

# Hybrid Predictive Models: When an Interpretable Model Collaborates with a Black-box Model

**Tong Wang**

*Tippie College of Business  
University of Iowa  
Iowa City, IA, USA*

TONG-WANG@UIOWA.EDU

**Qihang Lin**

*Tippie College of Business  
University of Iowa  
Iowa City, IA, USA*

QIHANG-LIN@UIOWA.EDU

**Editor:** Russ Greiner

## Abstract

Interpretable machine learning has become a strong competitor for black-box models. However, the possible loss of the predictive performance for gaining understandability is often inevitable, especially when it needs to satisfy users with diverse backgrounds or high standards for what is considered interpretable. This tension puts practitioners in a dilemma of choosing between high accuracy (black-box models) and interpretability (interpretable models). In this work, we propose a novel framework for building a Hybrid Predictive Model that integrates an interpretable model with *any* pre-trained black-box model to combine their strengths. The interpretable model substitutes the black-box model on a subset of data where the interpretable model is most competent, gaining transparency at a low cost of the predictive accuracy. We design a principled objective function that considers predictive accuracy, model interpretability, and model *transparency* (defined as the percentage of data processed by the interpretable substitute.) Under this framework, we propose two hybrid models, one substituting with association rules and the other with linear models, and design customized training algorithms for both models. We test the hybrid models on structured data and text data where interpretable models collaborate with various state-of-the-art black-box models. Results show that hybrid models obtain an efficient trade-off between transparency and predictive performance, characterized by *pareto frontiers*. Finally, we apply the proposed model on a real-world patients dataset for predicting cardiovascular disease and propose *multi-model Pareto frontiers* to assist model selection in real applications.

**Keywords:** interpretable machine learning, hybrid model, association rules, linear model, Pareto frontier

## 1. Introduction

As data from various fields have seen rapid growth in volume, variety, and velocity (Zikopoulos et al., 2012), machine learning models tend to grow bigger and more complicated in structure, especially with the outbreak of deep learning in the last decade. The predictive performance, however, is not obtained for free. Systems such as deep neural networks and ensemble models are *black-box* in nature, which means the models have a complex

and opaque decision-making process that is difficult for humans to understand. Therefore, in many heavily regulated industries such as judiciaries and healthcare, black-box models, despite their good predictive performance, often find it challenging to find their way into real-world deployment since domain experts are often reluctant to trust and adopt a model if they do not understand it. EU’s General Data Protection Regulation (GDPR) has called for the “right to explanation” (a right to information about individual decisions made by algorithms) (Parliament and of the European Union, 2016; Phillips, 2018) that requires human-understandable predictive processes of models.

To facilitate human understandability, various forms of interpretable models have been developed. For example, some models use a small number of cognitive chunks presented in a human-understandable way (e.g., linear, logical, etc.) to construct a model so users can comprehend how predictions are produced. Such models include rule-based models (Wang et al., 2017; Lakkaraju et al., 2017; Wang, 2018), sparse linear models (Zeng et al., 2017), case-based reasoning (Richter and Weber, 2016), etc. They are easy to understand, use, and improve compared to complex black-box models.

Yet, the benefits of interpretable machine learning models often come at a price of compromised predictive performance. While the state-of-the-art research in interpretable machine learning has pushed the limit of the models, and they can, sometimes, perform as well as black-box models, it is more often observed that interpretable models lose to black-box models on complicated tasks since they need to conform with domain-specific and user-specific need of interpretability. More importantly, in real-world practice, the level of model interpretability often varies a lot across different tasks and different users. In some domains, users only accept very simple models such as rule-base models with only a few rules, placing too much constraint on model complexity to achieve good predictive performance. In contrast, black-box models are trained to optimize a single objective, the predictive performance. In this paper, we focus on this more common situation where the best model is a black-box, and its competitors that satisfy user-defined interpretability only yield intolerable loss in predictive performance.

A popular solution to still gain interpretability in the above situation is to use black-box explainers, which generate post hoc explanations for a black-box model, either locally (Ribeiro et al., 2016) or globally (Adler et al., 2016; Lakkaraju et al., 2017). However, two concerning issues exist. First, explainers only approximate but do not characterize precisely the decision-making process of a black-box model, yielding an imperfect explanation fidelity. Second, ambiguity and inconsistency (Ross et al., 2017; Lissack, 2016) exist in explanations since different explanations can be generated for the same prediction, by different explainers or the same explainer with different parameters. Both issues result from the fact that explainers are not the decision-making process themselves and only approximate it afterward.

In this paper, we propose an alternative solution that introduces interpretability in a novel way. Our solution does not aim to explain a black-box but to substitute it on a subset of data where the interpretable model is comparably accurate. Therefore, the black-box model could be replaced by the simple model with no or low cost of predictive performance while gaining interpretability. If such an interpretable substitute can be found, we can design a hybrid predictive model where an input is first sent to the interpretable model to see if a prediction can be directly generated, if not, the black-box model will be activated. We call

the percentage of data processed by an interpretable model **transparency** of the model, which represents how often a prediction is interpretable. The goal of a hybrid model is to “squeeze” transparency from the decision-making process at a minimum cost of predictive performance.

To build a hybrid model, we need to address the following research questions. How to find such a sub-area in the feature space? What kind of interpretable models can be used as a local substitute for a complex model? And what is the form of the decision-making process? Finally, what does a system benefit from utilizing an interpretable local substitute?

To answer these questions, we design a general framework for hybrid predictive models and formulate a principled objective that considers predictive performance, the complexity of the interpretable substitute, and the transparency of the model. We instantiate this framework with the two most popular forms of interpretable models in this paper, association rules and linear models, to build two hybrid models with customized training algorithms. We apply both models to structured datasets and text data, where the interpretable models partially substitute state-of-the-art black-box models, including ensembles and neural networks. To evaluate the models, we use **pareto frontiers** that characterize the trade-off between transparency and accuracy. Pareto frontiers unify interpretable models, black-box models, and hybrid models, with black-box models located at transparency equal to zero, interpretable models at transparency equal to one, and hybrid models connecting the two extremes and spanning the entire spectrum of transparency. As shown by a user study in Pan et al. (2020), users’ choices of transparency and predictive performance vary greatly in terms of the tasks and models. Thus providing a flexible model that can be tuned to obtain different transparency and accuracy is important in real applications.

The proposed framework forms collaboration between an interpretable model and a black-box model, utilizing the strengths of them where they are most competent. Multi-model collaboration has been discussed in the previous literature (see related work), such as ensemble methods (Liaw et al., 2002) and gated experts systems (Berenji et al., 2004). Our framework is different in two aspects. First, the proposed hybrid model can work with **any** pre-trained black-box classifier. Second, our framework is **agnostic** to the black-box model, which means any implementational or algorithmic detail of the black-box model or even the type of model remains unknown during training. We only need the input and output of the black-box and the true labels of the input to train a hybrid model. This training mechanism will largely conceal information about the black-box system from its collaborator. This property is critical in building collaborations with proprietary systems that do not wish to share technical details. Our hybrid model framework can also be extended to human-model collaboration, where the black-box model can represent a human decision-maker by training on his prior data of decision making. Then we can construct an interpretable “assistant” to saving human labor and cost.

The rest of the paper is organized as follows. Section 2 reviews related work on interpretable machine learning. We present the framework of hybrid decision making in Section 3 and instantiate the model with association rules and linear models in Section 4 and Section 5, respectively, where we cast the two interpretable models into the hybrid framework and design customized training algorithms. In Section 6, we test the performance of the hybrid models on structured and text datasets from different domains where interpretabil-

ity is most pursued and discuss the results in detail. Section 7 discusses the application to predicting cardiovascular disease.

## 2. Related Work

### 2.1 Interpretable Models

Our work is broadly related to new methods for interpretable machine learning. An important line of research in interpretable machine learning is to develop models that are inherently interpretable and do not need external explainers; for example, models with a few cognitive chunks and simple logic. Previous works in this category include rule-based models such as rule sets (Wang et al., 2017; Rijnbeek and Kors, 2010; McCormick et al., 2011), rule lists (Yang et al., 2017; Angelino et al., 2017). Rule-based models check input against a set of ordered or unordered rules. If an instance is captured by a rule, the corresponding consequence of the rule is the prediction. Linear models such as sparse linear systems (Zeng et al., 2017; Ustun and Rudin, 2016; Koh et al., 2015) are also a canonical and popular form of interpretable models. They learn a possibly small set of non-zero weights for features and then compare the weighted sum of features with a threshold to determine the class labels. Other forms of linearly additive models are also interpretable, such as the generalized additive models (GAM) (Lou et al., 2012), which learn a shape function for each feature that represents how much a feature contributes to the output for different feature values. The shape functions can be learned via splines or trees. Recently, Agarwal et al. (2020) applies neural networks to learn each shape function and the model is called Neural Additive Model (NAM). The neural networks in a NAM are trained jointly and can learn arbitrarily complex relationships between their input feature and the output. NAM also utilizes the power of “black-box” models, but the presentation of the model is still interpretable.

When learning an interpretable model, the model complexity is regularized to improve interpretability and avoid overfitting. For rule-based models, model regularization often refers to reducing the number of rules, the number of features, the total number of conditions, etc (Wang, 2018; Lakkaraju et al., 2016). Sparse linear models regularize the number of non-zero coefficients or force the coefficients to be integers (Ustun and Rudin, 2016).

### 2.2 Black-box Explainers

For applications where an equally competent interpretable model cannot be found and a black-box model has to be chosen for higher accuracy, research in the past has developed various black-box explaining methods. These methods provide explanations, locally (Ribeiro et al., 2016) or globally (Adler et al., 2016; Lakkaraju et al., 2017) for a black-box model, providing insights by identifying key features, interactions of features (Tsang et al., 2017), etc. One of the pioneering work is LIME (Ribeiro et al., 2016) that explains the predictions of any classifier in an interpretable and faithful manner, by learning an interpretable linear model locally around the prediction. Others like Ribeiro et al. (2018); Lakkaraju et al. (2019) use association rules to explain a black-box system. More recently, research in deep learning has been connected strongly with interpretable machine learning (Goodfellow et al., 2014; Frosst and Hinton, 2017) and has contributed novel insights into representa-

tional issues. Despite the increasing interest in black-box explainers, there has been recent debate (Rudin, 2018) on potential issues of them since explainers only approximate but do not characterize precisely the decision-making process of a black-box model, often yielding an imperfect fidelity. In addition, there exists ambiguity and inconsistency (Ross et al., 2017; Lissack, 2016) in the explanation since there could be different explanations for the same prediction generated by different explainers, or by the same explainer with different parameters. Recent research discusses the ways to “fool” a black-box explainers (Lakkaraju and Bastani, 2020; Aivodji et al., 2019). These issues arise from the fact that the explainers only approximate in a post hoc way. They are not the decision-making process themselves.

Our work is fundamentally different from the research above. A hybrid predictive model is not a pure interpretable model. It utilizes the partial predictive power of black-box models to maximally preserve the predictive accuracy when gaining interpretability. It is also not an explainer model which only observes but does not participate in the decision process. Here, a hybrid predictive model strategizes the use of an interpretable model and a black-box model to make a prediction.

### 2.3 Model Hybrids

There exist a few singleton works that combine multiple models. One of the earliest works (Kohavi, 1996) proposed NBTree, which combines decision-tree classifiers and Naive-Bayes classifiers, Shin et al. (2000) introduced a system combining neural network and memory-based learning. Hua and Zhang (2006) combined SVM and logistic regression to forecast intermittent demand of spare parts, etc. A recent work (Wang et al., 2015) divides feature spaces into regions with sparse oblique tree splitting and assign local sparse additive experts to individual regions. Aside from the singletons, there has been a large body of continuous work on neural-symbolic or neural-expert systems (Garcez et al., 2015) pursued by a relatively small research community over the last two decades and has yielded several significant results (McGarry et al., 1999; Garcez et al., 2012; Taha and Ghosh, 1999; Towell and Shavlik, 1994). This line of research has been carried on to combine deep neural networks with expert systems to improve predictive performance (Hu et al., 2016).

Compared to the collaborative models discussed above, our model is distinct in that the proposed framework can work with *any* pre-trained black-box model. This black-box model could be a carefully calibrated, advanced model using confidential features or techniques. In addition, our framework is *agnostic* to the black-box model. Our training algorithm only needs predictions from the black-box model and does not need to alter the black-box during training or know any other information from it. This minimal requirement of information from the black-box collaborator renders much more flexibility in creating collaboration between different models, largely preserving confidential information from the more advanced partner.

## 3. Hybrid Predictive Model

We present a general framework for building a hybrid predictive model and formulate an objective function combining important properties of a predictive model. Let  $\mathcal{D} = \{(\mathbf{x}^{(n)}, y^{(n)})\}_{n=1}^N$  represent a set of training examples where  $\mathbf{x}^{(n)} \in \mathcal{X}$  is a set of attributes for instance  $n$  and  $y^{(n)} \in \mathcal{Y}$  is the target variable. In this paper we focus on binary classifi-

cation so  $y^{(n)} \in \{-1, 1\}$ . Let there be a black-box model  $f_b$  which can be any pre-trained model with unknown type. We are only provided with its predictions on the training set, denoted as  $\{\hat{y}_b^{(n)}\}_{n=1}^N$ .

Given  $\{(\mathbf{x}^{(n)}, y^{(n)}, \hat{y}_b^{(n)})\}_{n=1}^N$ , our goal is to build an interpretable model  $f_l$  that replaces  $f_b$  on a subset of data, forming a hybrid predictive model, denoted as  $f = \langle f_l, f_b \rangle$ . Let  $\mathcal{D}_l$  represent the subset of data processed by  $f_l$  and  $\mathcal{D}_b$  is the remaining data processed by  $f_b$ . We design the predictive process as below: an input  $\mathbf{x}$  is first sent to  $f_l$ . If a prediction can be generated,  $f_l$  will directly output a prediction  $\hat{y}_l$ , else,  $\mathbf{x}$  is sent to  $f_b$  to generate a prediction  $\hat{y}_b$ . See Figure 1 for an illustration.




Figure 1: A decision-making process of a hybrid predictive model.

We consider three properties critical when building a hybrid predictive model. *First*, we desire a model with high **predictive performance**. Since  $f_b$  is pre-given, the accuracy of a hybrid model is determined by two factors, the predictive accuracy of  $f_l$  and the collaboration of  $f_l$  and  $f_b$ , i.e., the partition of  $\mathcal{D}$  into  $\mathcal{D}_l$  and  $\mathcal{D}_b$ .  $f_l$  and  $f_b$  being completely different models allows the hybrid model to exploit the strengths of the two if the training examples are partitioned strategically, sending examples to the model which can obtain the maximum benefit. *Second*, we desire high **model interpretability** of  $f_l$ . Bringing interpretability into the decision-making process is one of the motivations of building a hybrid model. The definition of interpretability is model specific and usually refers to having a small size and low model complexity, often measured by the number of cognitive chunks (Doshi-Velez and Kim, 2017). For example, in rule-based models, the smallest cognitive chunks are conditions (feature-value pairs) in rules. The total number of conditions in a model represents the complexity of the model (Wang, 2018). For linear models, the smallest cognitive chunks are features, and therefore sparse linear models regularize the number of non-zero coefficients (Ning and Karypis, 2011). *Finally*, we want to maximize the amount of data sent to  $f_l$ , which represents how often a prediction is interpretable, i.e., captured by  $f_l$ . To do that we define a novel metric, **transparency**, which is the percentage of  $\mathcal{D}_l$  in  $\mathcal{D}$ .

