improves the performance of checklists.

We conduct a broad empirical study of trained checklist performance, generalization, optimality gap, and fairness in eight healthcare datasets [26, 32, 36, 52, 57]. Our results show that predictive checklists can perform as well as state-of-the-art classifiers on some datasets, and are able to inform practitioners of an optimality gap when datasets will not admit a checklist that is sufficiently accurate. We find that fairness gaps induced by checklists can be addressed with our added constraints, improving worst-case FPR disparity from 54.5% to 38.3% on our mortality prediction task. We give a demonstration of the flexibility and practical utility of predictive checklists through a detailed clinical application - creating a short-form model of the PCL Checklist for PTSD diagnosis.

The main contributions of this paper are:

1. We present a machine learning method to learn checklists from data. Our method allows practitioners to customize models to obey a wide range of real-world constraints. In addition, it pairs checklists with an optimality gap that can inform practitioners in model development.
2. We develop specialized techniques that improve the ability of our approach to train a checklist that performs well, and to pair this model with a small optimally gap. One of these techniques can be used to train checklists heuristically.
3. We provide a Python package to train and customize checklists using common open-source and commercial solvers such as CPLEX [22] and CBC [29].

2 Related Work

The use of checklist-style classifiers – e.g., classifiers with coefficients of ± 1 to Boolean conditions – dates to the work of Burgess [12]. The practice is often referred to as "unit weighting," and is motivated by observations that classifiers with unit weights may perform surprisingly well [see e.g., 10, 21, 25, 28]. The preference for checklists-style classifiers is pervasive in clinical applications [20] – see e.g., the TIMI Risk Score Antman et al. [3], which predicts mortality for patients with heart conditions, and the Adult ADHD Self-Report Scale [38], which has the same format as the PTSD checklist in Section 6.

Interpretable ML Our work is strongly related to methods that learn sparse linear classifiers with small integer coefficients [see e.g., 7, 14, 31, 37, 65, 66, 72]. We learn checklists from data by solving an integer program that can be viewed as a special case of an integer program to fit sparse linear classifiers with small integer coefficients presented in Ustun and Rudin [65], targeting constraints that binarize continuous and categorical features into items that can be included in a checklist [see also 13, 15, 68]. More generally, by restricting our focus to a special class of models, we are able to address the issues relevant to checklists, and develop specialized techniques to reduce computation.

Rule-Based Models Checklists are M-of-N rules — i.e., classifiers that predict $y = +1$ when M out of N conditions are true. Early work in machine learning used M-of-N rules in an auxiliary capacity – e.g., to explain the predictions of neural networks [64], or to use as features in decision trees [74]. More recent work has focused on learning M-of-N rules for standalone prediction – e.g. Chevaleyre et al. [18] describe a method to learn M-of-N rules by rounding the coefficients of linear SVMs to $\{-1, 0, 1\}$ using a randomized rounding procedure [see also 75]. Methods to learn disjunctions and conjunctions are also relevant since they correspond to 1-of-N rules and N-of-N rules, respectively. These methods effectively learn models by solving a special case of the ERM problem in (1) where $M = 1$ or $M = N$. Given that this problem is NP-hard, existing methods reduce computation by using algorithms that return approximate solutions – e.g., simulated annealing [69], set cover [49, 50], and ERM with a convex surrogate loss [23, 24, 47, 48]. These approaches could be used to learn checklists by solving the ERM problem in (1). However, they would not be able to handle the discrete constraints needed for binarization and customization. More generally, they would not be guaranteed to recover the most accurate checklists for a given dataset, nor pair models with a certificate of optimality that can be used to inform model development.

3 Methodology

We start with a dataset of n examples $(\boldsymbol{x}_i, y_i)_{i=1}^n$ where $\boldsymbol{x}_i = [x_{i1}, \ldots, x_{id}] \in \{0,1\}^d$ is a vector of d binary variables and $y_i \in \{0, 1\}$ is a label. Here, $\boldsymbol{x}_i$ are Boolean variables that could be used as

items in a checklist. We denote the indices of positive and negative examples as $I^+ = \{i \mid y_i = +1\}$ and $I^- = \{i \mid y_i = -1\}$ respectively, and let $n^+ = |I^+|$ and $n^- = |I^-|$. We denote the set of positive integers up to k as $[k] = \{1, \ldots, k\}$.

We use the dataset to learn a *predictive checklist* – i.e., a Boolean threshold rule that predicts $\hat{y} = +1$ when at least M of N items are checked. We represent a checklist as a linear classifier with the form:

$$\hat{y}_i = \text{sign}(\boldsymbol{\lambda}^\top \boldsymbol{x}_i \geq M).$$

Here $\boldsymbol{\lambda} = [\lambda_1, \ldots, \lambda_d] \in \{0,1\}^d$ is a coefficient vector and $\lambda_j = 1$ iff the checklist contains item j. We denote the number of items in a checklist with coefficients $\boldsymbol{\lambda}$ as $N = \sum_{j=1}^d \lambda_j$, and denote the threshold number of items that must be checked to assign a positive prediction as M. We learn checklists from data by solving an empirical risk minimization problem with the form:

$$\begin{aligned} \min_{\boldsymbol{\lambda}, M} \quad & l(\boldsymbol{\lambda}, M) + \epsilon_N N + \epsilon_M M \\ \text{s.t.} \quad & N = \|\boldsymbol{\lambda}\| \\ & M \in [N] \\ & \boldsymbol{\lambda} \in \{0,1\}^d \end{aligned} \tag{1}$$

Here, $l(\boldsymbol{\lambda}, M) = \sum_{i=1}^n \mathbb{1}[y_i \neq \hat{y}_i]$ counts the number of mistakes of a checklist with parameters $\boldsymbol{\lambda}$ and M. The parameters $\epsilon_N > 0$ and $\epsilon_M > 0$ are small penalties used to specify lexicographic preferences. We set $\epsilon_N < \frac{1}{nd}$ and $\epsilon_M < \frac{\epsilon_N}{d}$. These choices ensure that optimization will return the most accurate checklist, breaking ties between checklists that are equally accurate to favor smaller N, and breaking ties between checklists that are equally accurate and sparse to favor smaller M. Smaller values of N and M are preferable, as checklists with smaller N users check fewer items, and checklists with smaller M let users stop checking as soon as M items are checked (see Figure 3).

