
Gaining Free or Low-Cost Transparency with Interpretable Partial Substitute

Tong Wang¹

Abstract

This work addresses the situation where a black-box model with good predictive performance is chosen over its interpretable competitors, and we show interpretability is still achievable in this case. Our solution is to find an interpretable substitute on a subset of data where the black-box model is *overkill* or nearly *overkill* while leaving the rest to the black-box. This transparency is obtained at minimal cost or no cost of the predictive performance. Under this framework, we develop a Hybrid Rule Sets (HyRS) model that uses decision rules to capture the subspace of data where the rules are as accurate or almost as accurate as the black-box provided. To train a HyRS, we devise an efficient search algorithm that iteratively finds the optimal model and exploits theoretically grounded strategies to reduce computation. Our framework is *agnostic* to the black-box during training. Experiments on structured and text data show that HyRS obtains an effective trade-off between transparency and interpretability.

1. Introduction

The deployment of machine learning in real-world applications has led to a surge of interest in systems optimized not only for task performance but also model interpretability, especially when human experts are involved in the decision-making process. In many heavily regulated industries such as judiciaries and health care, understanding the decision-making process of an analytical model is not just a preference but often a matter of legal and ethic compliance (Doshi-Velez & Kim, 2017), as recently promoted by EU’s General Data Protection Regulation (GDPR) (Parliament & of the European Union, 2016).

Thanks to this increasing interest, interpretable machine learning has achieved unprecedented advancement, provid-

ing two solutions for facilitating human understandability.

The first solution is to develop models that are interpretable by themselves, such as rule-based models, scoring models, case-based models, etc., of small or reasonable sizes. They are self-contained and do not rely on other models to explain them. We call them *interpretable* models in this paper and call the decision-making process **transparent** since humans understand wholly and precisely how a decision is generated. This solution is favorable when interpretable models perform as well as or better than black-box models (Choi et al., 2016). However, due to possible constraints on the model complexity to achieve interpretability, the lose of predictive performance for choosing interpretable models is often inevitable (Wang et al., 2015), needing another solution to obtain interpretability in this case.

To deal with situations when an interpretable model is inadequate, and a black-box has to be chosen for better predictive performance, the second solution is proposed to develop models that explain black-boxes. These models generate post hoc explanations or approximations for a black-box either locally (Ribeiro et al., 2016) or globally (Adler et al., 2016; Lakkaraju et al., 2017), providing some insights into the black-box model by identifying key features or interactions of features (Tsang et al., 2017). However, two concerning issues exist. First, explainers only approximate but do not characterize exactly the decision-making process of a black-box model, yielding an imperfect explanation fidelity. Second, 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. Both issues result from the fact that the explainers only approximate in a post hoc way but are not the decision-making process themselves.

In this paper, we create an alternative solution to gain interpretability in the presence of a chosen black-box. We propose to use interpretable partial substitute to process a subset of data, where an interpretable model is adequate for producing predictions that are as good as the black box, i.e., where the black-box is *overkill*, to obtain free interpretability at no cost of the predictive performance; or, if the user is willing to trade some accuracy for transparency, our model can find the right subset of data at minimal cost of predic-

¹Department of Business Analytics, University of Iowa, Iowa, USA. Correspondence to: Tong Wang <tong-wang@uiowa.edu>.

tive performance. Thus, on this subset of data, the model gains transparency with 100% fidelity to replace otherwise non-perfect approximations by an explainer. We define the percentage of the subset *transparency* of the model.

To summarize, we design a novel framework that integrates an interpretable model with the black-box model into a sequential decision-making process. An input instance first goes through an interpretable model. If the model is competent, a prediction will directly be generated. Otherwise, the black box will be activated. We call the proposed model a *Hybrid Predictive Model*. Under this framework, we build Hybrid Rule Sets (HyRS) where we use association rules as interpretable local substitutes.