**Definition 1** *The transparency of a hybrid model  $f = \langle f_l, f_b \rangle$  on  $\mathcal{D}$  is the percentage of data processed by  $f_l$ , i.e.,  $\frac{|\mathcal{D}_l|}{|\mathcal{D}|}$ , denoted as  $\mathcal{E}(f, \mathcal{D})$ .*

Following this definition, a stand-alone interpretable model has transparency equal to one, and a black-box model has transparency equal to zero.

In light of the considerations above, we formulate a learning objective  $\Lambda(f_l, \mathcal{D}, \{\hat{y}_b^{(n)}\}_{n=1}^N)$  as a linear combination of the three metrics. Given training data  $\mathcal{D}$ , predictions of a pre-trained black-box on  $\mathcal{D}$ ,  $\{\hat{y}_b^{(n)}\}_{n=1}^N$ , and a pre-defined form of interpretable models  $\mathcal{F}_l$ , our

goal is to find the optimal  $f_l^*$  that

$$f_l^* \in \arg \min_{f_l \in \mathcal{F}_l} \Lambda(f_l, \mathcal{D}, \{\hat{y}_b^{(n)}\}_{n=1}^N). \quad (1)$$

Next, we instantiate the proposed framework with selected interpretable substitutes. An interpretable substitute needs to fulfill two objectives, to produce a human-understandable prediction and to partition the data. We choose two interpretable models as examples, association rules and linear models, corresponding to the model space of  $\mathcal{F}_{\text{rule}}$  and  $\mathcal{F}_{\text{linear}}$ , respectively. Both are popular forms of models in the interpretable machine learning literature. In the next two sections, we will cast the two interpretable models into the proposed framework and develop customized training algorithms optimizing the objective (1).

## 4. Hybrid Rule Set

In this section, we choose association rules for  $f_l$  and call the model Hybrid Rule Set. We use  $\mathcal{F}_{\text{rule}}$  to denote the model space for  $f_l$ . Rules are easy to understand for their simple logic and symbolic presentation. They also naturally handle the partition of data by separating examples according to if they satisfy the rules.

### 4.1 Model Formulation

We construct two rule sets for  $f_l$ , a *positive rule set*,  $\mathcal{R}_+$ , and a *negative rule set*,  $\mathcal{R}_-$ . If  $\mathbf{x}^{(n)}$  satisfies any rules in  $\mathcal{R}_+$ , it is classified as positive. Otherwise, if it satisfies any rules in  $\mathcal{R}_-$ , it is classified as negative. A decision produced from rules is denoted as  $\hat{y}_l^{(n)}$ . If  $\mathbf{x}^{(n)}$  does not satisfy any rules in  $\mathcal{R}_+$  or  $\mathcal{R}_-$ , it means  $f_l$  fails to decide on  $\mathbf{x}^{(n)}$ . Then  $\mathbf{x}^{(n)}$  is sent to the black-box model  $f_b$  to generate a decision  $\hat{y}_b^{(n)}$ . The predictive process of  $f$  is shown below.

$$\begin{aligned} &\text{if } \mathbf{x}^{(n)} \text{ obeys } \mathcal{R}_+, Y = 1 \\ &\text{else if } \mathbf{x}^{(n)} \text{ obeys } \mathcal{R}_-, Y = -1 \\ &\text{else } Y = f_b(\mathbf{x}) \end{aligned} \quad (2)$$

Here  $\mathcal{R}_+$  and  $\mathcal{R}_-$  are embedded in an “else if” logic. If an instance satisfies both  $\mathcal{R}_+$  and  $\mathcal{R}_-$ , it is classified as positive since  $\mathcal{R}_+$  is ranked higher. Define  $\mathcal{R} = (\mathcal{R}_+, \mathcal{R}_-)$ . We will use  $\mathcal{R}$  and  $f_l$  interchangeably for hybrid rule sets.

The association rules not only produce predictions but also naturally partition the feature space into  $\mathcal{D}_l$  and  $\mathcal{D}_b$ , covering a subset of positive instances with  $\mathcal{R}_+$  and a subset of negative instances with  $\mathcal{R}_-$ .

We define the coverage of rules below.

**Definition 2** A rule  $r$  covers  $\mathbf{x}^{(n)}$  if  $\mathbf{x}^{(n)}$  obeys the rule, denoted as  $\text{covers}(r, \mathbf{x}^{(n)}) = 1$ .

We say an instance  $\mathbf{x}^{(n)}$  obeys a rule  $r$  if  $\mathbf{x}^{(n)}$  satisfies all conditions in  $r$ .

**Definition 3** A set of rules  $R$  covers  $\mathbf{x}^{(n)}$  if  $\mathbf{x}^{(n)}$  obeys at least one rule in  $R$ , i.e.

$$\text{covers}(R, \mathbf{x}^{(n)}) = \mathbb{1} \left( \sum_{r \in R} \text{covers}(r, \mathbf{x}^{(n)}) \geq 1 \right). \quad (3)$$

**Definition 4** Given a data set  $\mathcal{D}$ , the support of a rule set  $\mathcal{R}$  in  $\mathcal{D}$  is the number of observations covered by  $\mathcal{R}$ , i.e.,

$$\text{support}(\mathcal{R}, \mathcal{D}) = \sum_{n=1}^N \text{covers}(\mathcal{R}, \mathbf{x}^{(n)}). \quad (4)$$

#### LEARNING OBJECTIVE

Now we cast learning objective  $\Lambda(f_t, \mathcal{D}, \{\hat{y}_b^{(n)}\}_{n=1}^N)$  to Hybrid Rule Sets. First, we measure the predictive performance with misclassification error, represented  $\ell(\langle \mathcal{R}, f_b \rangle, \mathcal{D})$ .

$$\begin{aligned} \ell(\langle \mathcal{R}, f_b \rangle, \mathcal{D}) &= \sum_{n=1}^N \left( \frac{(1 - y^{(n)})}{2} \text{covers}(\mathcal{R}_+, \mathbf{x}^{(n)}) \quad \triangleright \text{errors from } \mathcal{R}_+ \right. \\ &+ \frac{(1 + y^{(n)})}{2} \left( 1 - \text{covers}(\mathcal{R}_+, \mathbf{x}^{(n)}) \right) \text{covers}(\mathcal{R}_-, \mathbf{x}^{(n)}) \quad \triangleright \text{errors from } \mathcal{R}_- \\ &+ ((1 - \text{covers}(\mathcal{R}_+, \mathbf{x}^{(n)}))(1 - \text{covers}(\mathcal{R}_-, \mathbf{x}^{(n)})) \quad \triangleright \text{instances not covered by } \mathcal{R}_+ \text{ or } \mathcal{R}_- \\ &\times \left. \left( \frac{(1 + y^{(n)})}{2} \frac{(1 - \hat{y}_b^{(n)})}{2} + \frac{(1 - y^{(n)})}{2} \frac{(1 + \hat{y}_b^{(n)})}{2} \right) \right) / N. \quad \triangleright \text{errors from } f_b \end{aligned} \quad (5)$$

Next, we measure the interpretability of  $\mathcal{R}$ , denoted as  $\Omega(\mathcal{R})$ . There exists a list of choices from the literature (Wang et al., 2017; Lakkaraju et al., 2016; Wang, 2018). In this paper, we use the total number of conditions in rules, denoted as  $\Omega(\mathcal{R})$ . Regularizing the number of conditions not only increases interpretability but also naturally avoids overfitting. Without this regularization, the model will find long rules (having many conditions in a rule) with very small support but highly accurate, causing overfitting. Therefore, regularizing the number of conditions will benefit the model in both model interpretability and predictive performance. The design of interpretability is customizable. Other metrics such as the overlap of rules, lengths of rules, etc., can also be used here.

Finally, we compute the transparency of a model. Following Definition 1, the transparency of a hybrid rule set  $\langle \mathcal{R}, f_b \rangle$  is

$$\mathcal{E}(\langle \mathcal{R}, f_b \rangle, \mathcal{D}) = \frac{\text{support}(\mathcal{R})}{|\mathcal{D}|}. \quad (6)$$

We will write  $\ell(\langle \mathcal{R}, f_b \rangle, \mathcal{D})$  and  $\mathcal{E}(\langle \mathcal{R}, f_b \rangle, \mathcal{D})$  as  $\ell(\mathcal{R})$  and  $\mathcal{E}(\mathcal{R})$ , respectively, ignoring the dependence on  $\mathcal{D}$  and  $f_b$  for simpler notations since  $f_b$  and  $\mathcal{D}$  are fixed.

Thus, we derive an objective function combining predictive accuracy, model interpretability and transparency

$$\Lambda(\mathcal{R}) = \ell(\mathcal{R}) + \alpha_1 \Omega(\mathcal{R}) - \alpha_2 \mathcal{E}(\mathcal{R}), \quad (7)$$

and our goal is to find an optimal model  $\mathcal{R}^*$  such that

$$\mathcal{R}^* = (\mathcal{R}_+^*, \mathcal{R}_-^*) \in \arg \min_{\mathcal{R}} \Lambda(\mathcal{R}). \quad (8)$$

Here,  $\alpha_1, \alpha_2$  are non-negative coefficients. Tuning the parameters will produce models with different accuracy, interpretability, and transparency. For example, in an extreme case



when  $\alpha_2 \gg \alpha_1$ , the output will be a model that sends all data to  $f_l$ , producing a pure interpretable model. Then increasing  $\alpha_1$  will produce a simpler and sparser model, using fewer rules and fewer conditions. When  $\alpha_1 \gg \alpha_2$ , the model will force  $f_l$  to have no conditions and no rules, i.e., producing a pure black-box model.

**Producing Probabilistic Outputs** For dealing with more complicated predictive tasks, our model can be extended to produce probabilities by making small changes to the model formulation proposed above. To do that, we need the probabilistic outputs from the black-box model as well, represented by  $\{p_b(\mathbf{x}^{(n)})\}_{n=1}^N$  where  $p_b(\mathbf{x}^{(n)}) \in [0, 1]$  and  $p_b(\mathbf{x}^{(n)})$  represent the probability that instance  $\mathbf{x}^{(n)}$  is in the positive class. To provide a probability estimation for instances that are captured by rules, we follow the conventional way used in probability estimation in decision trees, that the probability of positive is the percentage of instances that are positive out of the total number of instances captured by the rule, which means,

$$\hat{p}_l(r) = \frac{\sum_{n=1}^N \text{covers}(r, \mathbf{x}^{(n)}) \mathbb{1}(y^{(n)} = 1)}{\sum_{n=1}^N \text{covers}(r, \mathbf{x}^{(n)})}, \text{ if } r \in \mathcal{R}_+ \quad (9)$$

If  $r \in \mathcal{R}_-$ , the coverage of  $r$  needs to exclude the coverage of  $\mathcal{R}_+$ , if there's any overlap. Thus

$$\hat{p}_l(r) = \frac{\sum_{n=1}^N (1 - \text{covers}(\mathcal{R}_+, \mathbf{x}^{(n)})) \text{covers}(r, \mathbf{x}^{(n)}) \mathbb{1}(y^{(n)} = 1)}{\sum_{n=1}^N (1 - \text{covers}(\mathcal{R}_+, \mathbf{x}^{(n)})) \text{covers}(r, \mathbf{x}^{(n)})}, \text{ if } r \in \mathcal{R}_-. \quad (10)$$

If an instance satisfies a rule  $r$ , then output  $\hat{p}_l(r)$ . In the case where an instance satisfies multiple rules, the instance is assigned to the positive rules, following the ‘‘if’’ logic. Then, if an instance satisfies multiple rules of the same sign, it is always assigned the probability of the most confident rule, i.e., the rule with the largest support. If instance  $\mathbf{x}^{(n)}$  is captured by the black-box, then the model outputs  $p_b(\mathbf{x}^{(n)})$ .

The loss function can be modified accordingly by changing minimizing error to minimizing negative AUC on the training set without changing anything else to the rest of the algorithm.

## 4.2 Training with Adaptive Stochastic Local Search

We describe a training algorithm to find an optimal solution  $\mathcal{R}^*$  to objective (8). Learning rule-based models is challenging because the solution space is a power set of the rule space. Fortunately, our objective function suggests theoretical bounds that can be exploited for reducing computation. We first present the algorithm structure, then detail the theoretically grounded strategies we embed into the algorithm, and finally describe the proposing step.

**Algorithm Structure** Our algorithm adopts the main structure of simulated annealing (Kirkpatrick et al., 1983). Given training examples,  $\mathcal{D}$ , black-box predictions  $\{\hat{y}_b^{(n)}\}_{n=1}^N$ , parameters  $\alpha_1, \alpha_2$ , base temperature  $C_0$  and the total number of iterations  $T$ , we design an adaptive stochastic local search algorithm. Each state corresponds to a tuple of a positive rule set and a negative rule set, indexed by the time stamp  $t$ , denoted as  $\mathcal{R}^{[t]} = (\mathcal{R}_+^{[t]}, \mathcal{R}_-^{[t]})$ . The starting state  $\mathcal{R}^{[0]}$  is initialized with two empty sets for  $\mathcal{R}_+^{[0]}$  and  $\mathcal{R}_-^{[0]}$ , respectively. The temperature is a function of time  $t$ ,  $C_0^{1-\frac{t}{T}}$  and it decreases with time. The neighboring

states are defined as rule sets that are obtained via a small change to the current model. At each iteration, the algorithm improves one of the three terms (accuracy, interpretability, and transparency) with approximately equal probabilities, by removing a rule from the current model or adding a rule to the current model. The proposed neighbor is then accepted with probability  $\exp(\frac{\Lambda(\mathcal{R}^{[t]})-\Lambda(\mathcal{R}^{[t+1]})}{C_0^{1-\frac{t}{T}}})$  which gradually decreases as the temperature cools down till eventually converging to the final output.

Next, we introduce the strategies and theoretical bounds to prune the search space before and during the search.

**Rule Space Pruning** We first derive a bound to prune the search space before the algorithm begins. Assume the rules in  $\mathcal{R}^*$  are drawn from candidate rules. We use FP-growth<sup>1</sup> to generate two sets of candidate rules,  $\Upsilon^+$  for positive rules and  $\Upsilon^-$  for negative rules. The algorithm will search within this rule space, so the complexity of the algorithm is directly determined by the size of the rule space. To reduce the search space, we derive a lower bound on the support of rules in  $\mathcal{R}^*$ .

**Theorem 5 (Bound on Support)**  $\forall r \in \mathcal{R}_+^*, support(r) \geq N\alpha_1; \forall r \in \mathcal{R}_-^*, support(r) \geq \frac{N\alpha_1}{1-\alpha_2}$ .

This means  $\mathcal{R}^*$  does not contain rules with too small a support. Rules that fail to satisfy the bounds can be safely deleted from the rule space without hurting the optimality. This theorem is used before the search begins to prune the rule space, which greatly reduces computation. On the other hand, removing rules with low support naturally helps prevent overfitting. The bound increases as  $\alpha_1$  increases, since  $\alpha_1$  represents the penalty for adding a rule. When  $\alpha_1$  is large, the model is encouraged to use fewer rules, thus each rule need to cover enough instances in order to qualify for a spot in the model, therefore they are encouraged to have large support. All proofs are in the supplementary material.