We recover a globally optimal solution to (1) by solving the integer program (IP):

$$\begin{aligned} \min_{\boldsymbol{\lambda}, z, M} \quad & l^+ + W^- l^- + \epsilon_N N + \epsilon_M M \\ \text{s.t.} \quad & B_i z_i \geq M - \sum_{j=1}^d \lambda_j x_{i,j} && i \in I^+ && \text{(2a)} \\ & B_i z_i \geq \sum_{j=1}^d \lambda_j x_{i,j} - M + 1 && i \in I^- && \text{(2b)} \\ & l^+ = \sum_{i \in I^+} z_i \\ & l^- = \sum_{i \in I^-} z_i \\ & N = \sum_{j=1}^d \lambda_j \\ & M \in [N] \\ & z_i \in \{0,1\} && i \in [n] \\ & \lambda_j \in \{0,1\} && j \in [d] \end{aligned}$$

Here, l^+ and l^- are variables that count the number of mistakes on positive and negative examples. The values of l^+ and l^- are computed using the mistake indicators $z_i = \mathbb{1}[\hat{y}_i \neq y_i]$, which are set to 1 through "Big-M" Constraints in (2a) and (2a).These constraints depend on the "Big-M" parameters B_i, which can be set to its tightest possible value $\max_\lambda M - \lambda^\top \boldsymbol{x}_i$. W^- is a user-defined parameter that reflects the relative cost of misclassifying a negative example. By default, we set $W^- = 1$, so that the objective minimizes training error.

Customization The IP in (2) can be customized to fit checklists that obey a wide range of real-world requirements on performance and model form. We list examples of constraints in Table 1.

Adaptive Binarization Practitioners often work with datasets that contain real-valued and categorical features. In such settings, rule-learning methods often require practitioners to binarize such features before training. Our approach allows practitioners to binarize features into items during training. This approach requires practitioners to first define T_j candidate items for each non-binary feature $u_{i,j}$. For example, given a real-valued feature $u_{i,j} \in \mathbb{R}$, the set of candidate items could

MODEL REQUIREMENT	EXAMPLE	CONSTRAINT
Model Size	Use ≤ 8 items	$N \leq 8$
Procedural	If checklist includes item about coughing, then include item about fever	$\lambda_{fever} \geq \lambda_{cough}$
Prediction	Predict $\hat{y}_i = +1$ for all patients with both fever and coughing	$\lambda_{fever} x_{i,fever} + \lambda_{cough} x_{i,cough} \geq M \; \forall \, i \in [n]$
Class-Based Accuracy	Max FPR $\leq \gamma\%$	$l^- \leq \lceil \gamma \cdot n^- \rceil$
Group Fairness	Max FPR disparity of $\gamma\%$ between males and females	$\left\lvert \frac{l_M^-}{n_M^-} - \frac{l_F^-}{n_F^-} \right\rvert \leq \gamma$

Table 1: Model requirements that can be addressed by adding constraints to IP (2). Each requirement can be directly encoded into the IP (2). The IP can then be solved using the same MIP solver to produce a most accurate checklist that obeys these constraints.

take the form $\{x_{i,j,1}, x_{i,j,2}, \ldots, x_{i,j,T_j}\}$ where $x_{i,j,t} = \mathbb{1}[u_{i,j} \geq v_{j,t}]$ and $v_{j,t}$ is a threshold. We would add the following constraint to IP (2) to ensure that the checklist only uses one of the T_j items: $\sum_{t \in [T_j]} \lambda_{t_j} \leq 1$. In this way, the IP would then choose a binary item that is aligned with the goal of maximizing accuracy. In practice, the set of candidate items can be produced automatically [e.g., using binning or information gain 54] or specified on the basis of clinical standards (as real-valued features like blood pressure and BMI have established thresholds). This approach is general enough to allow practitioners to optimize over all possible thresholds, though this may lead to overfitting. Our experiments in Section 5.2 show that dynamic binarization produces a meaningful improvement in performance for our model class.

Optimality Gap We solve IP (2) with a MIP solver, which finds a globally optimal solution through exhaustive search algorithms like branch-and-bound [70]. Given an instance of IP (2) and a time limit for when to stop the search, a solver returns: (i) the best feasible checklist found within the time limit – i.e., values of $\boldsymbol{\lambda}$ and M that achieve the best objective value $V^{\text{max}} = l(\boldsymbol{\lambda}, M)$; and (ii) a lower bound on the objective value – i.e., V^{min}, the minimal objective value for any feasible solution. These quantities are used to compute the *optimality gap*: $\varepsilon := 1 - \frac{V^{\text{min}}}{V^{\text{max}}}$. When the upper bound V^{max} matches the lower bound V^{min}, the solver returns a checklist with an optimality gap of $\varepsilon = 0$ – i.e., one that is *certifiably optimal*.

In our setting, the optimality gap reflects the worst-case difference in training error between the checklist that a solver returns and the most accurate checklist that we could hope to obtain by solving IP (2). Given a solution to IP (2) with an objective value of L and an optimality gap of ε, any feasible checklist has a training error of at least $\lceil (1-\varepsilon)L \rceil$. When ε is small, the most accurate checklist has a training error $\approx L$. Thus, if we are not satisfied with the performance of our model, we know that no checklist can perform better than L, and can make an informed decision to relax constraints, fit a classifier from a more complex hypothesis class, or not fit a classifier at all. When ε is large, our problem may admit a checklist with training error far smaller than L. If we are not satisfied with the performance of the checklist, we cannot determine whether this is because the dataset does not admit a sufficiently accurate model or because the optimization procedure has not been able to find it. This motivates the need to recover a certifiably optimal solution.

4 Algorithmic Improvements

In this section, we present specialized techniques to speed up the time required to find a feasible solution, and/or to produce a solution with a smaller optimality gap.

4.1 Submodular Heuristic

Our first technique is a submodular optimization technique that we use to generate initial feasible solutions for the IP in Formulation 2. We consider a knapsack cover problem. This problem can be solved to produce a checklist of N_{max} or fewer items that: (1) maximizes coverage of positive examples; (2) has a "budget" of negative examples B; and (3) uses at most one item from each of the feature groups $R_1, \ldots, R_T$ where each R_t denote a set of items derived from a particular real-valued feature and $\cup_{t=1}^{T} R_t = [d]$.

$$\max_{A} \quad f_M^+(A) \tag{3a}$$
$$\text{s.t.} \quad |A| \leq N_{max} \tag{3b}$$
$$|A \cap R_t| \leq 1 \qquad t \in [T] \tag{3c}$$

$$\sum_{j \in A} c(j) \leq B \tag{3d}$$

Here, $f_M^+ := \sum_{i \in I^+} \min(\sum_{j \in A} x_{i,j}, M)$ counts the positive examples that are covered by at least M items in A. Constraint (3b) limits the number of items in the checklist. Constraint (3c) requires the checklist to use at most one item from each feature group $R_1, \ldots, R_T$. Constraint (3d) controls the number of negative examples that are covered by at least M items in A. This is a knapsack constraint that assigns a cost to each item as $c(j) = \sum_{i \in I^-} x_{i,j}$ and limits the overall cost of A using the budget parameter $B > 0$.