This form of the model is motivated by how humans make decisions in many real-world situations. For example, when a doctor diagnoses a patient, if the patient is an irregular case with symptoms that do not match documented descriptions of any disease, then an experienced doctor or a consultation of several experts will be summoned, representing a complicated model (black-box model). However, if the patient demonstrates regular “textbook” symptoms, a diagnosis can be made right away via standard symptom matching and reasoning (interpretable model) by a simpler “model” such as a resident or a nurse. In these cases, residents and nurses can perform nearly well as experienced doctors since there are symptoms easy to explain. For a hospital, it is not economical to always request a consultation from several experts to treat simple cases. Thus, hospitals often stratify patients and send more complicated cases to more experienced doctors. Our model works similarly, and our goal is to construct an interpretable model that replaces the complicated black-box model on an appropriate subspace of data.

The benefits of a hybrid model include several aspects. *First*, it gains transparency of the model at no cost or minimal cost of predictive performance. Compared to explainers that provide post hoc analysis, an interpretable substitute is understandable by itself. There does not exist any ambiguity or inconsistency since an interpretable model is entirely faithful to itself. *Second*, rules only use a small set of features while the black-box model needs to use all. In some applications where features are costly to get (e.g., medical test results in hospitals), using a small set of features reduces cost. *Finally*, from an operational viewpoint, a black-box model can be a cumbersome system with high costs in operating and maintaining. Using a simpler interpretable model will save computing resources and time in real applications.

In this paper, we formulate a general framework and learning objective for building a hybrid model. The objective considers predictive accuracy, transparency, and interpretability. The model learns to capture the correct subset of data and achieves the best balance between sending enough data to the interpretable model and preserving predictive perfor-

mance. To train the model, we devise an efficient search algorithm that exploits the theoretical properties of the model to reduce computation complexity. We develop three bounds to reduce the search space, one applied before the search begins and two dynamic bounds during the search.

In the rest of the paper, we will review related work in Section 2. We present the general framework for learning a hybrid decision model in Section 3 and develop a Hybrid Rule Sets (HyRS) model under the proposed framework in Section 4. We design an efficient training algorithm that exploits theoretically grounded strategies for fast computation in Section 5. Finally, we provide a detailed experimental evaluation in Section 6. We conclude the paper in Section 7.

2. Related Work

Our work is broadly related to new methods for interpretable machine learning. There have been two lines of research in interpretable machine learning. The first is developing models that are interpretable stand-alone. Previous work in this category include rule-based models such as fuzzy rules (Alonso et al., 2018a), rule sets (Rijnbeek & Kors, 2010; McCormick et al., 2011; Dash et al., 2018b) and rule lists (Yang et al., 2017; Angelino et al., 2017)), scoring systems (Zeng et al., 2017; Ustun & Rudin, 2016; Koh et al., 2015), and etc. The second line of research is on developing explainer models that explain black-boxes locally (Ribeiro et al., 2016) or globally (Adler et al., 2016; Lakkaraju et al., 2017). One representative work is LIME (Ribeiro et al., 2016) that explains the predictions of any classifier by learning a linear model locally around the prediction. More recently, developments in deep learning have been connected strongly with interpretable machine learning and have contributed novel insights into representational issues. Recent works have proposed high-level symbolic representations used in knowledge representation (Yi et al., 2018).

Our work is fundamentally different from the research above. A hybrid decision model is not a pure interpretable model. It uses an interpretable substitute on a subset of data. It is also not a diagnostic model that only observes but does not participate in the decision process. Here, a hybrid decision model uses an interpretable and a black-box model simultaneously in decision making and utilizing the strength of interpretable models in producing understandable predictions.

There exist a few singleton works that combine multiple models. For example, (Kohavi, 1996) proposed NBTree which induces a hybrid of decision-tree classifiers and Naive-Bayes classifiers, (Shin et al., 2000) proposed a system combining neural network and memory-based learning, (Hua & Zhang, 2006) combined SVM and logistic regression to forecast intermittent demand of spare parts, etc. A recent work (Alonso et al., 2018b) builds a black-box oracle which outputs the simplest interpretable model to produce a sim-

ilar prediction a black-box model would generate. There exist potential inconsistency issues since there are multiple interpretable models likely to be selected. Another work (Wang et al., 2015) divides feature spaces into regions with sparse oblique tree splitting and assign local sparse additive experts to the individual areas. Our model is distinct in that the proposed framework can work with *any* black-box model and is *agnostic* to the model during training.