**Search Chain Bounding** We also derive two bounds that dynamically prune the search space during the search. The bounds are applied in each iteration to confine the Markov Chain within promising solution space. We first derive a bound on  $\Omega(\mathcal{R}^*)$ , the total number of conditions in  $\mathcal{R}^*$ . Let  $\lambda^{[t]}$  represent the best objective value found till time  $t$ , i.e.

$$\lambda^{[t]} = \min_{\tau \leq t} \Lambda(\mathcal{R}_{[\tau]}).$$

We claim

**Theorem 6 (Bound on Interpretability)**  $\Omega(\mathcal{R}^*) \leq \frac{\lambda^{[t]} + \alpha_2}{\alpha_1}$ .

This theorem says that the number of conditions in  $\mathcal{R}^*$  is upper bounded. The theorem implies that the Markov Chain only needs to focus on the solution space of small models. Therefore, in the proposing step, if the current state violates the bound, neighbors with more conditions should not be proposed, equivalent to pruning the neighbors.

Next we derive an upper bound on the transparency with the similar purpose of confining the Markov Chain.

**Theorem 7 (Bound on Transparency)**  $support(\mathcal{R}^*) \geq \frac{N(\alpha_1 - \lambda^{[t]})}{\alpha_2}$ .

---

1. FP-Growth is an off-the-shelf rule miner. Other rule miners such as Apriori or Eclat can also be used.

The theorem says the transparency of  $\mathcal{R}^*$  is lower bounded. Therefore when proposing the next state, if this theorem is violated, models with smaller transparency should not be proposed.

---

**Algorithm 1** Stochastic Local Search algorithm
 

---

```

1: Input:  $\{\hat{y}_b^{(n)}\}_{n=1}^N, \mathcal{D}, \alpha_1, \alpha_2, C_0$ 
2: Initialize:
3:  $\mathcal{R}^* = \mathcal{R}^{[0]} \leftarrow \emptyset$   $\triangleright$  start with a pure black-box model
4:  $\Upsilon_+ \leftarrow \text{FPGrowth}(\mathcal{D}, \text{minsupp} = N\alpha_1)$ 
5:  $\Upsilon_- \leftarrow \text{FPGrowth}(\mathcal{D}, \text{minsupp} = \frac{N\alpha_1}{1-\alpha_2})$ 
    $\triangleright$  mine candidate rules from  $\mathcal{D}$  using minimum support from Theorem 5
6: for  $t = 0 \rightarrow T$  do
7:    $\delta = \text{random}()$ 
8:   if  $\delta \leq \frac{1}{3}$  or  $\Omega(\mathcal{R}^{[t]}) \geq \frac{\lambda^{[t]} + \alpha_2}{\alpha_1}$  then
9:      $\mathcal{R}^{[t+1]} \leftarrow$  remove a rule from  $\mathcal{R}^{[t]}$   $\triangleright$  decrease the size of  $\mathcal{R}^{[t]}$  and use Theorem 6
10:  else if  $\delta \leq \frac{2}{3}$  or  $\text{support}(r)(\mathcal{R}^{[t]}) \leq \frac{N(\alpha_1 - \lambda^{[t]})}{\alpha_2}$  then
11:     $\mathcal{R}^{[t+1]} \leftarrow$  add a rule to  $\mathcal{R}$   $\triangleright$  increase the transparency and use Theorem 7
12:  else
13:     $\epsilon \leftarrow \{k | f^{[t]}(\mathbf{x}^{(k)}) \neq y^{(k)}\}$   $\triangleright$  indices of misclassified examples.
14:    if  $\text{covers}(\mathcal{R}^{[t]}, \mathbf{x}^{(\epsilon)}) = 1$  then
15:      if  $\epsilon$  is negative then
16:         $\mathcal{R}^{[t+1]} \leftarrow$  remove a rule from  $\mathcal{R}_+^{[t]}$  that covers instance  $\epsilon$ .
17:      else
18:         $\mathcal{R}^{[t+1]} \leftarrow$  add a rule to  $\mathcal{R}_+^{[t]}$  to cover  $\epsilon$  or remove a rule from  $\mathcal{R}_-^{[t]}$  that covers  $\epsilon$ 
19:      end if
20:    else
21:      if  $\epsilon$  is positive then
22:         $\mathcal{R}^{[t+1]} \leftarrow$  add a rule to  $\mathcal{R}_+^{[t]}$  and  $\mathcal{R}_-^{[t]}$  remains the same
23:      else
24:         $\mathcal{R}^{[t+1]} \leftarrow$  add a rule to  $\mathcal{R}_-^{[t]}$  and  $\mathcal{R}_+^{[t]}$  remains the same
25:      end if
26:    end if
27:  end if
28:   $\mathcal{R}^{[t+1]}$  is accepted with probability  $\exp(\frac{\Lambda(\mathcal{R}^{[t]}) - \Lambda(\mathcal{R}^{[t+1]})}{C_0^{\frac{t}{T}}})$   $\triangleright$  accept the proposal with
   certain probability
29:   $\mathcal{R}^* = \arg \min_{\mathcal{R}^{[t+1]}, \mathcal{R}^*} \Lambda(\mathcal{R})$   $\triangleright$  update the best solution
30: end for
31: Output:  $\mathcal{R}^*$ 

```

---

Both bounds in [Theorem 6](#) and [Theorem 7](#) become smaller as  $\lambda^{[t]}$  continuously gets smaller along the search steps. Exploiting the theorems in the search algorithm, the algorithm checks the intermediate solution at each iteration and pull the search chain back to the promising area (models of sizes smaller than  $\frac{\lambda^{[t]} + \alpha_2}{\alpha_1}$  and transparency larger than

$\frac{N(\alpha_1 - \lambda^{[t]})}{\alpha_2}$ ) whenever the bounds are violated, equivalent to pruning the search space at each iteration.

The algorithm is presented in Algorithm 1<sup>2</sup>. Here we detail the proposing step. To propose a neighbor, at each iteration, we choose to improve one of the three terms (accuracy, interpretability, and transparency) with approximately equal probabilities. With probability  $\frac{1}{3}$ , or when the bound in Theorem 6 is violated, we aim to decrease the size of  $\mathcal{R}^{[t]}$  (improve interpretability) by removing a rule from  $\mathcal{R}^{[t]}$  (line 8 - 9). With probability  $\frac{1}{3}$  or when the lower bound on transparency (Theorem 7) is violated, we aim to increase the coverage of  $\mathcal{R}^{[t]}$  (improve transparency) by adding a rule to  $\mathcal{R}^{[t]}$  (line 10-11). Finally, with another probability  $\frac{1}{3}$ , we aim to decrease the classification error (improve accuracy) (line 13-26). To decrease the misclassification error, at each iteration, we draw an example from examples misclassified by the current model (line 13). If the example is covered by  $\mathcal{R}^{[t]}$  (line 14), it means it was sent to the interpretable model but was covered by the wrong rule set. If the instance is negative (line 14), we remove a rule from the positive rule set that covers it. If the instance is positive, it means it is covered by the negative rule set. Then we either add a rule to the positive set to cover it or remove a rule in the negative rule set that covers it, re-routing it to  $f_b$ . If the example is not covered by  $\mathcal{R}^{[t]}$  (line 20), it means it was previously sent to the black-box model but misclassified, since we cannot alter the black-box model, we add a rule to the positive or negative rule set (consistent with the label of the instance) to cover the example. When choosing a rule to add or remove, we first evaluate the rules using precision, which is the percentage of correctly classified examples in the rule. Then we balance between exploitation, choosing the best rule, and exploration, selecting a random rule, to avoid getting stuck in a local minimum.

## 5. Hybrid *Linear Models*

Next, we instantiate the hybrid framework with linear models and call the model Hybrid Linear Model. Our goal is to learn  $f_\iota = f_{\text{linear}} = (\mathbf{w}, \theta_+, \theta_-)$ , where  $\mathbf{w}$  is the coefficient of a linear model and  $\theta_+$  and  $\theta_-$  ( $\theta_+ \geq \theta_-$ ) are two thresholds, in order to form  $f = \langle f_{\text{linear}}, f_b \rangle$  as:

$$f(\mathbf{x}) = \begin{cases} 1 & \text{if } \mathbf{w}^\top \mathbf{x} \geq \theta_+ \\ -1 & \text{if } \mathbf{w}^\top \mathbf{x} \leq \theta_- \\ f_b(\mathbf{x}) & \text{otherwise} \end{cases} \quad (11)$$

### 5.1 Model Formulation

We cast the optimization for a hybrid linear model into the hybrid decision-making framework and re-formulate the learning objective.

#### LEARNING OBJECTIVE

The (in-sample) predictive performance characterizes the fitness of the model to the training data. Since  $f_b$  is pre-given, the predictive performance is determined by two factors, the accuracy of  $f_{\text{linear}}$  on instances with  $\mathbf{w}^\top \mathbf{x} \geq \theta_+$  or  $\mathbf{w}^\top \mathbf{x} \leq \theta_-$  and the accuracy of  $f_b$  on

2. Code is available at <https://github.com/wangtongada/HyRS>

the remaining examples. We wish to obtain a good partition of data  $\mathcal{D}$  by assigning  $f_b$  and  $f_{\text{linear}}$  to different region of the data such that the strength of  $f_b$  and  $f_{\text{linear}}$  are properly exploited. Second, we include the gap  $\theta_+ - \theta_-$  as a penalty term in the objective to account for transparency of a hybrid linear model. The smaller gap between  $\theta_+$  and  $\theta_-$  there is, more data are classified by the linear model. In the most extreme case where  $\theta_+ - \theta_- = 0$ , all data are sent to the linear model, and the hybrid linear model is reduced to a pure linear classifier, i.e., transparency equals one. Finally, we also need to consider model regularization in the objective. As the weight for the sparsity enforcing regularization term increases, the model encourages using a smaller number of features which increases the interpretability of the model as well as preventing overfitting.

Combining the three factors discussed above, we formulate the learning objective for hybrid linear model as:

$$F^* := \min_{\mathbf{w}, \theta_+ \geq \theta_-} \{F(\mathbf{w}, \theta_+, \theta_-) := \mathcal{L}(\mathbf{w}, \theta_+, \theta_-; f_b) + \alpha_1 r(\mathbf{w}) + \alpha_2(\theta_+ - \theta_-)\}, \quad (12)$$

where  $\mathcal{L}(\mathbf{w}, \theta_+, \theta_-; f_b)$  is the loss function defined on the training set,  $\theta_+ - \theta_-$  is a penalty term to increase the transparency of  $f$ ,  $r$  is a convex and closed regularization term (e.g.  $\|\mathbf{w}\|_1$ ,  $\frac{1}{2}\|\mathbf{w}\|_2^2$  or an indicator function of a constraint set), and  $\alpha_1 \in \mathbb{R}_+$  and  $\alpha_2 \in \mathbb{R}$  are coefficients which balance the importance of the three components in (12).

Next, we define  $\mathcal{L}(\mathbf{w}, \theta_+, \theta_-; f_b)$  in details. According to (11), the hybrid model  $f$  will misclassify an instance  $\mathbf{x}^{(n)}$  in two scenarios. When  $y_b^{(n)} = 1$ ,  $\mathbf{w}^\top \mathbf{x} = \theta_-$  becomes the active decision boundary in the sense that  $\mathbf{x}^{(n)}$  with  $\mathbf{w}^\top \mathbf{x}^{(n)} \geq \theta_-$  is predicted as positive (if  $\mathbf{w}^\top \mathbf{x}^{(n)} \geq \theta_+$ , it is predicted as positive by the linear model; if  $\theta_- \leq \mathbf{w}^\top \mathbf{x}^{(n)} \leq \theta_+$ , it is predicted as positive by the black-box since  $y_b^{(n)} = 1$ ) and  $\mathbf{x}^{(n)}$  with  $\mathbf{w}^\top \mathbf{x} < \theta_-$  is classified as negative according to (11). When  $y_b^{(n)} = -1$ ,  $\mathbf{w}^\top \mathbf{x} = \theta_+$  becomes the active decision boundary in the sense that  $\mathbf{x}^{(n)}$  with  $\mathbf{w}^\top \mathbf{x}^{(n)} \geq \theta_+$  is predicted as positive and  $\mathbf{x}^{(n)}$  with  $\mathbf{w}^\top \mathbf{x}^{(n)} < \theta_+$  is classified as negative according to (11) (if  $\mathbf{w}^\top \mathbf{x}^{(n)} \leq \theta_-$ , it is predicted as negative by the linear model; if  $\theta_- \leq \mathbf{w}^\top \mathbf{x}^{(n)} \leq \theta_+$ , it is predicted as negative by the black-box since  $y_b^{(n)} = -1$ ). See Figure 2 for an illustration of the decision boundaries.

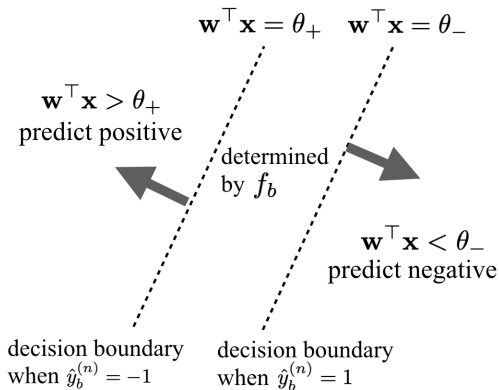


Figure 2: The partitioning of the data space by two decision boundaries

To define the losses associated to  $f$  in these two scenarios, we partition the data into two subsets based on the labels predicted by the black-box  $\{y_b^{(n)}\}_{n=1}^N$ . In particular, we define

$$I_b^+ = \{n | y_b^{(n)} = 1\} \text{ and } I_b^- = \{n | y_b^{(n)} = -1\}. \quad (13)$$

The loss function in (12) over the dataset  $\mathcal{D}$  is then defined as

$$\mathcal{L}(\mathbf{w}, \theta_+, \theta_-; f_b) = \frac{1}{N} \sum_{n \in I_b^+} \phi(y^{(n)}(\mathbf{w}^\top \mathbf{x}^{(n)} - \theta_-)) + \frac{1}{N} \sum_{n \in I_b^-} \phi(y^{(n)}(\mathbf{w}^\top \mathbf{x}^{(n)} - \theta_+)), \quad (14)$$

where  $\phi(z) : \mathbb{R} \rightarrow \mathbb{R}$  is a non-increasing convex closed loss function which can be one of those commonly used in linear classification such as

$$\begin{cases} \text{hinge loss} & \phi(z) = (1 - z)_+ \\ \text{smoothed hinge loss} & \phi(z) = \frac{1}{2}(1 - z)_+^2 \\ \text{logistic loss} & \phi(z) = \log(1 + \exp(-z)). \end{cases} \quad (15)$$

Note that  $I_b^+$  and  $I_b^-$  form a partition of  $\{1, 2, \dots, N\}$  and each data point corresponds to one loss term in (14). The intuition of this loss function is as follows. Take a data point  $\mathbf{x}^{(n)}$  with  $y^{(n)} = 1$  and  $y_b^{(n)} = 1$  as an example. Our hybrid model (11) will classify  $\mathbf{x}^{(n)}$  correctly as long as  $\mathbf{x}^{(n)}$  does not fall in the “negative half space”, namely,  $\mathbf{w}^\top \mathbf{x}^{(n)} > \theta_-$ . Hence, with the non-increasing property of  $\phi$ , the loss term  $\phi(y^{(n)}(\mathbf{w}^\top \mathbf{x}^{(n)} - \theta_-))$  will encourage a positive value of  $\mathbf{w}^\top \mathbf{x}^{(n)} - \theta_-$ . On the other hand, for a data point  $\mathbf{x}^{(n)}$  with  $y^{(n)} = 1$  and  $y_b^{(n)} = -1$ , our hybrid model will classify  $\mathbf{x}^{(n)}$  correctly only when  $\mathbf{x}^{(n)}$  falls in the “positive half space”, namely,  $\mathbf{w}^\top \mathbf{x}^{(n)} \geq \theta_+$ . We use the loss term  $\phi(y^{(n)}(\mathbf{w}^\top \mathbf{x}^{(n)} - \theta_+))$  to encourage a positive value of  $\mathbf{w}^\top \mathbf{x}^{(n)} - \theta_+$ . The true label  $y^{(n)} = 1$  in both examples above but the similar interpretation applies to the case of  $y^{(n)} = -1$ .