The optimization problem in (3) can be solved using a submodular minimization algorithm given that f_M^+ is monotone submodular (see Appendix A.1 for a proof) and that constraints (3b) to (3d) are special kinds of matroid constraints. We solve (3) with a submodular minimization algorithm adapted from Badanidiyuru and Vondrák [5] (Appendix A.2). The procedure takes as input M, N_{max} and B and outputs a set of checklists that use up to N_{max} items. This procedure can run within 1 second. Given a problem where we would want to train a checklist that has at most $N_{\max}$ items, we would solve this problem $|\mathcal{B}|N_{\max}$ times, varying $M \in [N_{max}]$ and $B \in \mathcal{B}$.

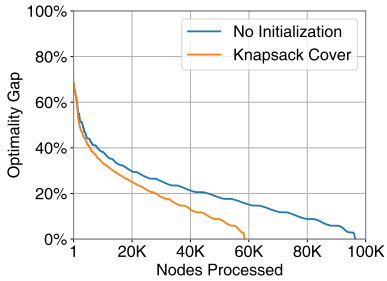

(a)

Algorithm 1: Sequential Training over N and M

Input: maximum # of items N_{max}, objective value function $V(\cdot)$
pool $\leftarrow \emptyset$
for $i \in [N_{max}]$, $j \in [i]$ **do**
  $S \leftarrow \{C | C \in \text{pool}, N(C) \leq i, M(C) \leq j\}$
  $C_{init} \leftarrow \text{argmin}_{C \in S} f(C)$
  $C_{new} \leftarrow$ solve IP (2) with initial solution C_{init} subject to $N \leq i$ and $M \leq j$
  Add $\{C_{new}\}$ to pool
end
return *pool*

(b)

Figure 1: (a) Performance profile when solving IP (2) with and without initialization. We show the optimality gap vs. the # of branch-and-bound nodes when solving IP (2) for the `heart` dataset where $N \leq 6$ and $M \leq 2$. (b) algorithm for sequential training over N and M.

In Figure 1a, we show the effect of initializing the IP with checklists produced using the submodular heuristic. As shown, these solutions allow us to initialize the search with a good upper bound, allowing us to ultimately recover a certifiably optimal solution faster. Our empirical results in Section 5.2 show that checklists from our submodular heuristic outperform checklists fit using existing baselines. This highlights the value of this approach as a heuristic to fit checklists (in that it can produce checklists that perform well), and as a technique for initialization (in that it can produce checklists that can be used to initialize IP (2) with a small upper bound).

4.2 Path Algorithms

Our second technique is a path algorithm for sequential training. Given an instance of the checklist IP (2), this algorithm sequentially solves smaller instances of IP (2). The algorithms exploits the fact that small instances can be solved to optimality quickly and thus "ruled out" of the feasible region of subsequent instances. The resulting approach can reduce the size of the branch-and-bound tree and improve the ability to recover a checklist with a small optimality gap. These algorithms can wrap around any procedure to solve IP (2).

In Algorithm 1, we present a path algorithm to train checklists over the full range of model sizes from 1 to $N^{\max}$. The algorithm recovers a solution to a problem with a hard model size limit by solving smaller instances with progressively larger limits of $N_{\max}$ and $M_{\max}$. We store the best feasible solution from each iteration in a pool, which is then used to initialize future iterations. If the best feasible solution is certifiably optimal, we add a constraint to constrict the search space for subsequent instances. This allows us to reduce the size of the branch-and-bound tree, resulting in a faster solution in each iteration.

5 Experimental Results

In this section, we benchmark the performance of our methods on clinical prediction tasks. We attach all code to reproduce our results to our submission, and include supporting material in Appendix C.

5.1 Setup

Data. We consider eight clinical classification tasks shown in Table 2. For each task, we create a classification dataset by using one of the following techniques to binarize features:

Fixed: We convert continuous and ordinal features into threshold indicator variables using the median, and convert categorical feature into an indicator for the most common category.

Adaptive: We convert continuous and ordinal features into 4 threshold indicator variables using quintiles as thresholds, and convert each categorical feature with a one-hot encoding.

Optbinning: We use the method proposed by Navas-Palencia [54] to binarize all features.

We process each dataset to oversample the minority class to equalize the number of positive and negative examples due to class imbalance. We split each dataset into a 80% training and 20% test set.

Experiments. We use the training set to fit predictive checklists that use at most $N \leq 8$ items and that are required to select at most 1 item from a feature group. We compare the performance of the following methods for creating checklists:

MIP: We fit a predictive checklist by solving an instance of IP (2) with the appropriate constraints. We solve this problem using CPLEX 12.10 [22] paired with the computational improvements in Section 4 on a 2.4 GHz CPU with 16 GB RAM for ≤ 10 minutes.

Cover: We fit a predictive checklist using the submodular heuristic in Section 4.1. We set $N_{max} = 8$, and vary $M \in [N_{max}]$ and $B \in \{kn^{+} | k \in \{\frac{1}{3}, \frac{1}{2}, 1, 2, 3\}\}$.

Unit: We fit a predictive checklist using unit weighting [10]. We fit L_1-regularized logistic regression models for L_1 penalties $\in [10^{-5}, 10]$. We convert each model into an array of checklists by including items with positive coefficients and the complements of items with negative coefficients. We convert each logistic regression model into multiple checklists by setting $M \in [N]$.

Evaluation. We report results for MIP checklists where $M = 1$, which we refer to as MIP_OR since it corresponds to an OR rule. We also report performance for L_1-regularized logistic regression (LR) and XGBoost (XGB) as baselines to evaluate performance. We train models with these methods using all features in an adaptive-binarized dataset to obtain an informal upper bound on checklist performance.

For each method and each dataset, we report the performance of models that achieve the lowest training error with $N \leq 8$ items. We compare the performance of *Fixed* and *Adaptive* binarization here, and report performance of *Optbinning* in Appendix C.2.

5.2 Results

In Table 3, we report performance for each method on each dataset. We show the performance effects of varying N in Figure 2, and provide examples of predictive checklists trained using our method in Figure 3 and Appendix C.3. In what follows, we discuss these results.

Dataset	d_T	d_{pct}	n	Prediction Task	Reference
adhd	5	20	632	Patient diagnosed with ADHD	[52]
cshock	40	52	3,456	Patient has cardiogenic shock	[36]
cshock2	40	52	15,254	Patient has cardiogenic shock	[57]
kidney	17	80	1,760	Patient with kidney failure died after renal replacement therapy	[36]
mortality	50	484	36,684	Patient died in hospital	[36]
ptsd	20	80	1,106	Patient has a PCL-5 PTSD diagnosis	[52]
heart	13	82	330	Patient has heart disease	[26]
readmit	42	94	16,732	Patient re-admitted to hospital within 30 days	[32]

Table 2: Clinical prediction tasks used in Section 5. All datasets are publically available. Here, d_T is the number of feature groups, and d_{pct} is the number of features after binarization using the adaptive method.