The loss function (14) is defined by assuming the black-box model has provided binary labels  $\{y_b^{(n)}\}_{n=1}^N$  with each  $y_b^{(n)} \in \{1, -1\}$ . However, it is possible that the black-box model only provides “soft” labels, namely, a set of probabilities  $\{p_b^{(n)}\}_{n=1}^N$  with  $p_b^{(n)}$  representing the predicted probability that  $\mathbf{x}^{(n)}$  is in the positive class. If that is the case, we can easily modify the loss function (14) using  $p_b^{(n)}$  and  $1 - p_b^{(n)}$  as the soft membership of  $\mathbf{x}^{(n)}$  in sets  $I_b^+$  and  $I_b^-$ . In particular, we can obtain the following loss function which is the weighted version of (14):

$$\mathcal{L}(\mathbf{w}, \theta_+, \theta_-; f_b) = \frac{1}{N} \sum_{n=1}^N p_b^{(n)} \phi(y^{(n)}(\mathbf{w}^\top \mathbf{x}^{(n)} - \theta_-)) + \frac{1}{N} \sum_{n=1}^N (1 - p_b^{(n)}) \phi(y^{(n)}(\mathbf{w}^\top \mathbf{x}^{(n)} - \theta_+)). \quad (16)$$

**Producing Probabilistic Outputs** Similar to the hybrid rule set method introduced in Section 4, the hybrid linear model can also produce the predicted probabilities of instances being in the positive and negative classes. Let  $\mathbf{x}$  be a testing instance. Let  $p_b(\mathbf{x})$  be the predicted probability of  $\mathbf{x}$  being positive by the black-box model. According to the rule (11), it is intuitive to use  $\mathbf{w}^\top \mathbf{x} = (\theta_+ + \theta_-)/2$ , namely, the hyper-plane right in the middle of the two linear decision boundaries, as the decision boundary of a purely linear model. A standard approach to produce probabilistic outputs from a purely linear model is to apply

---

**Algorithm 2** Smoothing Accelerated Proximal Gradient Method for (12) (Nesterov, 2007, 2013)

---

- 1: **Input:** The initial solution  $\mathbf{w}^{[0]} \in \text{dom}(r)$ ,  $(\theta_+^{[0]}, \theta_-^{[0]})$  satisfying  $\theta_+^{[0]} \geq \theta_-^{[0]}$ , step length parameters  $\eta > 0$  and  $\alpha_0 \in (0, 1)$ , the smoothing parameter  $\mu \geq 0$  ( $\mu = 0$  when  $\phi$  is smooth) and the number of iterations  $T$ .
- 2: Set  $(\widehat{\mathbf{w}}^{[0]}, \widehat{\theta}_+^{[0]}, \widehat{\theta}_-^{[0]}) = (\mathbf{w}^{[0]}, \theta_+^{[0]}, \theta_-^{[0]})$
- 3: **for**  $t = 1, 2, \dots, T$  **do**
- 4:

$$\begin{aligned} \mathbf{w}^{[t]} &= \arg \min_{\mathbf{w} \in \mathbb{R}^d} \left\langle \nabla_{\mathbf{w}} \mathcal{L}_\mu(\widehat{\mathbf{w}}^{[t-1]}, \widehat{\theta}_+^{[t-1]}, \widehat{\theta}_-^{[t-1]}), \mathbf{w} \right\rangle + \frac{1}{2\eta} \left\| \mathbf{w} - \widehat{\mathbf{w}}^{[t-1]} \right\|_2^2 + \alpha_1 r(\mathbf{w}) \\ (\theta_+^{[t]}, \theta_-^{[t]}) &= \arg \min_{\theta_+ \geq \theta_-} \left\{ \begin{array}{l} [\nabla_{\theta_+} \mathcal{L}_\mu(\widehat{\mathbf{w}}^{[t-1]}, \widehat{\theta}_+^{[t-1]}, \widehat{\theta}_-^{[t-1]}) + \alpha_2] \theta_+ \\ + [\nabla_{\theta_-} \mathcal{L}_\mu(\widehat{\mathbf{w}}^{[t-1]}, \widehat{\theta}_+^{[t-1]}, \widehat{\theta}_-^{[t-1]}) - \alpha_2] \theta_- \\ + \frac{1}{2\eta} \left\| \begin{pmatrix} \theta_+ \\ \theta_- \end{pmatrix} - \begin{pmatrix} \widehat{\theta}_+^{[t-1]} \\ \widehat{\theta}_-^{[t-1]} \end{pmatrix} \right\|_2^2 \end{array} \right\} \end{aligned}$$

- 5: Compute  $\alpha_t \in (0, 1)$  from the equation  $\alpha_t^2 = (1 - \alpha_t)\alpha_{t-1}^2$  and set  $\beta_t = \frac{\alpha_{t-1}(1 - \alpha_{t-1})}{\alpha_{t-1}^2 + \alpha_t}$ .
- 6:

$$\begin{pmatrix} \widehat{\mathbf{w}}^{[t]} \\ \widehat{\theta}_+^{[t]} \\ \widehat{\theta}_-^{[t]} \end{pmatrix} = \begin{pmatrix} \mathbf{w}^{[t]} \\ \theta_+^{[t]} \\ \theta_-^{[t]} \end{pmatrix} + \beta_t \left[ \begin{pmatrix} \mathbf{w}^{[t]} \\ \theta_+^{[t]} \\ \theta_-^{[t]} \end{pmatrix} - \begin{pmatrix} \mathbf{w}^{[t-1]} \\ \theta_+^{[t-1]} \\ \theta_-^{[t-1]} \end{pmatrix} \right]$$

7: **end for**

- 8: **Output:**  $\mathbf{w}^{[t]}, \theta_+^{[t]}, \theta_-^{[t]}$
- 

the logistic transformation to the linear formula. Hence, when  $\mathbf{w}^\top \mathbf{x} \geq \theta_+$  or  $\mathbf{w}^\top \mathbf{x} \leq \theta_-$ , namely, when  $\mathbf{x}$  is assigned to the linear model, we use  $\frac{\exp(\mathbf{w}^\top \mathbf{x} - (\theta_+ + \theta_-)/2)}{1 + \exp(\mathbf{w}^\top \mathbf{x} - (\theta_+ + \theta_-)/2)}$  as the predicted probability of  $\mathbf{x}$  being positive. When  $\mathbf{x}$  is assigned to the black-box model, we still use  $p_b(\mathbf{x})$  as the output probability. As a result, we produce the predicted probability of being positive by the hybrid linear model as

$$\begin{cases} \frac{\exp(\mathbf{w}^\top \mathbf{x} - (\theta_+ + \theta_-)/2)}{1 + \exp(\mathbf{w}^\top \mathbf{x} - (\theta_+ + \theta_-)/2)} & \text{if } \mathbf{w}^\top \mathbf{x} \geq \theta_+ \text{ or } \mathbf{w}^\top \mathbf{x} \leq \theta_- \\ p_b(\mathbf{x}) & \text{otherwise} \end{cases} \quad (17)$$

## 5.2 Model Training

With the loss function defined in (14), the hybrid model can be trained by solving the convex minimization problem (12) for which many efficient optimization techniques are available in literature including subgradient methods (Nemirovski et al., 2009; Duchi et al., 2011), accelerated gradient methods (Nesterov, 2013; Beck and Teboulle, 2009), primal-dual methods (Nemirovski, 2004; Chambolle and Pock, 2011) and many stochastic first-order methods based on randomly sampling over coordinates or data (Johnson and Zhang, 2013; Shalev-Shwartz and Zhang, 2014; Duchi et al., 2011; Zhang and Xiao, 2017). The choice of algorithms for (12) depends on various characteristics of the problem such as smoothness, strong convexity and data size.

Since numerical optimization is not the focus of this paper, we will simply utilize the accelerated proximal gradient method (APG) by Nesterov (2013) to solve (12) when  $\phi$  is smooth (e.g. the second and third cases in (15)). When  $\phi$  is non-smooth (e.g. the first cases in (15)), we apply the well-known smoothing technique (Nesterov, 2007) to create a smooth approximation for  $\phi$  and then solve the smooth problem using APG. To describe the smoothing technique, we denote by  $\phi^*(\alpha) : \mathbb{R} \rightarrow \mathbb{R} \cup \{+\infty\}$  the Fenchel conjugate of  $\phi$ , namely,  $\phi^*(\alpha) = \sup_z \{\alpha z - \phi(z)\}$ . It is well-known in convex analysis that, given our assumptions on  $\phi$ , we have  $\phi(z) = \sup_\alpha \{z\alpha - \phi^*(\alpha)\}$ . Let  $h(\alpha) : \mathbb{R} \rightarrow \mathbb{R}$  be a 1-strongly convex function with respect to the Euclidean norm. Given a smoothing parameter  $\mu \geq 0$ , the smooth approximation for  $\phi$  can be constructed as  $\phi_\mu := \sup_\alpha \{z\alpha - \phi^*(\alpha) - \mu h(\alpha)\}$ . According to Nesterov (2007), the function  $\phi_\mu$  is smooth and its gradient is  $\nabla \phi_\mu(z) = \arg \min_\alpha \{z\alpha - \phi^*(\alpha) - \mu h(\alpha)\}$  which has a closed-form for most commonly used  $\phi$ . According to Nesterov (2007), we have the relationship  $\phi_\mu(z) \leq \phi(z) \leq \phi_\mu(z) + \mu \max_{\alpha \in \text{dom}(\phi^*)} h(\alpha)$  where  $\text{dom}(\phi^*)$  is typically compact for non-smooth  $\phi$ .<sup>3</sup> These facts suggest that, to solve (12) with non-smooth  $\phi$ , we can approximate  $\phi$  by  $\phi_\mu$  with a small  $\mu$  and then apply the APG method (Nesterov, 2013) to the smooth problem

$$\min_{\mathbf{w}, \theta_+ \geq \theta_-} \mathcal{L}_\mu(\mathbf{w}, \theta_+, \theta_-; f_b) + \alpha_1 r(\mathbf{w}) + \alpha_2 (\theta_+ - \theta_-), \quad (18)$$

where

$$\mathcal{L}_\mu(\mathbf{w}, \theta_+, \theta_-; f_b) = \frac{1}{N} \sum_{n \in I_b^+} \phi_\mu(y^{(n)}(\mathbf{w}^\top \mathbf{x}^{(n)} - \theta_-)) + \frac{1}{N} \sum_{n \in I_b^-} \phi_\mu(y^{(n)}(\mathbf{w}^\top \mathbf{x}^{(n)} - \theta_+))$$

if the black-box returns binary predicted labels, and

$$\mathcal{L}_\mu(\mathbf{w}, \theta_+, \theta_-; f_b) = \frac{1}{N} \sum_{n=1}^N p_b^{(n)} \phi_\mu(y^{(n)}(\mathbf{w}^\top \mathbf{x}^{(n)} - \theta_-)) + \frac{1}{N} \sum_{n=1}^N (1 - p_b^{(n)}) \phi_\mu(y^{(n)}(\mathbf{w}^\top \mathbf{x}^{(n)} - \theta_+)) \quad (19)$$

if the black-box returns predicted probabilities.

The returned solution for (18) will also be a good solution to (12) when  $\mu$  is small. Note that when  $\mu = 0$ , we have  $\phi_\mu(z) = \phi(z)$  and (18) is reduced to (12). This smoothing APG method is presented as Algorithm 2. When  $\phi$  is smooth, one just needs to choose  $\mu = 0$ . For simplicity of notation, we have suppressed the dependence of  $\mathcal{L}_\mu(\mathbf{w}, \theta_+, \theta_-; f_b)$  on  $f_b$ .

The two subproblems in Step 4 and Step 5 in Algorithm 2 can be solved in closed forms for most commonly used  $r$  (e.g.,  $r(\mathbf{w}) = \|\mathbf{w}\|_1$ ). In Algorithm 2, users need to provide a step length parameter  $\eta$ . In order to guarantee the convergence in theory, this parameter can be chosen as  $\eta = \frac{1}{L_\mu}$ . Here,  $L_\mu$  is the Lipchitz continuity constant<sup>4</sup> of  $\nabla \mathcal{L}_\mu$  which is in the order of  $O(\frac{1}{\mu})$  if  $\phi$  is non-smooth and  $\mu > 0$  and is identical to the Lipchitz continuity constant  $\nabla \mathcal{L}$  if  $\phi$  is smooth and  $\mu = 0$ . When  $\eta$  is chosen to be this theoretical value, Algorithm 2 can ensure  $F(\mathbf{w}^{[t]}, \theta_+^{[t]}, \theta_-^{[t]}) - F^* \leq \epsilon$  in  $T = O(\frac{1}{\epsilon})$  iterations when  $\phi$  is non-smooth and in  $T = O(\frac{1}{\sqrt{\epsilon}})$  iterations when  $\phi$  is smooth. See Nesterov (2007, 2013) for the details of the convergence analysis. Although a theoretical value for  $\eta$  has a closed form,

3. In this paper, we use  $\text{dom}(f)$  to represent the domain of a function  $f$ .

4. The explicit expression of  $L_\mu$  can be found in Nesterov (2007).



it is generally too conservative and leads to slow convergence in practice. To address this issue, a line search scheme, for example the one proposed in Lin and Xiao (2015), can be applied to adaptively choose  $\eta$  during the iterations, which can improve the convergence speed in practice significantly while still preserving the theoretical convergence property. We implement APG with this adaptive line search scheme in our numerical experiments.

## 6. Experiments

We perform a detailed experimental evaluation of hybrid models on public data from domains where interpretability is critical, including healthcare, judiciaries, and business problems. On these datasets, the two forms of interpretable models will collaborate with popular black-box models including ensemble models and neural networks and solve problems on structured data and text data. We first discuss the evaluation metrics, then describe experimental setup and finally analyze the experimental results.

### 6.1 Evaluation Metrics

The goal for building a hybrid model is to provide interpretable and accurate predictions on as many data points as possible. Thus we propose two metrics below to evaluate the predictive performance, interpretability, and transparency.

**Model Complexity** The two types of models we choose (rules and linear models) have naturally interpretable forms. To provide a consistent evaluation of model complexity for hybrid rule sets and hybrid linear models, we measure the number of smallest **cognitive units** in a model. For rule sets, we count the total number of conditions (feature-value pairs) in association rules. Note that we choose the total number of conditions instead of the number of features because a feature could be used multiple times by different rules in a model. For linear models, we count the number of non-zero coefficients. Both of them represent the amount of units a user needs to comprehend in order to understand a model.

**Pareto Frontier** To evaluate the trade-off between transparency and predictive performance, such as accuracy, we use Pareto frontiers. A Pareto Frontier represents how much predictive performance a model needs to trade to obtain a certain level of transparency. None of the models on the same Pareto frontier dominates others, i.e., higher in both transparency and accuracy. Thus a Pareto frontier captures the natural trade-off obtained by hybrid models between transparency and accuracy. See Figure 3 for an illustration. The curves provide a practical and straightforward way for model selection. Instead of choosing between two discrete choices of either black-box models or interpretable models, now users can find a middle ground where partial transparency and good predictive performance can co-exist, providing one with a wider range of models. Users can choose the best operating point on this curve, based on one’s desired transparency as well as the tolerable loss in predictive performance.

### 6.2 Structured Datasets

We use four structured datasets that are publicly available. 1) *census* (Kohavi, 1996) ( $48,842 \times 14$ ) predicts if an individual’s annual income is greater than  $50K$  based on the individual’s demographic information, education background, and work information. 2) *ju-*

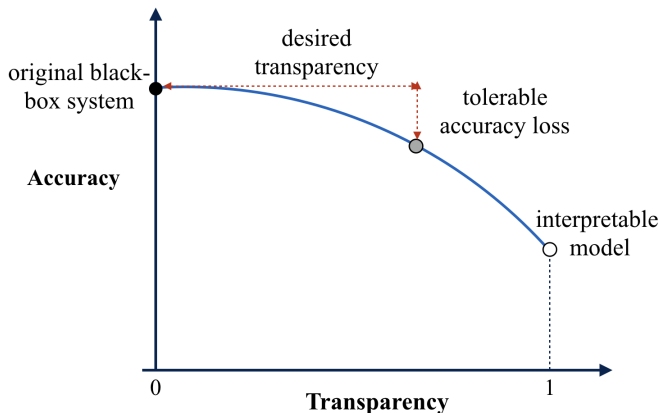


Figure 3: Pareto frontier of a hybrid model.

*venile* (Osofsky, 1995) (4,023 observations and 55 reduced features) studies the consequences of juvenile exposure to violence. The dataset was collected via a survey. 3) *coupon* (2017  $\times$  90 binarized features) predicts if a customer will accept a coupon recommended by an in-vehicle recommender system. 4) *stop-and-frisk* (80,755  $\times$  27) predicts if someone will be frisked by a police officer, using data from NYPD Stop, Question, and Frisk database from 2014 to 2016.

**Baselines** We process each dataset by binarizing all categorical features. For each structured dataset, we build three canonical black-box models as baselines and also later as input for hybrid models, Random Forests (Liaw et al., 2002), AdaBoost (Freund and Schapire, 1995) and extreme gradient boosting trees (XGBoost) (Chen and Guestrin, 2016). We partition each dataset into 80% training and 20% testing. We do 5-fold cross-validation for parameter tuning on the training set to choose the best parameters and then evaluate the best model on the test set. To provide a baseline for interpretability and accuracy comparison, we also build interpretable models. The methods we choose here include traditional decision trees CART (Breiman, 2017), C4.5 (Quinlan, 2014) and C5.0 (Pandya and Pandya, 2015) and more recent rule-based classifiers, Bayesian Rule Sets (BRS) (Wang et al., 2017) and Scalable Bayesian Rule List (Yang et al., 2017). BRS and SBRL are two recent representative methods which have proved to achieve simpler models with competitive predictive accuracy compared to the older rule-based classifiers. We also include Lasso (Tibshirani, 1996), generalized additive model (GAM) (Servén and Brummitt, 2018) and neural additive model (NAM) (Agarwal et al., 2020) to provide a baseline for hybrid linear models.

**Implementations** We use R packages (Hornik et al., 2007; Quinlan, 2004) to build interpretable methods except for BRS which has the code publicly available<sup>5</sup>. For C4.5 and C5.0, we tune the minimum number of samples at splits. For BRS, we set the maximum length of rules to 3 as recommended by the paper (Wang et al., 2017). For BRS, there are parameters  $\alpha_+$ ,  $\beta_+$ ,  $\alpha_-$ ,  $\beta_-$  that govern the likelihood of the data. We set  $\beta_+$ ,  $\beta_-$  to 1 and vary  $\alpha_+$ ,  $\alpha_-$  from  $\{100, 1000, 10000\}$ . For SBRL, there are hyperparameters  $\lambda$  for the

5. <https://github.com/wangtongada/BOA>

expected length of the rule list and  $\eta$  for the expected cardinality of the rules in the optimal rule list. We vary  $\lambda$  from  $\{5, 10, 15, 20\}$  and  $\eta$  from  $\{1, 2, 3, 4, 5\}$ . To preserve interpretability, all interpretable baselines have the number of rules under 30 (for decision trees we count the number of leaves). We used the python package for GAM (Servén and Brummitt, 2018) and the automatic parameter tuning via a grid search. We also used the public code <sup>6</sup> for NAM (Agarwal et al., 2020), where we tuned the number of layers in the neural networks from 1 to 3 and the number of basis functions from 5 to 40 with a step size of 5. Then we use python packages (Pedregosa et al., 2011) for the black-box models. For random forest, we tune the number of trees and the number of features for each tree. For Adaboost, we set the number of estimators to 800 and tune the maximum number of features and the maximum depth of trees. For XGBoost, we set the number of base learners to be 800 and set the learning rate to 0.01. Then we tune the maximum depth of trees, maximum features used by a tree and the minimum samples that need to exist in a leaf. For Lasso, we tuned the regularizing parameter  $\lambda$ . Models with the best cross-validated validation performance are selected.

The predictive performance of the black-box and interpretable models are reported in Table 1. On all four datasets, black-box models outperform interpretable models the exception of Lasso and GAM on juvenile dataset. But Lasso uses 303 non-zero coefficients and GAM needs 303 shape functions, making them hardly interpretable. To compare interpretability against the hybrid models, we report the number of conditions for rule-based models and the number of non-zero coefficients for Lasso in Figure 4.

Table 1: Accuracy of interpretable models on test sets

Datasets	<i>Black-box Models</i>			<i>Interpretable Models</i>							
	RF	Adaboost	XGBoost	CART	C4.5	C5.0	BRS	SBRL	Lasso	GAM	NAM
<b>Census</b>	.856	.856	.862	.843	.850	.846	.806	.820	.833	.849	.812
<b>Juvenile</b>	.904	.903	.904	.886	.884	.881	.881	.886	.909	.906	.895
<b>Coupon</b>	.755	.774	.703	.670	.678	.675	.675	.654	.665	.654	.661
<b>Stop&amp;Frisk</b>	.702	.695	.696	.678	.677	.679	.672	.677	.677	.686	.647

**Building Hybrid Models** Given the black-box models built above, we create hybrid rule sets and hybrid linear models. For hybrid rule sets, we set the maximum length of rules to be 4. We choose  $\alpha_1$  from  $[0.001, 0.06]$  and choose  $\alpha_2$  from  $[0.001, 0.5]$  to obtain a list of models. For hybrid linear models, we choose  $\alpha_1$  from  $\{0.001, 0.01, 0.1\}$  in order to select a small number of features. We choose  $\alpha_2$  from  $[-0.1, 0.1]$  to obtain a list of models. When building the hybrid linear model for all datasets, we choose  $\phi$  to be the logistic loss in (15), choose  $r(x) = \|x\|_1$ , and run Algorithm 2 with  $\mu = 0$  until either the number of iterations reaches 1,000 or the norm of the proximal gradient is no more than 0.0001, which is a standard way to terminate a gradient-based method. At termination, the changes of the objective value in Algorithm 2 is less than 0.001%, indicating the convergence of the algorithm.

We report the run time in Table 2. Hybrid linear models are much more Pareto since the search algorithm is based on gradient descent and computationally convenient. Hybrid

6. [https://github.com/nickfrosst/neural\\_additive\\_models](https://github.com/nickfrosst/neural_additive_models)

rule sets, on the other hand, involve evaluating intermediate solutions and proposing next states, and therefore they are computationally heavier.

Table 2: Means and standard deviations of training time for hybrid models

	Census	Juvenile	Coupon	Stop-and-Frisk
Hybrid Rule Set	10m 55s $\pm$ 45s	2m 10s $\pm$ 21s	3m 25s $\pm$ 6s	41m 50s $\pm$ 5s
Hybrid Linear Model	1m 5s $\pm$ 1s	26s $\pm$ 1s	52s $\pm$ 2s	2m 32s $\pm$ 3s

**Pareto Frontiers Analysis** We plot the Pareto frontiers for  $\langle f_{\text{rules}}, f_b \rangle$  and  $\langle f_{\text{linear}}, f_b \rangle$  in Figure 4. The Y-axis represents the accuracy evaluated on test sets. (See Figure 7 in the Appendix for Pareto frontiers evaluating AUC). For easier comparison, we set the range of the Y-axis to be the same for each dataset (row). The black-box baseline models are located at transparency equal to zero in each plot, represented by a black square. The hybrid models span the entire spectrum of transparency with accuracy slowly decreasing as transparency increases till reaching an interpretable model at transparency equal to one, represented by an empty circle. Besides, we show baseline interpretable models at the same row for each dataset. All of them are located at transparency equal to one.

The Pareto frontiers are concave with different decaying rates for the four datasets. For the census dataset, the curves remain almost flat as transparency is less than 50%. This suggests that transparency is gained at nearly no cost of predictive performance if the transparency is not too large. In addition, this **free** transparency is obtained using very few conditions for hybrid rule sets, only 4 or 5 conditions (1 or 2 rules in total), compared to the interpretable baselines using a lot more conditions but achieving lower accuracy overall. For the juvenile dataset, it is interesting to observe the accuracy even increases, although slightly, when 80% of the data are sent to an interpretable model, then drops rapidly if continuing sending data to interpretable models. This means the majority of the data can be easily explained by a couple of short rules or simple linear models, and the rest 20% need to be handled by a more advanced black-box model. If forcing everything to interpretable models, not only the performance will decrease, but the model complexity will also increase significantly. We notice that for this dataset, the highest accuracy is achieved by Lasso, 0.91. But Lasso uses 303 non-zero coefficients, turning it into a practically black-box model for having so many cognitive units. The results on census and juvenile datasets proved our hypothesis that even a black-box model is globally better, there might exist a subspace where a very simple model can do just as well or even slightly better, since simpler models are less likely to overfit. Finally, for the last two datasets, coupon recommendation and stop-and-frisk prediction, the Pareto frontier shows a clear trade-off between transparency and accuracy, with the accuracy steadily decreases as transparency increases.

We observe that hybrid linear models generally use more cognitive units (non-zero coefficients) in a model than hybrid rule sets (conditions). This is because linear models need more features to accurately define a decision boundary in the feature space. A slight change in the coefficients may affect the partitioning of the entire data. Therefore, regularizing the number of coefficients too much will hurt the performance. Thus the model has to keep many non-zero coefficients. On the other hand, hybrid rule sets use association rules, which are hyper-cubes in the data space and each are defined by a couple of features only. It

# HYBRID PREDICTIVE MODEL

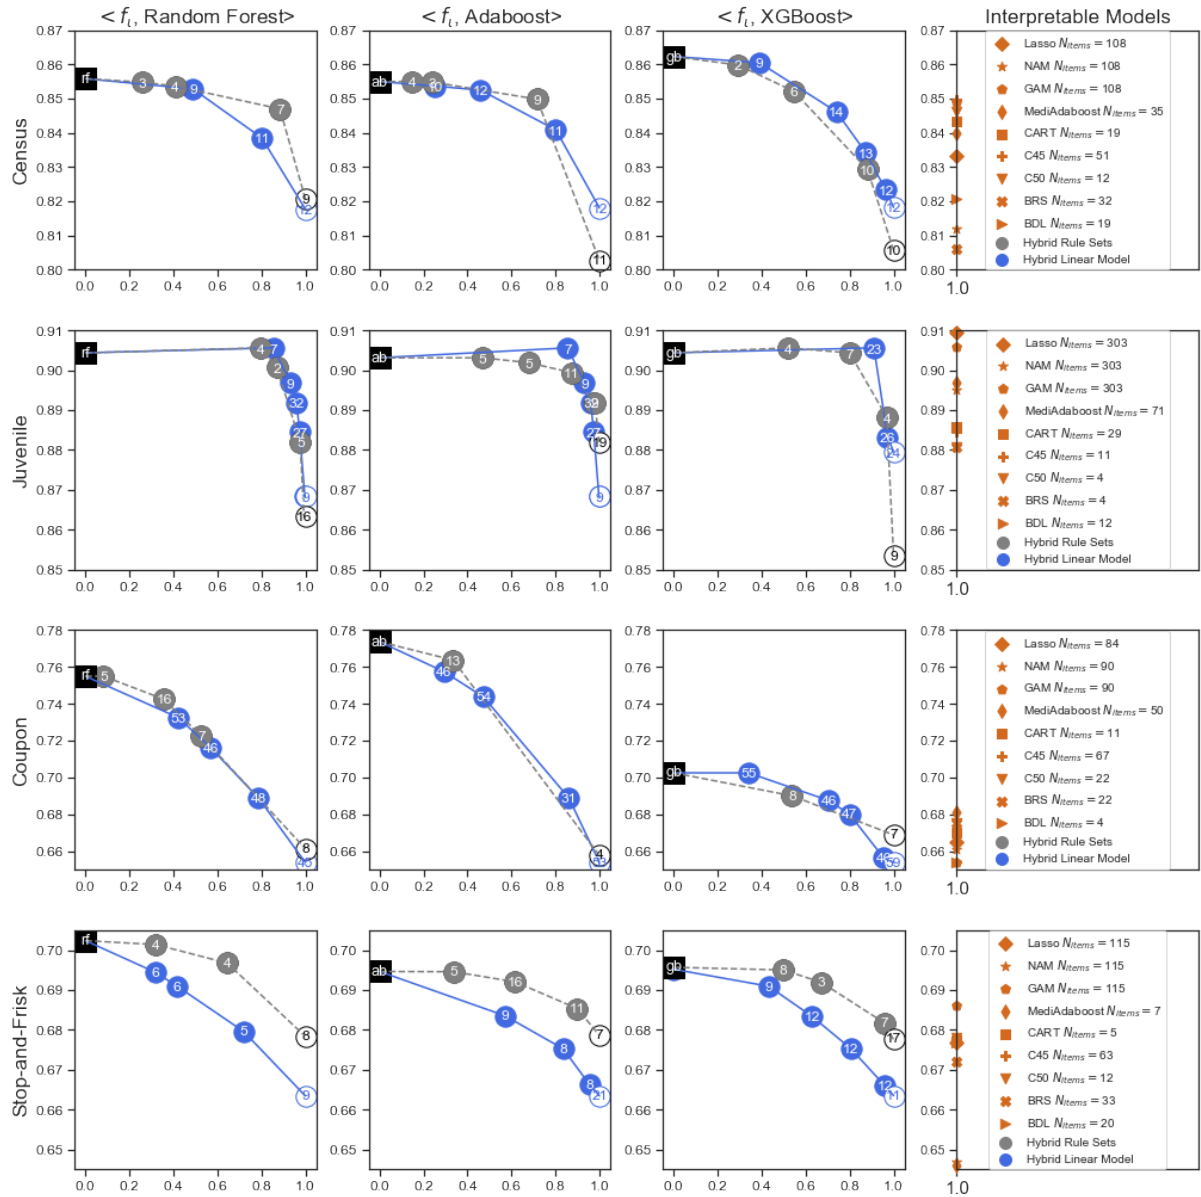


Figure 4: Pareto frontiers of hybrid models. The X-axis represents the transparency, and the Y-axis represents the accuracy on the test set. The value inside each circle is the number of items in a hybrid model. “rf”, “ab” and “gb” in the black squares represent Random Forest, Adaboost, and XGBoost.

has been shown in the literature that allowing rules to have length up to 4 can already accurately place it on correct subsets and achieve high predictive performance (Wang et al., 2017).

In addition to evaluating the binary outputs using accuracy, we also evaluate the probabilistic outputs using AUC as the metric. We repeat the same process described above to produce the AUC - transparency trade-off curves. The curves demonstrate similar pattern that the rules are able to carve out a significant amount of feature space without losing predictive performance on some data. Please see the plot in the appendix for more details.