		Fixed Binarization				Adaptive Binarization				Lower Bounds	
Dataset	**Metric**	Unit	Cover	MIP_OR	MIP	Unit	Cover	MIP_OR	MIP	LR	XGB
adhd	test error (%)	12.7	12.7	12.7	12.7	1.6	7.1	7.1	**0.8**	0.8	1.6
$n = 594$	train error (%)	11.7	11.7	11.7	11.7	1.0	4.3	4.3	**0.4**	0.2	1.0
$d_{pct} = 20$	gap (%)	-	-	0.0	0.0	-	-	0.0	0.0	-	-
cshock	test error (%)	30.8	**28.9**	35.1	29.9	29.0	30.8	33.5	30.8	29.2	29.3
$n = 2237$	train error (%)	31.6	30.2	35.6	29.3	28.4	31.3	34.0	**28.0**	23.3	3.7
$d_{pct} = 52$	gap (%)	-	-	36.5	72.7	-	-	0.0	73.8	-	-
cshock2	test error (%)	23.7	24.2	30.8	23.6	**21.6**	25.1	31.0	22.0	21.7	21.7
$n = 8815$	train error (%)	26.1	24.9	29.3	24.1	23.4	24.1	28.7	**22.5**	21.0	8.2
$d_{pct} = 51$	gap (%)	-	-	27.0	81.0	-	-	38.2	80.9	-	-
kidney	test error (%)	36.8	39.7	38.4	33.9	32.8	38.4	33.3	**30.2**	30.2	25.1
$n = 1722$	train error (%)	33.4	34.7	36.6	30.8	34.6	36.3	33.6	**28.5**	27.0	0.5
$d_{pct} = 80$	gap (%)	-	-	30.4	77.1	-	-	61.0	75.4	-	-
mortality	test error (%)	34.4	35.7	35.7	**29.8**	33.1	36.9	36.4	30.4	25.2	27.9
$n = 21139$	train error (%)	32.8	37.7	37.7	**28.9**	34.4	36.5	38.1	29.6	22.9	8.2
$d_{pct} = 484$	gap (%)	-	-	69.5	86.3	-	-	34.7	84.0	-	-
ptsd	test error (%)	9.5	10.4	13.5	10.4	9.0	11.3	12.2	**7.7**	6.8	5.0
$n = 873$	train error (%)	8.3	9.0	13.3	7.6	8.3	10.7	12.9	**5.4**	5.2	0.0
$d_{pct} = 80$	gap (%)	-	-	0.0	0.0	-	-	0.0	57.0	-	-
readmit	test error (%)	38.7	40.1	37.5	37.6	35.4	35.4	33.8	**32.7**	34.0	36.9
$n = 9766$	train error (%)	36.0	37.2	36.4	35.4	34.0	34.1	33.1	**33.0**	31.5	16.7
$d_{pct} = 94$	gap (%)	-	-	20.7	84.1	-	-	41.0	78.5	-	-
heart	test error (%)	24.4	30.5	31.7	23.2	25.6	20.7	29.3	**19.5**	24.4	20.7
$n = 303$	train error (%)	12.1	13.3	20.2	11.3	12.5	14.5	19.8	**10.5**	11.7	0.8
$d_{pct} = 82$	gap (%)	-	-	0.0	0.0	-	-	0.0	31.2	-	-

Table 3: Performance results of all methods on all datasets. We report the training error, test error, and optimality gap for the classifier that minimizes training error and satisfies all constraints.

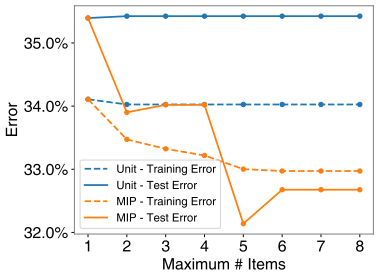

Figure 2: Performance of checklists with $\leq N$ items on `readmit` fit using our approach and unit weighting.

Predict 30-day readmission if 3+ items are checked	
# admissions past year ≥ 2	☐
bed days ≥ 3	☐
length of stay ≥ 5 days	☐
chronic or uncontrolled pain	☐
palliative care	☐
substance abuse past month	☐

Figure 3: Predictive checklist fit for `readmit` using MIP (train/test error = 33.0%/32.7%). We show sample checklists for all other datasets in Appendix C.3.

On Performance Our results in Table 3 show that our method MIP outperforms existing techniques to fit checklists (Unit) on 8/8 datasets in terms of training error and 7/8 datasets in terms of test error. As shown in Figure 2, this holds across all values of N. The performance loss of checklists relative to more complex model classes (LR and XGB) varies greatly between datasets. Checklists trained on `heart` and `readmit` perform nearly as well as more complex classifiers. On `mortality` and `kidney`, however, the loss may exceed 5%. From our results in Table 3, we observe that adaptive binarization achieves a better training error than fixed binarization for 7/8 datasets, highlighting the potential to improve performance by binarizing during training time (see Section 3 for a discussion).

The checklists in Appendix C.3 highlight how smaller values of M can produce checklists that are easier to use. However, larger values of M yield in performance gains across all datasets (see e.g., results for `heart` where the difference in test error between MIP_OR and MIP is 9.8%).

On Generalization Our results show that checklist performance on the training set tends to generalize well to the test set. This holds for all checklists, including Unit, Cover, MIP_OR, and MIP, as expected given that checklists belong to a simple model class. Generalization can help us customize checklists in an informed way since we can assume that any changes to the training loss (e.g. with the addition of a constraint) would result in a similar change in the test set. In this regime, practitioners can evaluate the impact of model requirements on *predictive* performance – since they can measure differences in training error between checklists that obey different requirements, and expect these differences in training error as a reliable proxy for differences in test error.

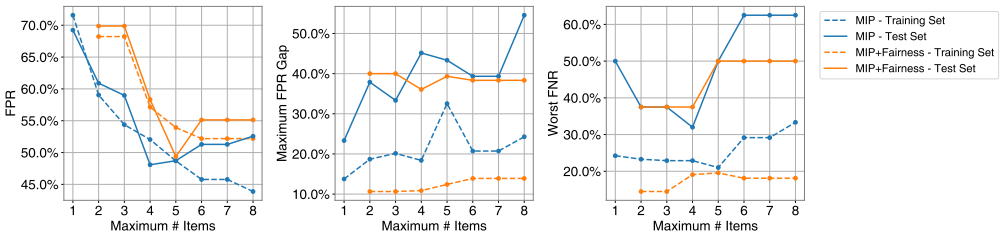

Figure 4: FPR (left), maximum FPR disparity (center) and worst-case FNR (right) versus number of items for checklists for mortality prediction given CRRT using `kidney`. Group fairness constraints effectively limit the FPR disparity and worst-case FNR during training, at the cost of overall performance.