We show an example of hybrid models below and another one in the appendix.

**Example: Hybrid Models for the Census Dataset** The census dataset predicts if an individual’s annual income is greater than \$50K, based on his/her demographic information (gender, age, marital status, race, etc), education background (degree and the number of years of education) and work information (occupation, hours for work per week, etc). We

Table 3:  $\langle f_{\text{rule}}, \text{Adaboost} \rangle$  for juvenile delinquency dataset. The transparency is 71.6% and the accuracy is 85.0%, evaluated on the test set.

	Rules	Model
<b>If</b>	capital gain = 0 <b>and</b> capital loss = 0 <b>and</b> number of years of education $\leq 13$ <b>and</b> degree $\neq$ Bachelor’s <b>and</b> occupation $\neq$ Tech-support <b>OR</b> age $\leq 41$ <b>and</b> capital gain = 0 <b>and</b> hours of work per week $< 40$ <b>and</b> marital status $\neq$ Married-civ-spouse $\rightarrow Y = -1$ (income $< \$50\text{K}$ )	} $\mathcal{R}_-$
<b>Else</b>	$Y = f_b(\mathbf{x})$	

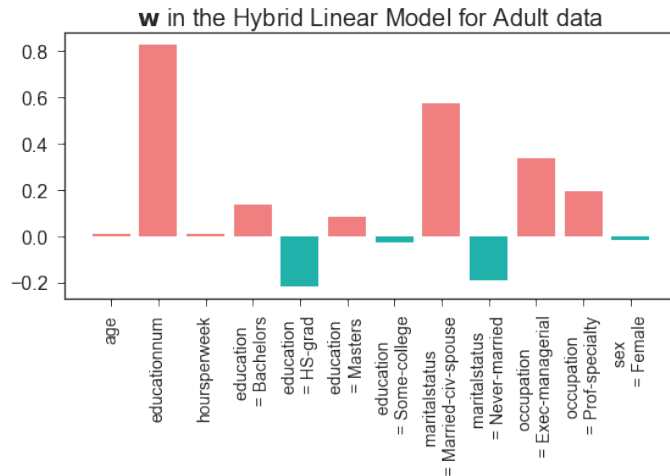


Figure 5: Coefficients in a  $\langle f_{\text{linear}}, \text{Adaboost} \rangle$  model for Census dataset. The model has transparency of 39.0% and accuracy of 86.1%.  $\theta_+ = 1.98, \theta_- = 0.70$ .

choose a hybrid rule set model that substitute Adaboost with reasonably large transparency and accuracy. We show the model in Table 3. The positive rule set is empty in this model. The negative rule set consists of two rules that capture negative instances. They cover 71.6% of the data and the overall accuracy is 85.0%.

We also show a hybrid linear model that also substitutes Adaboost model in Figure 5. The model identifies the most positive features include the number of years of education, if the marital status is married to a civilian spouse, and if the occupation is executive managerial or professional specialty; the most negative features are if education is high-school graduate and the marital status is never married.

### 6.3 Text Classification

We apply the hybrid models to two text mining tasks. One task is to predict movies’ financial success based on its professional critics’ reviews and the other is to classify whether a Yelp review is helpful. Since the obtained Pareto frontiers follow very similar trend as those for structured data and to prevent providing duplicate insights, we move the results, as well as detailed description of experimental setup, model training and parameter tuning to Appendix.

## 7. Cardiovascular Disease Prediction

In this section, we apply the proposed model to a real-world dataset to show that a hybrid model can be useful in practice. The context we choose is medical consultation /recommendation, where a medical staff provides medical advice and predictions based on a patient’s health conditions and medical history. Here we provide professional advice about whether a patient will develop cardiovascular disease. Let there be a pre-trained decision-maker  $f_b$ , and our goal is to build an interpretable model to assist  $f_b$  to provide understandable predictions. In this context,  $f_b$  can be a highly accurate black-box machine learning model such that an interpretable model we construct is for partial transparency while not losing too much accuracy. To extend the application, we can also assume the  $f_b$  represents a human decision-maker, which we can train  $f_b$  based on the human’s prior decisions. In such a case, an interpretable model  $f_l$  is developed to replace human labor for lower-cost, faster, while interpretable recommendations.

The dataset we use is a real-world patient dataset derived from an Atherosclerosis Risk in Communities Study (ARIC) study. These data are available upon request from BioLINCC.

**Data Description** The dataset consists of a total of 12,907 patients, for whom we have known cardiovascular disease(CVD) outcomes over ten years. Among the 12,907 patients, 8.5% have CVD. There are 110 defined features for each patient. The features include patients’ medical history such as whether having a probable myocardial infarction (MI), definite MI, suspect MI, definite fatal coronary heart disease (CHD), possible fatal CHD, stroke, etc, and their existing health condition such as blood pressure, heart rate, etc. In addition, there are features describing patients’ life habits such as consumption of vegetables, fruits, and protein, respectively, whether they smoke, how often they sport, etc. Finally, the features also include patients’ demographic information<sup>7</sup>.

**Experiment Design** This set of experiments is meant to more closely reflect a real-world scenario and, as such, is guided by a CVD specialist, either a human decision-maker or a black-box model. To simulate such a CVD specialist, we partition the data into 80% training

7. Please refer to <https://biolincc.nhlbi.nih.gov/studies/aric/> for the full list of features.

and 20% testing. We run three black-box models (random forest, Adaboost, and XGBoost) and choose the best performing model on the test set as the CVD specialist, which, in a real-world situation, can represent either a human specialist or a machine learning model. Since the data is highly unbalanced (8.5% positive), we choose AUC as the evaluation metric and use the probabilistic output from  $f_b$  to train hybrid models. Random forest achieved the highest AUC, 0.766, while Adaboost achieved an AUC of 0.763, and XGBoost obtained 0.753. Thus we choose Random forest to be  $f_b$ . In this experiment, we rule sets as the interpretable models. Since interpretability often needs to conform with domain-specific needs, here we assume patients or the human specialist set the upper bound of the number of items in a model to be 25. Thus the interpretable models and hybrid models we produce need to be smaller than that to be considered “interpretable”. The parameter setting and tuning is similar to that in Section 6.

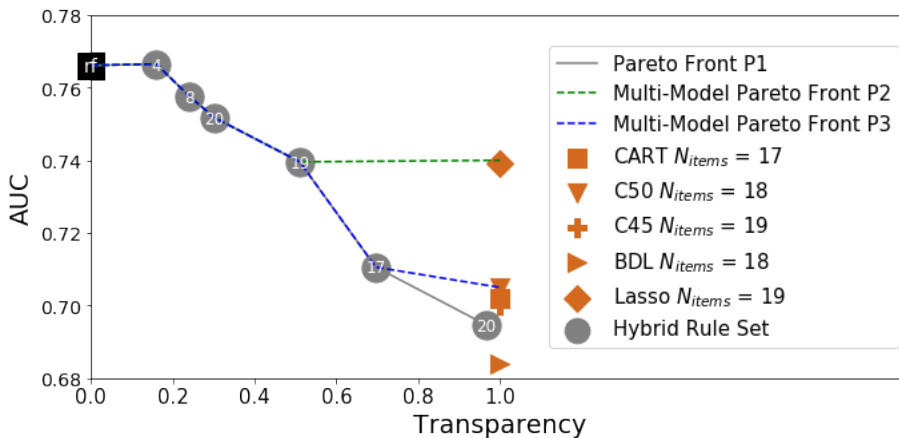


Figure 6: Pareto frontiers of hybrid models  $\langle f_L, \mathbf{rf} \rangle$ . The X-axis represents the transparency and the Y-axis represents the AUC on the test set. The value inside each circle is the number of conditions in a hybrid model.

**Results** We show in Figure 6 the Pareto fronts of the obtained hybrid rule sets. We show one of the models with 19 conditions in Table 4. This hybrid rule set consists of five rules, two of which capture positive cases and three capture negative cases. In cases where an instance satisfies multiple rules, the instance is assigned to the positive rules, following the “if” logic. Then, if an instance satisfies multiple rules of the same sign, it is always assigned the probability of the most confident rule, i.e., the rule with the largest support.

To compare with the hybrid models, we apply other completely transparent models discussed in Section 6. Note that Mediboost and BRS are excluded since they do not provide probabilistic outputs. We also exclude GAM as it produces shape functions for all 110 features, which violate our requirement of interpretability in this application setting. Since the data is highly unbalanced, to produce comparable results, we upsample the positive class 10 times to obtain a more balanced data and then train the baseline models.



Table 4:  $\langle f_{\text{rule}}, \text{XGboost} \rangle$  for predicting cardiovascular disease. The transparency is 51% and the accuracy is 74%, evaluated on the test set.

	Rules	Probability	Support
<b>If</b>	diabetes = 1 <i>and</i> VWF Value $\geq 98.0$ <i>and</i> high blood pressure med in past 2 weeks <i>and</i> peripheral artery disease =0, <b>then</b> <i>OR</i> no plaque or shadowing on either side of internal carotid <i>and</i> FVC (liters) $< 4.70$ <i>and</i> diastolic blood pressure $\geq 76$ <i>and</i> waist-to-hip ratio $\geq 0.86$ , <b>then</b>	$p_+ = 0.29$	3%
<b>Else if</b>	<i>OR</i> hypertension = 0 <i>and</i> Apolipoprotein AI (mg/dL) $\geq 124$ <i>and</i> LVHSCR01 $< 1016$ , <b>then</b> <i>OR</i> consumed dark or grain breads $\geq 1$ <i>and</i> never smoked cigarettes <i>and</i> systolic blood pressure $< 114.0$ mmHg1 <i>and</i> triglycerides in SI Units $< 1.38$ , <b>then</b> <i>OR</i> HDL cholesterol in mg/dL $\geq 45.26$ <i>and</i> number of cigarettes $< 5/\text{day}$ <i>and</i> hypertension = 0 <i>and</i> waist-to-hip ratio $< 0.86$	$p_+ = 0.13$ $p_+ = 0.03$ $p_+ = 0.02$	20% 21% 10.2%
<b>Else</b>	$Y = f_b(\mathbf{x})$	$p_+ = 0.01$	11% 49%

**Multi-model Pareto Front** If users only select from hybrid models, the AUC - transparency trade-off they get is characterized by Pareto Front  $P_1$  in Figure 6. However, in real-world applications, users often use many different methods to generate multiple models for model selection. We report the performance of interpretable baselines in Figure 6. To account for multi-model selection and different users’ preferences, we define *Multi-model Pareto Front* for practical use when users have different preferences for model interpretability. A multi-model Pareto front represents a set of models (except the black-box) produced by different methods that satisfy a user’s need for interpretability.

In this application, if a user considers a Lasso with 19 non-zero coefficients is a viable solution that meets his interpretability standard, then hybrid rule sets which only achieve less than 50% transparency but lower AUC are completely dominated by Lasso, those with 17 conditions and 20 conditions each, on the right bottom of Figure 6. This means, in a real-world application, hybrid rule sets with lower accuracy than the most accurate interpretable baseline, Lasso in this case, will never be considered since they are dominated by the interpretable models in both transparency and interpretability. Thus, the multi-model Pareto front for the users is  $P_2$ , represented by the green dashed line in Figure 6. On the other hand, some users consider multiplications and summations too complicated when using a linear model. They would prefer a list of rules that they can look through quickly without the need to do too much calculation in their heads. Then one would remove Lasso from consideration. The multi-model Pareto front for these users is, therefore, the Pareto front  $P_3$ , represented by the blue dashed line in Figure 6. Hybrid models provide more model choices for users such that they can weigh their considerations for transparency and predictive performance and choose a model that is most suitable for their applications.

## 8. Conclusion and Future Directions

We proposed a novel Hybrid Predictive Model that integrates an interpretable model with **any** black-box model for classification. The interpretable model substitutes the black-box

on a subset of data to gain transparency at an efficient cost of predictive performance. The model is trained to jointly optimize predictive performance, interpretability, and transparency, with carefully designed training algorithms to achieve the best balance among the three. Hybrid models provide more choices for users. Instead of choosing between black-box models (transparency equal to zero) and interpretable models (transparency equal to one), hybrid models span the entire spectrum of transparency, and users can choose the best operating point based on the desired transparency and tolerable loss in accuracy. We also propose multi-model Pareto fronts in the paper that allow users to choose from multiple methods, based on their own definition of interpretability and the need of predictive performance. Providing such flexibility for users is important as shown by a recent user study on users’ choices of models with different transparency and predictive performance (Pan et al., 2020).

Our model can be combined seamlessly with a black-box explainer, working as a pre-step for a black-box explainer: we first construct an interpretable model to capture and explain a sub-region with 100% fidelity and leave the rest of the data to a black-box model to predict and an explainer to provide post hoc explanations.

### Future Directions

For future research, there are several directions worth pursuing. First, we can explore using other forms of interpretable models, such as case-based models: if an instance is close enough to a prototype, the prediction is determined by the prototype; otherwise, it is sent to a black-box model. In addition, other regularization terms of interpretable models can be included in the objective function, such as using fewer features or adding a fairness component that penalizes disparity in different subgroups of data.

In this paper, the black-box is pre-trained, and only the outputs are provided during training. It will be interesting to co-train a black-box and an interpretable model. We hypothesize that the co-training is likely to achieve a better balance among accuracy, interpretability, and transparency, but the computation complexity will increase significantly. It remains a question whether the benefits of co-training are worth the effort.

Finally, the current models only work on structured data and text data. Association rules and linear models are not designed for raw image processing or other problems where the features themselves are not semantically interpretable. The notions of interpretability used in image classification tend to be completely different from those for structured data where each feature is separately meaningful. There are some recent attempts on building interpretable models for image classification (Chen et al., 2018; Zhang et al., 2018) and it will be interesting to combine these models with black-box neural nets such as CNN or ResNet for image classification, following the proposed hybrid framework.

## Appendix A: Proofs for Theorems

**Proof** (of Theorem 5) To prove that an optimal model does not contain any rules with support smaller than a threshold, we prove by contradiction that if there's a model that contains such a rule  $z$ , removing it will always decrease the objective value, thus violating the optimality assumption.

We start with positive rules. For a rule  $r \in \mathcal{R}_+^*$ , we define

$$\mathcal{R}_{\setminus r}^* = \{z \in \mathcal{R}_+^*, z \neq r\}. \quad (20)$$

We want to find conditions on the support such that the following inequality always holds.

$$\Lambda(\mathcal{R}_{\setminus r}^*) \leq \Lambda(\mathcal{R}^*), \quad (21)$$

where

$$\Lambda(\mathcal{R}^*) = \ell(\mathcal{R}^*) + \alpha_1 \cdot \Omega(\mathcal{R}^*) - \alpha_2 \cdot \frac{\text{support}(\mathcal{R}^*)}{N} \quad (22)$$

$$\Lambda(\mathcal{R}_{\setminus r}^*) \leq \ell(\mathcal{R}_{\setminus r}^*) + \alpha_1 \cdot (\Omega(\mathcal{R}^*) - 1) - \alpha_2 \cdot \frac{\text{support}(\mathcal{R}_{\setminus r}^*)}{N}. \quad (23)$$

Since we want inequality (21) to hold for any  $r \in \mathcal{R}^*$ , we upper bound  $\Lambda(\mathcal{R}_{\setminus r}^*)$  by upper bounding  $\ell(\mathcal{R}_{\setminus r}^*)$  and  $\text{support}(\mathcal{R}_{\setminus r}^*)$ .

$$\ell(\mathcal{R}_{\setminus r}^*) \leq \ell(\mathcal{R}^*) + \frac{\text{support}(\mathcal{R}_{\setminus r}^*)}{N}, \quad (24)$$

with the minimum achieved when instances originally covered by  $r$  are all incorrectly classified after removing  $r$ .

$$\text{support}(\mathcal{R}_{\setminus r}^*) \leq \text{support}(\mathcal{R}^*) \quad (25)$$

with the minimum achieved when all instances originally covered by  $r$  are now covered by  $\mathcal{R}_-^*$ , therefore does not change the coverage of  $\mathcal{R}^*$  overall.

Plugging formula (24) and (25) into equation (23) and combine it with (22) and (21) yields

$$\text{support}(r) \leq N\alpha_1. \quad (26)$$

Thus, if  $\text{support}(r) \leq N\alpha_1$ , removing it from a  $\mathcal{R}^*$  will produce a better model. Therefore, such rules do not exist in an optimal model  $\mathcal{R}^*$ .

Then we follow the similar steps to prove for negative rules. We define  $\mathcal{R}_{\setminus r}^*$  as a set of rules where  $r$  is removed from  $\mathcal{R}_-^*$ . The proofs here use the same steps from inequality (20) to (24). The effective coverage, however, equals to  $\text{support}(\mathcal{R}^*) - \text{support}(\mathcal{R}_{\setminus r}^*)$ . Thus

$$\text{support}(r) \leq \frac{N\alpha_1}{1 - \alpha_2}. \quad (27)$$

To summarize,  $\mathcal{R}_+^*$  does not contain any rules with  $\text{support}(r) \leq N\alpha_1$  and  $\mathcal{R}_-^*$  does not contain any rules with  $\text{support}(r) \leq \frac{N\alpha_1}{1 - \alpha_2}$ . ■

**Proof** (of Theorem 6) We choose the optimal model found till time  $t$  to be the benchmark to compare with  $\mathcal{R}^*$ . Since  $\mathcal{R}^* \in \min \Lambda(\mathcal{R})$ ,

$$\lambda^{[t]} \geq \Lambda(\mathcal{R}^*), \quad (28)$$

Now we lower bound  $\Lambda(\mathcal{R}^*)$ , following equation (21)

$$\Lambda(\mathcal{R}^*) \geq 0 + \alpha_1 \Omega(\mathcal{R}^*) - \alpha_2 \quad (29)$$

Combining inequality (28) and (29) yields

$$\Omega(\mathcal{R}^*) \leq \frac{\lambda^{[t]} + \alpha_2}{\alpha_1}. \quad (30)$$

■

**Proof** (of Theorem 7) We again compare  $\mathcal{R}^*$  with the best model we found till time  $t$  and

$$\lambda^{[t]} \geq \Lambda(\mathcal{R}^*), \quad (31)$$

Then we lower bound  $\Lambda(\mathcal{R}^*)$ ,

$$\Lambda(\mathcal{R}^*) \geq 0 + \alpha_1 \Omega(\mathcal{R}^*) - \alpha_2 \frac{\text{support}(\mathcal{R}^*)}{N} \quad (32)$$

If  $\mathcal{R}^* \neq \emptyset$ , then  $\Omega(\mathcal{R}^*) \geq 1$ , then

$$\lambda^{[t]} \geq \alpha_1 - \alpha_2 \frac{\text{support}(\mathcal{R}^*)}{N}, \quad (33)$$

Thus

$$\text{support}(\mathcal{R}^*) \geq \frac{N(\alpha_1 - \lambda^{[t]})}{\alpha_2} \quad (34)$$

■

## Appendix B: Supplementary Experimental Results

**Part I: AUC Evaluation of Hybrid Models on Structured datasets** In order to evaluate the AUC obtained by the hybrid models, we first generate the predicted probability of being positive for each testing instance. We used the formulas (9) and (10) as the probabilities by the hybrid rule set method and used the formula (17) as the probabilities by the hybrid linear method. Then we sort all testing instances according to their probabilities of being positive and calculated AUCs. The result based on AUC can be found in Figure 7.

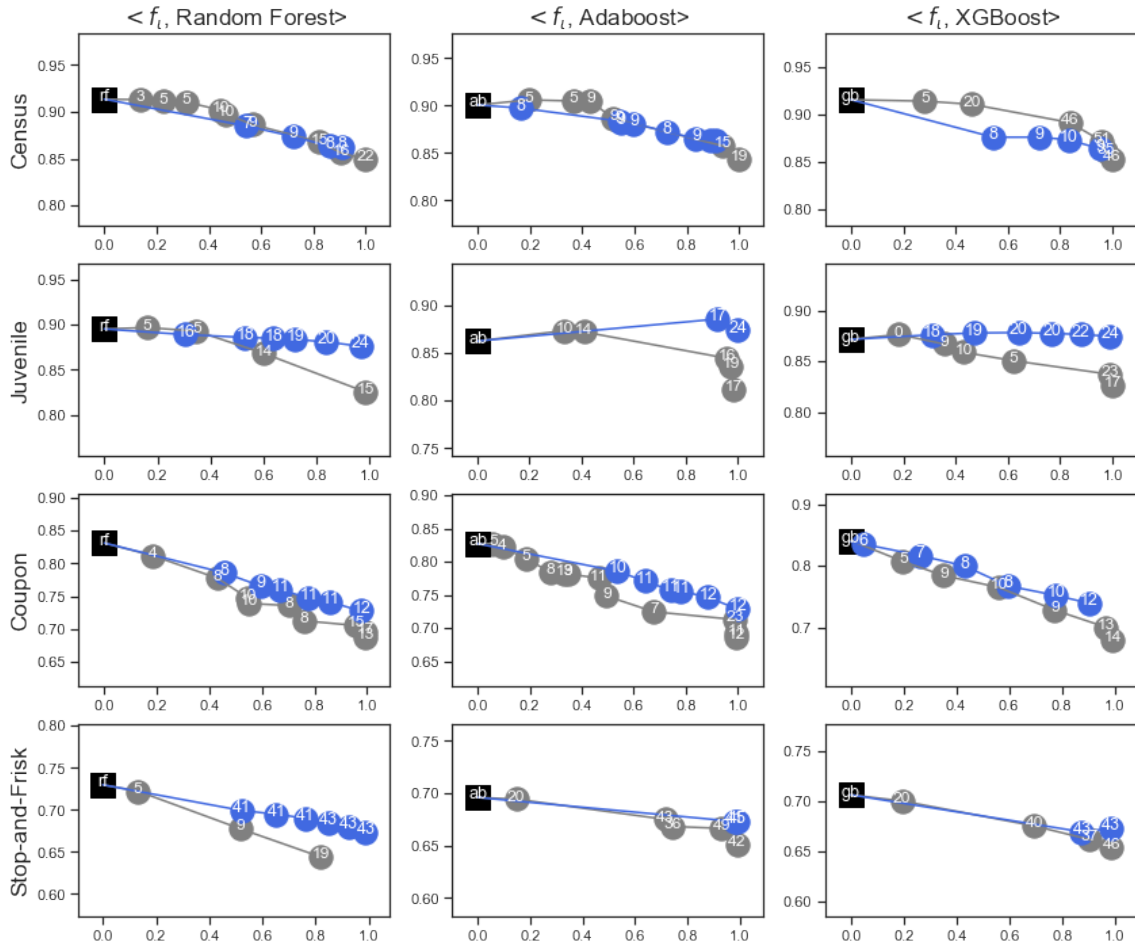


Figure 7: Pareto fronts of hybrid models. The X-axis represents the transparency and the Y-axis represents the AUC on the test set. The value inside each circle is the number of items in a hybrid model. “rf”, “ab” and “gb” in the black squares represent Random Forest, Adaboost, and XGBoost, respectively.

**Part II: Examples of Hybrid Models for the Juvenile Delinquency Dataset** We show hybrid models built from the juvenile delinquency data. The dataset was collected from a survey where children were asked questions regarding their exposure to violence from family members, friends, or other sources in their community and use that to predict if s/he will commit delinquency. Therefore, each feature is a question from the survey and the corresponding values are respondents’ answer to the questions. For example, question 18 asks “has anyone (including family members or friends) ever attacked you with a gun, knife or some other weapon?” and the answer is “Yes”, “No”, “Refused to answer”, or “Not sure”.

We first show rule sets substitutes in an XGBoost model in Table 5. The positive rule set contains one rule with three conditions, and the negative rule set contains one rule with four conditions. They collectively cover and explain 80% of the data without losing any predictive performance from XGBoost alone.

Table 5:  $\langle f_{\text{rule}}, \text{XGboost} \rangle$  for juvenile delinquency dataset. The transparency is 80.0% and the accuracy is 90.6%, evaluated on the test set.

	Rules	Model	$p^+$
<b>If</b>	has anyone (including family members or friends) ever attacked you with a gun, knife or some other weapon? (Q18A) = “Yes” <b>and</b> (Q18C) has anyone (including family members or friends) ever threatened you with a gun or knife, but didn’t actually shoot or cut you = “Yes” <b>and</b> has your friends ever purposely damaged or destroyed property that did not belong to them ? (Q48AA)= “Yes” → $Y = 1$ (will commit delinquency)	} $\mathcal{R}_+$	
<b>Else if</b>	has anyone (including family members or friends) ever physically attacked you without a weapon, but you thought you were trying to kill or seriously injure you? (Q18B) $\neq$ “Yes” <b>and</b> has your friends ever sold hard drugs such as heroin, cocaine and LSD? (Q48AG) $\neq$ “Yes” <b>and</b> has your friends ever stolen something worth more than \$50? (Q48AH) $\neq$ “Yes” <b>and</b> has your friends ever purposely damaged or destroyed property that did not belong to them (Q48BA) $\neq$ “All of them” → $Y = -1$ (will not commit delinquency)		
<b>Else</b>	$Y = f_b(\mathbf{x})$	$f_b$	

For comparison, we show in Figure 8 the hybrid linear model that substitutes Adaboost, with the same accuracy as the hybrid rule set but larger transparency, 85.5%.

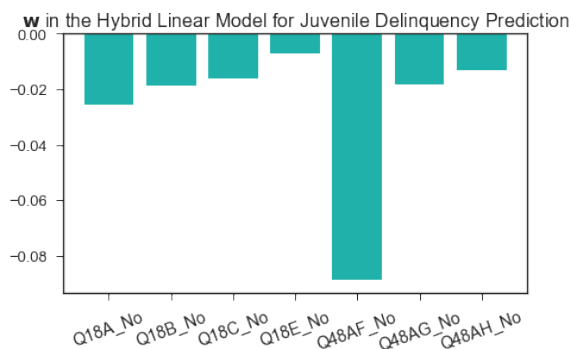


Figure 8: Coefficients in a  $\langle f_{\text{linear}}, \text{Adaboost} \rangle$  model. The model has transparency of 85.5% and accuracy of 90.6%.  $\theta_+ = 0.81, \theta_- = -0.14$ . See Appendix B for the questions in the figure.

The two models are very consistent in choosing the features. Both models include question 18A, 18B, 18C, 48AG and 48AH and there is slight difference in using a couple of other questions. The questions are listed below.

- Q18A: has anyone (including family members or friends) ever attacked you with a gun, knife or some other weapon, regardless of when it happened or whether you ever reported to police
- Q18B: has anyone (including family members or friends) ever physically attacked you without a weapon, but you thought they were trying to kill or seriously injure you?
- Q18C: has anyone (including family members or friends) ever threatened you with a gun or knife, but didn't actually shoot or cut you?
- Q18E: has anyone (including family members or friends) ever beat you up with their fists so hard that you were hurt pretty bad?
- Q48AF: has your friends ever broken into a vehicle or building to steal something?
- Q48AG: has your friends sold hard drugs such as heroin, cocaine, and LSD?
- Q48AH: has your friends stolen something worth more than \$50 ?

**Part III: Text Classification** We apply the hybrid models to two text mining tasks. For each dataset, we convert text data into bag-of-words, and each feature is a word.

**Movie Success Prediction** The goal of this task is to predict the financial success of a movie based on its professional critics' reviews. We frame this problem into a binary classification task. In particular, we define as "Success" a movie that produced a gross profit of over 30% above its production budget and define as "Failure" otherwise. The dataset is publicly available (Joshi et al., 2010)<sup>8</sup> and contains reviews on 1,718 movies. We have 616 movies (230 successes and 386 failures) remained after eliminating the movies that did not have the required financial information (production budget and gross profit). We randomly partition the data into training (492 movies), and testing (124 movies) sets. After merging all reviews on the same movies into one document, we then create the document-term matrix (DTM) with the frequencies of terms after removing punctuation, numbers, and stopwords and applying stemming. In the DTM, we only use the terms with frequency at least 100 and appearing in at least five documents. The obtained DTM contains 2,069 columns (terms). We also use two black-box models for this task. The first model is the XGBoost model using the same hyper-parameters as in the previous task. The second model is the same neural network as in the previous task except that there are 64 neurons in the first hidden layers. The hyperparameters of the black-box models are selected from a finite grid using hold-out validation.

**Review Usefulness Classification Experimental Setup and Details** In this task, the goal is to predict whether a consumer review in Yelp is useful. The dataset we use is a subset of the publically available Yelp Challenge data<sup>9</sup> in 2017. The subset we select contains all the reviews in the original dataset that are about the businesses in the category of "Doctors". On Yelp, consumers can vote "useful" for a review (but cannot vote "not useful"). We define a useful review as one that had received at least one useful vote before

---

8. [www.cs.cmu.edu/~ark/movie\\$-data/](http://www.cs.cmu.edu/~ark/movie$-data/)

9. <https://www.yelp.com/dataset/challenge>

it was collected into this dataset. With this definition, we obtain 25,901 reviews where 12,970 are not useful, and 12,931 are useful. We randomly partition the data into training (20,721 reviews), and testing (5,180 reviews) sets. We then create the document-term matrix (DTM) in terms of TF-IDF after removing punctuation, numbers, and stopwords and applying stemming. In the DTM, we only use the terms with frequency at least 50. The obtained DTM contains 2,315 columns (terms). We used two different black-box models  $f_b$  for this task. The first model is the extreme gradient boosting trees (XGBoost) (Chen and Guestrin, 2016) with the maximum depth of a tree being 5. The second model is a three-layer fully connected neural network, where the activation function is the sigmoid function, and there are 128 and 16 neurons in the first and second hidden layers. The hyperparameters of the black-box models are selected from a finite grid using hold-out validation.

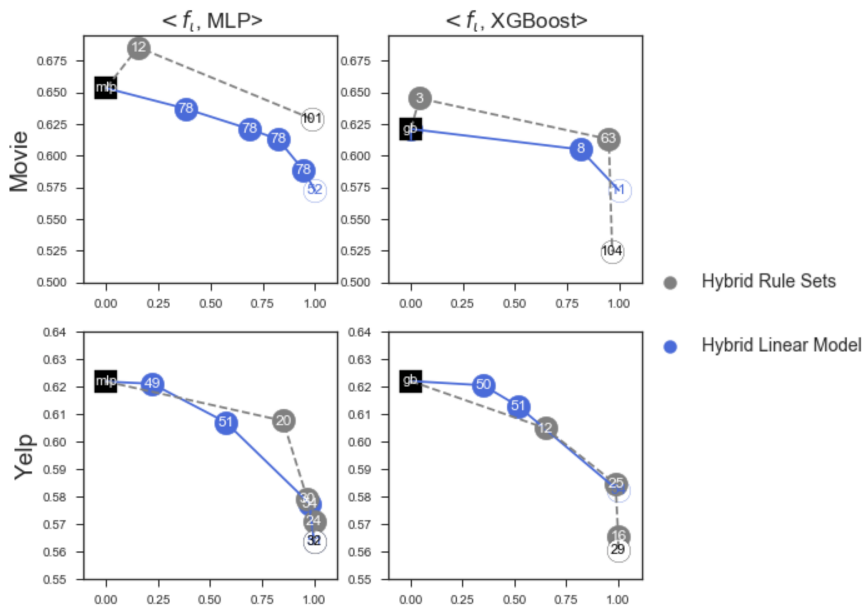


Figure 9: Pareto frontiers of hybrid models. The X-axis represents the transparency and the Y-axis represents the accuracy on the test set. The value inside each circle is the number of items in a hybrid model.