Predict Mortality Given CRRT if 3+ Items are Checked	
Age $\geq$ 66.0 years	☐
AST $\geq$ 162.6 IU/L	☐
Blood pH $\leq$ 7.29	☐
MCV $\geq$ 99.0 fl	☐
Norepinephrine $\geq$ 0.1 mcg/kg/min	☐
Platelets $\leq$ 65.0 $\times 10^3/\mu L$	☐
RDW $\geq$ 19.2%	☐
Time in ICU $\geq$ 14.1 hours	☐

	FNR	FPR	Worst FNR	Max FPR Gap
Training	20.0%	43.9%	33.3%	24.3%
Test	22.2%	52.6%	62.5%	54.5%

(a) MIP

Predict Mortality Given CRRT if 2+ Items are Checked	
ALT $\geq$ 16.0 IU/L	☐
Bicarbonate $\leq$ 17.0 mmol/L	☐
Blood pH $\leq$ 7.22	☐
Norepinephrine $\geq$ 0.1 mcg/kg/min	☐
RDW $\geq$ 19.2%	☐
Time in ICU $\geq$ 117.3 hours	☐

	FNR	FPR	Worst FNR	Max FPR Gap
Training	17.5%	52.2%	18.1%	13.9%
Test	19.6%	55.1%	50.0%	38.3%

(b) MIP + Fairness Constraints

Figure 5: Checklists trained to predict mortality given CRRT (a) without group fairness constraints and (b) with group fairness constraints. Checklist (a) has better overall FPR, but has worse FNR for the worst-case group, and exhibits disparities in FPR across groups. Checklist (b) obeys group fairness constraints that attenuate the training FPR gap to 15% and the worst-case FNR to 20%.

On the Value of an Optimality Gap One of the benefits of our approach is that it pairs each checklist with an optimality gap, which can inform model development by identifying settings where we cannot train a checklist that is sufficiently accurate (see Section 3). On `heart`, we train a checklist with a training error of 10.5% and an optimality gap of 31.2%, which suggests that there may exist a checklist with a training error of 7.2%. This optimality gap is not available for checklists trained with a heuristic (like unit weighting or domain expertise), so we would not be able to attribute poor model performance to a difficult problem, or to the heuristic being ineffective.

On Fairness We analyze the fairness of MIP checklists trained on the `kidney` dataset. The goal is to predict in-hospital mortality in patients with kidney failure after the administration of Continuous Renal Replacement Therapy (CRRT) in an ICU setting. In screening checklists, a false negative is worse than a false positive, so we train checklists that minimize FPR (i.e. setting $W^- = n^-$ in IP 2) subject to a constraint on FNR at 20%. We find that our checklists exhibit performance disparities across sex and race – consistent with the tendency of simpler models to exhibit performance disparities [39].

We next train a checklist with constraints that limit disparities in training FPR and FNR over 6 intersectional groups $g \in G = \{\texttt{Male}, \texttt{Female}\} \times \{\texttt{White}, \texttt{Black}, \texttt{Other}\}$. Specifically, 30 constraints that limit the FPR disparity across groups to 15%, and 6 constraints to cap the max FNR for each group to 20%. This is unique compared prior fair classification work [73] because it mitigates unfairness over intersectional groups in an exact way [i.e., no convex relaxations 46]. We find that we can limit performance disparities across all values of N (see Figures 4 and 5, and Appendix C.4), observing that MIP with group fairness constraints exhibits lower FPR disparities over groups, e.g., reducing the test set FPR gap between Black females and White males from 54.5% to 30.6%, but has larger overall FPR. By recovering models along this trade-off, domain experts can make informed decisions about model deployment.

6 Learning a Short-Form Checklist for PTSD Screening

In this section, we demonstrate the practical benefits of our approach by learning a short-form checklist for Post-Traumatic Stress Disorder (PTSD) diagnosis.

Background The *PTSD Checklist for DSM-5* (PCL-5) is a self-report screening tool for PTSD [8, 71]. The PCL-5 consists of 20 questions that assess the severity of symptoms of PTSD DSM-5. Patients respond to each question with answers of *Not at all*, *A little bit*, *Moderately*, *Quite a bit*, or *Extremely*. Given these responses, the PCL assigns a provisional diagnosis of PTSD by counting the number of responses of *Moderately* or more frequent across four clusters of questions. Patients with a provisional diagnosis are then assessed for PTSD by a specialist. The PCL-5 can take over 10 minutes to complete, which limits its use in studies that include the PCL-5 in a battery of tests (i.e., using the provisional diagnosis to evaluate the prevalence of PTSD and its effect as a potential confounder), and has led to the development of short-form models [see e.g., 9, 43, 76].

Problem Formulation Our goal is to create a *short-form* version of the PCL-5 – i.e., a checklist that assigns the same provisional diagnoses as PCL-5 using only a subset of the 20 questions. We train our model using data from the AURORA study [52], which contains PCL-5 responses of U.S. patients who have visited the ED following a traumatic event. Here, $y_i = +1$ if the patient is assigned a provisional PTSD diagnosis. We encode the response for each question into four binary variables $x_{i,k,l} = \mathbb{1}[q_k \geq l]$ where q_k is the response to each question and $l \in [4]$ is denotes its response. Our final dataset contains $d = 80$ binary variables for $n = 873$ patients.

Since this checklist would be used to screen patients, a false negative diagnosis is less desirable than a false positive diagnosis. We train a checklist that minimizes FPR and that obeys the following operational constraints: (1) limit FNR to 5%; (2) pick one threshold for each PCL-5 question; (3) use at most 8 questions. We make use of the same methods and techniques as in Section 5.2.

Results Our results show that short-form checklists fit using our method outperform other checklists (see Figure 6). In Figure 7, we display a sample $N = 8, M = 4$ checklist trained with our method which reproduces the PCL-5 diagnosis with high accuracy, offering an alternative to the PCL-5 with simpler binary items and less than half the number of questions.

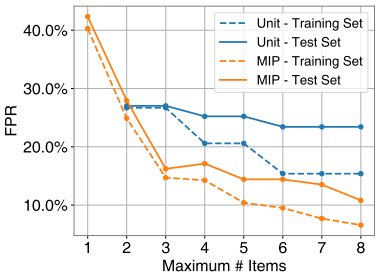

Figure 6: FPR versus number of items for short-forms of the PCL-5. MIP outperforms Unit across all values of N.