We train the two types of hybrid models for the tasks above and tune parameters in both methods to obtain a list of models. When training hybrid rule sets, we binarize the data and consider both the existence and the non-existence of words (considering the negation of features), thus doubling the features. When training hybrid linear models, we continue to choose  $\phi$  to be the hinge loss and  $r(x) = \|x\|_1$  and choose the parameters in Algorithm 2 as in Section 6.2. We vary  $\alpha_2$  from 0 to 1 to show the trade-off between transparency and accuracy. For each value of  $\alpha_2$ , we search  $\alpha_1$  over a grid using cross-validation to research a good balance between the classification accuracy on the validation set and the number of non-zero coefficients in  $\mathbf{w}$ . In particular, we choose  $\alpha_1$  such that the number of non-zero coefficients is no more than 80 in order to obtain an interpretable model. The Pareto frontiers



evaluated on the test set are shown in Figure 9. For the Yelp review dataset, the curves for the two models are very similar. They decrease slowly when transparency is small and then more quickly when transparency is larger. For the movie review dataset, it is interesting to observe that the hybrid rule sets models are able to outperform the black-box by a little when transparency is small. This is because rules are less likely to overfit especially the black-box model uses thousands of features when the rules only uses a very small subset.

We examine hybrid models for the yelp review data set. We choose models that collaborate with XGBoost and with reasonably large transparency. Table 6 shows the hybrid rule set model. The positive rule set contains two rules and the negative rule set contains one rule. The rules describe which words need to be included or cannot be included in order to be classified as positive or negative. Figure 10 shows the words selected by the linear model and their corresponding coefficients. Note that the coefficients are all positive in this model, indicating each word listed will contributing to predicting positive (a useful review). However, this does not mean that the linear model only captures the positive class. The threshold for predicting negative is 0.16, which means if a review does not contain any of the words in Figure 10, or only a few words with low weight, then  $\mathbf{w}^\top \mathbf{x}$  will marginally above 0, but less than 0.16 and the instance is predicted as negative. Only when the scores are large enough, for example, containing many keywords from the list, the model predicts positive. If  $\mathbf{w}^\top \mathbf{x}$  is between 0.16 and 0.82, the XGBoost model is activated to make a prediction.

Table 6:  $\langle f_{\text{rule}}, \text{XGboost} \rangle$  for juvenile delinquency dataset. The transparency is 64.9% and the accuracy is 60.5%, evaluated on the test set.

	Rules	Model
<b>If</b>	the review contains “wouldnt” <b>and</b> “ever” <b>and</b> doesn’t contain “left” <b>OR</b> the review contains “age” <b>and</b> doesn’t contain “worth” <b>OR</b> the review contains “let” <b>and</b> doesn’t contain “occas” or “lab” → $Y = 1$ (review is useful)	} $\mathcal{R}_+$
<b>Else if</b>	the review doesn’t contains “seem” or “night” or “cant” or “saw” → $Y = -1$ (review is unuseful)	
<b>Else</b>	$Y = f_b(\mathbf{x})$	$f_b$




Figure 10: Coefficients in a  $\langle f_{\text{linear}}, \text{XGBoost} \rangle$ . The accuracy is 61.3%, and the transparency is 51.7%.  $\theta_+ = 0.82, \theta_- = 0.16$ .

## References

- Philip Adler, Casey Falk, Sorelle A Friedler, Gabriel Rybeck, Carlos Scheidegger, Brandon Smith, and Suresh Venkatasubramanian. Auditing black-box models for indirect influence. In *ICDM*, pages 1–10. IEEE, 2016.
- Rishabh Agarwal, Nicholas Frosst, Xuezhou Zhang, Rich Caruana, and Geoffrey E Hinton. Neural additive models: Interpretable machine learning with neural nets. *arXiv preprint arXiv:2004.13912*, 2020.
- Ulrich Aivodji, Hiromi Arai, Olivier Fortineau, Sébastien Gambs, Satoshi Hara, and Alain Tapp. Fairwashing: the risk of rationalization. *arXiv preprint arXiv:1901.09749*, 2019.
- Elaine Angelino, Nicholas Larus-Stone, Daniel Alabi, Margo Seltzer, and Cynthia Rudin. Learning certifiably optimal rule lists. In *SIGKDD*, pages 35–44. ACM, 2017.
- Amir Beck and Marc Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM journal on imaging sciences*, 2(1):183–202, 2009.
- Hamid Berenji, Yan Wang, David Vengerov, Reza Langari, and Mo Jamshidi. Using gated experts in fault diagnosis and prognosis. In *2004 IEEE International Conference on Fuzzy Systems (IEEE Cat. No. 04CH37542)*, volume 1, pages 463–467. IEEE, 2004.
- Leo Breiman. *Classification and regression trees*. Routledge, 2017.
- Antonin Chambolle and Thomas Pock. A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of mathematical imaging and vision*, 40(1):120–145, 2011.
- Chaofan Chen, Oscar Li, Alina Barnett, Jonathan Su, and Cynthia Rudin. This looks like that: deep learning for interpretable image recognition. *arXiv preprint arXiv:1806.10574*, 2018.
- Tianqi Chen and Carlos Guestrin. Xgboost: A scalable tree boosting system. In *SIGKDD*, pages 785–794. ACM, 2016.
- Finale Doshi-Velez and Been Kim. A roadmap for a rigorous science of interpretability. *arXiv preprint arXiv:1702.08608*, 2017.
- John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12(Jul): 2121–2159, 2011.
- Yoav Freund and Robert E Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. In *European conference on computational learning theory*, pages 23–37. Springer, 1995.
- Nicholas Frosst and Geoffrey Hinton. Distilling a neural network into a soft decision tree. *arXiv preprint arXiv:1711.09784*, 2017.

- A d’Avila Garcez, Tarek R Besold, Luc De Raedt, Peter Földiak, Pascal Hitzler, Thomas Icard, Kai-Uwe Kühnberger, Luis C Lamb, Risto Miikkulainen, and Daniel L Silver. Neural-symbolic learning and reasoning: contributions and challenges. In *Proceedings of the AAAI Spring Symposium on Knowledge Representation and Reasoning: Integrating Symbolic and Neural Approaches, Stanford*, 2015.
- Artur S d’Avila Garcez, Krysia B Broda, and Dov M Gabbay. *Neural-symbolic learning systems: foundations and applications*. Springer Science & Business Media, 2012.
- Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014.
- K Hornik, A Zeileis, T Hothorn, and C Buchta. Rweka: an r interface to weka. *R package version*, pages 03–4, 2007.
- Zhiting Hu, Xuezhe Ma, Zhengzhong Liu, Eduard Hovy, and Eric Xing. Harnessing deep neural networks with logic rules. *arXiv preprint arXiv:1603.06318*, 2016.
- Zhongsheng Hua and Bin Zhang. A hybrid support vector machines and logistic regression approach for forecasting intermittent demand of spare parts. *Applied Mathematics and Computation*, 181(2):1035–1048, 2006.
- Rie Johnson and Tong Zhang. Accelerating stochastic gradient descent using predictive variance reduction. In *Advances in neural information processing systems*, pages 315–323, 2013.
- Mahesh Joshi, Dipanjan Das, Kevin Gimpel, and Noah A Smith. Movie reviews and revenues: An experiment in text regression. In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 293–296. Association for Computational Linguistics, 2010.
- Scott Kirkpatrick, C Daniel Gelatt, and Mario P Vecchi. Optimization by simulated annealing. *science*, 220(4598):671–680, 1983.
- Hian Chye Koh, Wei Chin Tan, and Chwee Peng Goh. A two-step method to construct credit scoring models with data mining techniques. *International Journal of Business and Information*, 1(1), 2015.
- Ron Kohavi. Scaling up the accuracy of naive-bayes classifiers: A decision-tree hybrid. In *KDD*, volume 96, pages 202–207, 1996.
- Himabindu Lakkaraju and Osbert Bastani. ” how do i fool you?” manipulating user trust via misleading black box explanations. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, pages 79–85, 2020.
- Himabindu Lakkaraju, Stephen H Bach, and Jure Leskovec. Interpretable decision sets: A joint framework for description and prediction. In *SIGKDD*. ACM, 2016.
- Himabindu Lakkaraju, Ece Kamar, Rich Caruana, and Jure Leskovec. Interpretable & explorable approximations of black box models. *arXiv preprint arXiv:1707.01154*, 2017.

- Himabindu Lakkaraju, Ece Kamar, Rich Caruana, and Jure Leskovec. Faithful and customizable explanations of black box models. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, pages 131–138, 2019.
- Andy Liaw, Matthew Wiener, et al. Classification and regression by randomforest. *R news*, 2(3):18–22, 2002.
- Qihang Lin and Lin Xiao. An adaptive accelerated proximal gradient method and its homotopy continuation for sparse optimization. *Computational Optimization and Applications*, 60(3):633–674, 2015.
- M Lissack. Dealing with ambiguity—the ‘black box’ as a design choice. *SheJi (forthcoming)*, 2016.
- Yin Lou, Rich Caruana, and Johannes Gehrke. Intelligible models for classification and regression. In *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 150–158, 2012.
- Tyler McCormick, Cynthia Rudin, and David Madigan. A hierarchical model for association rule mining of sequential events: An approach to automated medical symptom prediction. 2011.
- Kenneth McGarry, Stefan Wermter, and John MacIntyre. Hybrid neural systems: from simple coupling to fully integrated neural networks. *Neural Computing Surveys*, 2(1):62–93, 1999.
- Arkadi Nemirovski. Prox-method with rate of convergence  $o(1/t)$  for variational inequalities with lipschitz continuous monotone operators and smooth convex-concave saddle point problems. *SIAM Journal on Optimization*, 15(1):229–251, 2004.
- Arkadi Nemirovski, Anatoli Juditsky, Guanghui Lan, and Alexander Shapiro. Robust stochastic approximation approach to stochastic programming. *SIAM Journal on optimization*, 19(4):1574–1609, 2009.
- Yurii Nesterov. Smoothing technique and its applications in semidefinite optimization. *Mathematical Programming*, 110(2):245–259, 2007.
- Yurii Nesterov. *Introductory lectures on convex optimization: A basic course*, volume 87. Springer Science & Business Media, 2013.
- Xia Ning and George Karypis. Slim: Sparse linear methods for top-n recommender systems. In *2011 IEEE 11th International Conference on Data Mining*, pages 497–506. IEEE, 2011.
- Joy D Osofsky. The effect of exposure to violence on young children. *American Psychologist*, 50(9):782, 1995.
- Danqing Pan, Tong Wang, and Satoshi Hara. Interpretable companions for black-box models. *arXiv preprint arXiv:2002.03494*, 2020.

- Rutvija Pandya and Jayati Pandya. C5. 0 algorithm to improved decision tree with feature selection and reduced error pruning. *International Journal of Computer Applications*, 117(16):18–21, 2015.
- Parliament and Council of the European Union. General data protection regulation, 2016.
- Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. Scikit-learn: Machine learning in python. *Journal of machine learning research*, 12(Oct):2825–2830, 2011.
- Mark Phillips. International data-sharing norms: from the oecd to the general data protection regulation (gdpr). *Human genetics*, pages 1–8, 2018.
- J Ross Quinlan. Data mining tools see5 and c5. 0. <http://www.rulequest.com/see5-info.html>, 2004.
- J Ross Quinlan. *C4. 5: programs for machine learning*. Elsevier, 2014.
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. Why should i trust you?: Explaining the predictions of any classifier. In *SIGKDD*, pages 1135–1144. ACM, 2016.
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. Anchors: High-precision model-agnostic explanations. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- Michael M Richter and Rosina O Weber. *Case-based reasoning*. Springer, 2016.
- Peter R Rijnbeek and Jan A Kors. Finding a short and accurate decision rule in disjunctive normal form by exhaustive search. *Machine learning*, 80(1):33–62, 2010.
- Andrew Slavin Ross, Michael C Hughes, and Finale Doshi-Velez. Right for the right reasons: Training differentiable models by constraining their explanations. *arXiv preprint arXiv:1703.03717*, 2017.
- Cynthia Rudin. Please stop explaining black box models for high stakes decisions. *arXiv preprint arXiv:1811.10154*, 2018.
- D. Servén and C. Brummitt. pygam: Generalized additive models in python. *Zenodo*, 2018.
- Shai Shalev-Shwartz and Tong Zhang. Accelerated proximal stochastic dual coordinate ascent for regularized loss minimization. In *International Conference on Machine Learning*, pages 64–72, 2014.
- Chung-Kwan Shin, Ui Tak Yun, Huy Kang Kim, and Sang Chan Park. A hybrid approach of neural network and memory-based learning to data mining. *IEEE Transactions on Neural Networks*, 11(3):637–646, 2000.
- Ismail A Taha and Joydeep Ghosh. Symbolic interpretation of artificial neural networks. *IEEE Transactions on knowledge and data engineering*, 11(3):448–463, 1999.

- Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288, 1996.
- Geoffrey G Towell and Jude W Shavlik. Knowledge-based artificial neural networks. *Artificial intelligence*, 70(1-2):119–165, 1994.
- Michael Tsang, Dehua Cheng, and Yan Liu. Detecting statistical interactions from neural network weights. *arXiv preprint arXiv:1705.04977*, 2017.
- Berk Ustun and Cynthia Rudin. Supersparse linear integer models for optimized medical scoring systems. *Machine Learning*, 102(3):349–391, 2016.
- Jialei Wang, Ryohei Fujimaki, and Yosuke Motohashi. Trading interpretability for accuracy: Oblique treed sparse additive models. In *SIGKDD*, pages 1245–1254. ACM, 2015.
- Tong Wang. Multi-value rule sets for interpretable classification with feature-efficient representations. In *Advances in Neural Information Processing Systems*, pages 10835–10845, 2018.
- Tong Wang, Cynthia Rudin, Finale Doshi-Velez, Yimin Liu, Erica Klampfl, and Perry MacNeille. A bayesian framework for learning rule sets for interpretable classification. *The Journal of Machine Learning Research*, 18(1):2357–2393, 2017.
- Hongyu Yang, Cynthia Rudin, and Margo Seltzer. Scalable bayesian rule lists. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 3921–3930. JMLR. org, 2017.
- Jiaming Zeng, Berk Ustun, and Cynthia Rudin. Interpretable classification models for recidivism prediction. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 180(3):689–722, 2017.
- Quanshi Zhang, Ying Nian Wu, and Song-Chun Zhu. Interpretable convolutional neural networks. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- Yuchen Zhang and Lin Xiao. Stochastic primal-dual coordinate method for regularized empirical risk minimization. *The Journal of Machine Learning Research*, 18(1):2939–2980, 2017.
- Paul C Zikopoulos, Chris Eaton, Dirk DeRoos, Thomas Deutsch, and George Lapis. *Understanding big data: Analytics for enterprise class hadoop and streaming data*. Mcgraw-hill New York, 2012.