Predict PTSD if 4+ Items are Checked	
Avoiding thinking about experience $\geq$ moderately	☐
Avoiding activities or situations $\geq$ moderately	☐
Blame for stressful experience $\geq$ a little bit	☐
Feeling distant or cut off $\geq$ quite a bit	☐
Irritable or angry outbursts $\geq$ moderately	☐
Loss of interest in activities $\geq$ moderately	☐
Repeated disturbing dreams $\geq$ moderately	☐
Trouble experiencing positive feelings $\geq$ moderately	☐

Figure 7: Short-form of the PCL-5 trained with MIP. The checklist has a train/test FPR of 6.6%/10.8%, and a train/test FNR of 4.7%/6.3%. It reproduces the PCL-5 diagnosis on the training/test sets with 94.1/90.9% accuracy, and has a 70.8% optimality gap. We show condensed PCL-5 questions due to space constraint.

7 Conclusion

In this work, we developed a machine learning method to learn predictive checklists from data by solving an IP to fit sparse linear classifiers with small integer coefficients and adaptive binarization. More broadly, our work highlights a promising approach to fit models to obey constraints related to qualities like interpretability, fairness [6, 17], and safety [1]. Using our approach, practitioners can potentially co-design checklists alongside clinicians – by encoding requirements into the model and reporting their effects on predictive performance [19, 35].

We caution against the use of items in a checklist as causal factors for outcomes (i.e., questions in the short-form PCL are not "drivers" of PTSD), and against the use of our method solely due to its interpretability. Recent work has argued against the use of models simply because they are "interpretable" [27, 45] and highlights the importance of validating claims surrounding interpretability by user studies [41, 58]. Interesting directions for future work include: learning checklists for risk prediction [66]; optimizing the order of items to minimize response time; and studying interventions that further improve predictive performance in settings where checklists fail to perform well.

References

[1] Dario Amodei, Chris Olah, Jacob Steinhardt, Paul Christiano, John Schulman, and Dan Mané. Concrete problems in ai safety. *arXiv preprint arXiv:1606.06565*, 2016.

[2] Elaine Angelino, Nicholas Larus-Stone, Daniel Alabi, Margo Seltzer, and Cynthia Rudin. Learning certifiably optimal rule lists for categorical data. *Journal of Machine Learning Research*, 18:1–78, 2018.

[3] Elliott M Antman, Marc Cohen, Peter JLM Bernink, Carolyn H McCabe, Thomas Horacek, Gary Papuchis, Branco Mautner, Ramon Corbalan, David Radley, and Eugene Braunwald. The TIMI risk score for unstable angina/non–ST elevation MI. *The Journal of the American Medical Association*, 284(7):835–842, 2000.

[4] BM Astrachan, M Harrow, D Adler, L Brauer, A Schwartz, C Schwartz, and G Tucker. A checklist for the diagnosis of schizophrenia. *The British Journal of Psychiatry*, 121(564): 529–539, 1972.

[5] Ashwinkumar Badanidiyuru and Jan Vondrák. Fast algorithms for maximizing submodular functions. In *Proceedings of the twenty-fifth annual ACM-SIAM symposium on Discrete algorithms*, pages 1497–1514. SIAM, 2014.

[6] Solon Barocas, Moritz Hardt, and Arvind Narayanan. *Fairness and Machine Learning*. fairmlbook.org, 2018. http://www.fairmlbook.org.

[7] Lieven Billiet, Sabine Van Huffel, and Vanya Van Belle. Interval coded scoring: a toolbox for interpretable scoring systems. *PeerJ Computer Science*, 4:e150, 2018.

[8] Christy A Blevins, Frank W Weathers, Margaret T Davis, Tracy K Witte, and Jessica L Domino. The posttraumatic stress disorder checklist for dsm-5 (pcl-5): Development and initial psychometric evaluation. *Journal of traumatic stress*, 28(6):489–498, 2015.

[9] Paul D Bliese, Kathleen M Wright, Amy B Adler, Oscar Cabrera, Carl A Castro, and Charles W Hoge. Validating the primary care posttraumatic stress disorder screen and the posttraumatic stress disorder checklist with soldiers returning from combat. *Journal of consulting and clinical psychology*, 76(2):272, 2008.

[10] Philip Bobko, Philip L Roth, and Maury A Buster. The usefulness of unit weights in creating composite scores: A literature review, application to content validity, and meta-analysis. *Organizational Research Methods*, 10(4):689–709, 2007.

[11] David Brodie and Richard Wells. An evaluation of the utility of three ergonomics checklists for predicting health outcomes in a car manufacturing environment. In *Proc. of the 29th Annual Conference of the Human Factors Association of Canada*. Citeseer, 1997.

[12] Ernest W Burgess. Factors determining success or failure on parole. *The workings of the indeterminate sentence law and the parole system in Illinois*, pages 221–234, 1928.

[13] Emilio Carrizosa, Belen Martin-Barragan, and Dolores Romero Morales. Binarized support vector machines. *INFORMS Journal on Computing*, 22(1):154–167, 2010.

[14] Emilio Carrizosa, Amaya Nogales-Gómez, and Dolores Romero Morales. Strongly agree or strongly disagree?: Rating features in support vector machines. *Information Sciences*, 329: 256–273, 2016.

[15] Emilio Carrizosa, Amaya Nogales-Gómez, and Dolores Romero Morales. Clustering categories in support vector machines. *Omega*, 66:28–37, 2017.

[16] Ken Catchpole and Stephanie Russ. The problem with checklists. *BMJ quality & safety*, 24(9): 545–549, 2015.

[17] Simon Caton and Christian Haas. Fairness in machine learning: A survey. *arXiv preprint arXiv:2010.04053*, 2020.

[18] Yann Chevaleyre, Frederic Koriche, and Jean-Daniel Zucker. Rounding methods for discrete linear classification. In *Proceedings of the 30th International Conference on Machine Learning (ICML-13)*, pages 651–659, 2013.

[19] Evangelia Christodoulou, Jie Ma, Gary S Collins, Ewout W Steyerberg, Jan Y Verbakel, and Ben Van Calster. A systematic review shows no performance benefit of machine learning over logistic regression for clinical prediction models. *Journal of clinical epidemiology*, 110:12–22, 2019.

[20] Robyn Clay-Williams and Lacey Colligan. Back to basics: checklists in aviation and healthcare. *BMJ quality & safety*, 24(7):428–431, 2015.

[21] Jacob Cohen. Things i have learned (so far). In *Annual Convention of the American Psychological Association, 98th, Aug, 1990, Boston, MA, US; Presented at the aforementioned conference.* American Psychological Association, 1992.

[22] IBM ILOG Cplex. V12. 1: User's manual for cplex. *International Business Machines Corporation*, 46(53):157, 2009.

[23] Sanjeeb Dash, Dmitry M Malioutov, and Kush R Varshney. Screening for learning classification rules via boolean compressed sensing. In *Acoustics, Speech and Signal Processing (ICASSP), 2014 IEEE International Conference on*, pages 3360–3364. IEEE, 2014.

[24] Sanjeeb Dash, Dmitry M Malioutov, and Kush R Varshney. Learning interpretable classification rules using sequential rowsampling. In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3337–3341. IEEE, 2015.

[25] Robyn M Dawes. The robust beauty of improper linear models in decision making. *American psychologist*, 34(7):571–582, 1979.

[26] Robert Detrano, Andras Janosi, Walter Steinbrunn, Matthias Pfisterer, Johann-Jakob Schmid, Sarbjit Sandhu, Kern H Guppy, Stella Lee, and Victor Froelicher. International application of a new probability algorithm for the diagnosis of coronary artery disease. *The American journal of cardiology*, 64(5):304–310, 1989.

[27] Finale Doshi-Velez and Been Kim. Towards a rigorous science of interpretable machine learning. *arXiv preprint arXiv:1702.08608*, 2017.

[28] Hillel J Einhorn and Robin M Hogarth. Unit weighting schemes for decision making. *Organizational behavior and human performance*, 13(2):171–192, 1975.

[29] John Forrest, Ted Ralphs, Stefan Vigerske, Lou Hafer, Bjarni Kristjansson, JP Fasano, E Straver, M Lubin, HG Santos, R Lougee, et al. coin-or/cbc: Version 2.9. 9. *URL http://dx. doi. org/10.5281/zenodo*, 1317566, 2018.

[30] Brigid M Gillespie and Andrea Marshall. Implementation of safety checklists in surgery: a realist synthesis of evidence. *Implementation Science*, 10(1):137, 2015.

[31] Sharad Goel, Justin M Rao, Ravi Shroff, et al. Precinct or prejudice? understanding racial disparities in new york city's stop-and-frisk policy. *Annals of Applied Statistics*, 10(1):365–394, 2016.

[32] Jeffrey L Greenwald, Patrick R Cronin, Victoria Carballo, Goodarz Danaei, and Garry Choy. A novel model for predicting rehospitalization risk incorporating physical function, cognitive status, and psychosocial support using natural language processing. *Medical care*, 55(3): 261–266, 2017.

[33] Brigette Hales, Marius Terblanche, Robert Fowler, and William Sibbald. Development of medical checklists for improved quality of patient care. *International Journal for Quality in Health Care*, 20(1):22–30, 2008.

[34] Alex B Haynes, Thomas G Weiser, William R Berry, Stuart R Lipsitz, Abdel-Hadi S Breizat, E Patchen Dellinger, Teodoro Herbosa, Sudhir Joseph, Pascience L Kibatala, Marie Carmela M Lapitan, et al. A surgical safety checklist to reduce morbidity and mortality in a global population. *New England Journal of Medicine*, 360(5):491–499, 2009.

[35] Sungsoo Ray Hong, Jessica Hullman, and Enrico Bertini. Human factors in model interpretability: Industry practices, challenges, and needs. *Proceedings of the ACM on Human-Computer Interaction*, 4(CSCW1):1–26, 2020.

[36] Alistair EW Johnson, Tom J Pollard, Lu Shen, H Lehman Li-Wei, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G Mark. Mimic-iii, a freely accessible critical care database. *Scientific data*, 3(1):1–9, 2016.

[37] Jongbin Jung, Connor Concannon, Ravi Shroff, Sharad Goel, and Daniel G Goldstein. Simple rules to guide expert classifications. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 183(3):771–800, 2020.

[38] Ronald C Kessler, Lenard Adler, Minnie Ames, Olga Demler, Steve Faraone, EVA Hiripi, Mary J Howes, Robert Jin, Kristina Secnik, Thomas Spencer, et al. The world health organization adult adhd self-report scale (asrs): a short screening scale for use in the general population. *Psychological medicine*, 35(2):245, 2005.

[39] Jon Kleinberg and Sendhil Mullainathan. Simplicity creates inequity: implications for fairness, stereotypes, and interpretability. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 807–808, 2019.

[40] Heidi S Kramer and Frank A Drews. Checking the lists: A systematic review of electronic checklist use in health care. *Journal of biomedical informatics*, 71:S6–S12, 2017.

[41] Isaac Lage, Emily Chen, Jeffrey He, Menaka Narayanan, Been Kim, Samuel J Gershman, and Finale Doshi-Velez. Human evaluation of models built for interpretability. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing*, volume 7, pages 59–67, 2019.

[42] Himabindu Lakkaraju, Stephen H Bach, and Jure Leskovec. Interpretable decision sets: A joint framework for description and prediction. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 1675–1684, 2016.

[43] Ariel J Lang and Murray B Stein. An abbreviated ptsd checklist for use as a screening instrument in primary care. *Behaviour research and therapy*, 43(5):585–594, 2005.

[44] Lorelei Lingard, Glenn Regehr, Beverley Orser, Richard Reznick, G Ross Baker, Diane Doran, Sherry Espin, John Bohnen, and Sarah Whyte. Evaluation of a preoperative checklist and team briefing among surgeons, nurses, and anesthesiologists to reduce failures in communication. *Archives of surgery*, 143(1):12–17, 2008.

[45] Zachary C Lipton. The mythos of model interpretability: In machine learning, the concept of interpretability is both important and slippery. *Queue*, 16(3):31–57, 2018.

[46] Michael Lohaus, Michael Perrot, and Ulrike Von Luxburg. Too relaxed to be fair. In *International Conference on Machine Learning*, pages 6360–6369. PMLR, 2020.

[47] Dmitry Malioutov and Kush Varshney. Exact rule learning via boolean compressed sensing. In *Proceedings of The 30th International Conference on Machine Learning*, pages 765–773, 2013.

[48] Dmitry M Malioutov, Kush R Varshney, Amin Emad, and Sanjeeb Dash. Learning interpretable classification rules with boolean compressed sensing. In *Transparent Data Mining for Big and Small Data*, pages 95–121. Springer, 2017.

[49] Mario Marchand and John Shawe-Taylor. Learning with the set covering machine. In *ICML*, pages 345–352, 2001.

[50] Mario Marchand and John Shawe Taylor. The set covering machine. *The Journal of Machine Learning Research*, 3:723–746, 2003.

[51] Robert Mauro, Asaf Degani, Loukia Loukopoulos, and Immanuel Barshi. The operational context of procedures and checklists in commercial aviation. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, volume 56, pages 758–762. SAGE Publications Sage CA: Los Angeles, CA, 2012.

[52] Samuel A McLean, Kerry Ressler, Karestan Chase Koenen, Thomas Neylan, Laura Germine, Tanja Jovanovic, Gari D Clifford, Donglin Zeng, Xinming An, Sarah Linnstaedt, et al. The aurora study: a longitudinal, multimodal library of brain biology and function after traumatic stress exposure. *Molecular psychiatry*, 25(2):283–296, 2020.

[53] Keith E Morse, Steven C Bagley, and Nigam H Shah. Estimate the hidden deployment cost of predictive models to improve patient care. *Nature medicine*, 26(1):18–19, 2020.

[54] Guillermo Navas-Palencia. Optimal binning: mathematical programming formulation. *arXiv preprint arXiv:2001.08025*, 2020.

[55] Janki Patel, Kamran Ahmed, Khurshid A Guru, Fahd Khan, Howard Marsh, Mohammed Shamim Khan, and Prokar Dasgupta. An overview of the use and implementation of checklists in surgical specialities–a systematic review. *International Journal of Surgery*, 12(12): 1317–1323, 2014.

[56] Emma Pierson, David M Cutler, Jure Leskovec, Sendhil Mullainathan, and Ziad Obermeyer. An algorithmic approach to reducing unexplained pain disparities in underserved populations. *Nature Medicine*, 27(1):136–140, 2021.

[57] Tom J Pollard, Alistair EW Johnson, Jesse D Raffa, Leo A Celi, Roger G Mark, and Omar Badawi. The eicu collaborative research database, a freely available multi-center database for critical care research. *Scientific data*, 5(1):1–13, 2018.

[58] Forough Poursabzi-Sangdeh, Daniel G Goldstein, Jake M Hofman, Jennifer Wortman Wortman Vaughan, and Hanna Wallach. Manipulating and measuring model interpretability. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–52, 2021.

[59] Peter Pronovost, Dale Needham, Sean Berenholtz, David Sinopoli, Haitao Chu, Sara Cosgrove, Bryan Sexton, Robert Hyzy, Robert Welsh, Gary Roth, et al. An intervention to decrease catheter-related bloodstream infections in the icu. *New England Journal of Medicine*, 355(26): 2725–2732, 2006.

[60] Hajo Reijers, Henrik Leopold, and Jan Recker. Towards a science of checklists. In *Proceedings of the 50th Hawaii International Conference on System Sciences*, pages 5773–5782. University of Hawaii, 2017.

[61] Daniel L Stufflebeam. Guidelines for developing evaluation checklists: the checklists development checklist (cdc). *Kalamazoo, MI: The Evaluation Center Retrieved on January*, 16:2008, 2000.

[62] Ø Thomassen, Guttorm Brattebø, Eirik Søfteland, HM Lossius, and J-K Heltne. The effect of a simple checklist on frequent pre-induction deficiencies. *Acta anaesthesiologica scandinavica*, 54(10):1179–1184, 2010.

[63] Øyvind Thomassen, Ansgar Storesund, Eirik Søfteland, and Guttorm Brattebø. The effects of safety checklists in medicine: a systematic review. *Acta Anaesthesiologica Scandinavica*, 58(1): 5–18, 2014.

[64] Geoffrey G Towell and Jude W Shavlik. Extracting refined rules from knowledge-based neural networks. *Machine learning*, 13(1):71–101, 1993.

[65] Berk Ustun and Cynthia Rudin. Supersparse Linear Integer Models for Optimized Medical Scoring Systems. *Machine Learning*, 102(3):349–391, 2016.

[66] Berk Ustun and Cynthia Rudin. Learning optimized risk scores. *Journal of Machine Learning Research*, 20(150):1–75, 2019.

[67] Darshali A Vyas, Leo G Eisenstein, and David S Jones. Hidden in plain sight—reconsidering the use of race correction in clinical algorithms, 2020.

[68] Tong Wang. Multi-value rule sets. *arXiv preprint arXiv:1710.05257*, 2017.

[69] Tong Wang, Cynthia Rudin, Finale Doshi-Velez, Yimin Liu, Erica Klampfl, and Perry MacNeille. Or's of and's for interpretable classification, with application to context-aware recommender systems. *arXiv preprint arXiv:1504.07614*, 2015.

[70] Laurence A Wolsey. *Integer Programming*, volume 42. Wiley New York, 1998.

[71] Jennifer H Wortmann, Alexander H Jordan, Frank W Weathers, Patricia A Resick, Katherine A Dondanville, Brittany Hall-Clark, Edna B Foa, Stacey Young-McCaughan, Jeffrey S Yarvis, Elizabeth A Hembree, et al. Psychometric analysis of the ptsd checklist-5 (pcl-5) among treatment-seeking military service members. *Psychological assessment*, 28(11):1392, 2016.

[72] Mike Wu, Michael Hughes, Sonali Parbhoo, Maurizio Zazzi, Volker Roth, and Finale Doshi-Velez. Beyond sparsity: Tree regularization of deep models for interpretability. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.

[73] Muhammad Bilal Zafar, Isabel Valera, Manuel Gomez-Rodriguez, and Krishna P Gummadi. Fairness constraints: A flexible approach for fair classification. *J. Mach. Learn. Res.*, 20(75): 1–42, 2019.

[74] Zijian Zheng. Constructing x-of-n attributes for decision tree learning. *Machine learning*, 40 (1):35–75, 2000.

[75] Jean-Daniel Zucker, Yann Chevaleyre, and Dao Van Sang. Experimental analysis of new algorithms for learning ternary classifiers. In *2015 IEEE RIVF International Conference on Computing & Communication Technologies*, pages 19–24. IEEE, 2015.

[76] Kelly L Zuromski, Berk Ustun, Irving Hwang, Terence M Keane, Brian P Marx, Murray B Stein, Robert J Ursano, and Ronald C Kessler. Developing an optimal short-form of the ptsd checklist for dsm-5 (pcl-5). *Depression and anxiety*, 36(9):790–800, 2019.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes]
   (c) Did you discuss any potential negative societal impacts of your work? [Yes]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A] We do not include any theoretical results.
   (b) Did you include complete proofs of all theoretical results? [N/A] We do not include any theoretical results.
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes]
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [No]
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes] See Table 2.
   (b) Did you mention the license of the assets? [Yes] See Appendix C.1.
   (c) Did you include any new assets either in the supplemental material or as a URL? [Yes] We include our code in the supplemental material.
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [Yes] See Appendix C.1.
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [Yes] See Appendix C.1.
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]