3. A General Framework for Hybrid Models

We present a general framework for building a hybrid model and define a principled objective function. We start with a set of training examples $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^N$ where $\mathbf{x}_i \in \mathcal{X}$ is a tuple of attributes and $y_i \in \{1, 0\}$ is the class label. Let $f = \langle f_l, f_b \rangle$ represent a hybrid model that consists of an interpretable model f_l and a black-box model f_b . f is agnostic to f_b and only needs the predictions of f_b on \mathcal{D} as input, denoted as $\mathcal{Y}_b = \{\hat{y}_{b_i}\}_i^N$.

A critical issue in designing a hybrid model is how to automatically distribute data to f_l and f_b . This is equivalent to creating a partition of the dataset \mathcal{D} to \mathcal{D}_l and \mathcal{D}_b , corresponding to training examples sent to f_l and f_b , respectively. We design the predictive process as below: an input instance \mathbf{x}_k is first sent to the interpretable model f_l . If a prediction can be made, an output $\hat{y}_{l,k}$ is directly generated. Otherwise, it is sent to f_b to generate a prediction $\hat{y}_{b,k}$. $\hat{y}_{l,k} \in \mathcal{D}_l$ and $\hat{y}_{b,k} \in \mathcal{D}_b$. See the predictive process in Figure 1.




Figure 1. A predictive process of a hybrid model.

Our goal is to construct an interpretable model f_l to be combined with f_b , with the objective that considers three fundamental properties: 1) The **predictive accuracy**. Since f_b is pre-given, the predictive performance of f depends on the predictive accuracy of f_l independently and the collaboration of f_l and f_b , i.e., the partition of \mathcal{D} to \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 both models if the training examples are partitioned strategically, sending examples to the model which can predict them correctly. In some circumstance, the combination of a weak and a strong model can yield performance better than the strong model alone. 2) The **interpretability** of f_l . Bringing interpretability into the decision process is one of the motivations for building a hybrid model. Therefore, small size and low complexity are much-desired properties of f_l . The defini-

tion of interpretability is model specific and usually refers to using a small number of cognitive chunks (Doshi-Velez & Kim, 2017). 3) The **transparency** of the hybrid model. This is a new metric we propose for the hybrid framework to capture the percentage of data that are processed by f_l , i.e., 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 $\mathcal{T}(f, \mathcal{D})$.

We formulate the learning objective for building a hybrid decision model as a linear combination of the three metrics above. This framework unifies interpretable models and black-box models: interpretable models have transparency of one, and black-box models have transparency of zero.

4. Hybrid Rule Sets Model

Under the proposed framework, we instantiate a hybrid decision model. Here we take a significant step towards interpretability by choosing rules for f_l . Rules are easy to understand due to their simple logic and symbolic presentation. They also naturally handle the partition of data by separating examples according to if they satisfy the rules.

Now we present the Hybrid Rule Sets (HyRS) model. A HyRS model consists of two sets of rules. The first set of rules captures positive instances, called the *positive rule set* and denoted as \mathcal{R}_+ . The second set of rules captures negative instances, called the *negative rule set* and denoted as \mathcal{R}_- . If \mathbf{x}_k satisfies any positive rules, it is classified as positive. Otherwise, if it satisfies any negative rules, it is classified as negative. Denote $\mathcal{R} = \mathcal{R}_+ \cup \mathcal{R}_-$. A decision produced from \mathcal{R} is denoted as $\hat{y}_{l,k}$. If \mathbf{x}_k does not satisfy any rules in \mathcal{R}_+ or \mathcal{R}_- , it means f_l fails to decide on \mathbf{x}_k . Then \mathbf{x}_k is sent to the black-box model f_b to generate a decision $f_b(\mathbf{x}_i)$. \mathcal{D}_l is a set of instances sent to f_l . In the context of a HyRS model, we use f_l and \mathcal{R} interchangeably when we refer to the interpretable model. We summarize the decision-making process below.

if \mathbf{x}_i obeys \mathcal{R}_+ , $Y = 1$
else if \mathbf{x}_i obeys \mathcal{R}_- , $Y = 0$
else $Y = f_b(\mathbf{x}_i)$

We show an example of a HyRS model in Table 1 learned from a heart disease dataset from UCI ML repository (Lichman, 2013). In this model, there are two rules in \mathcal{R}_+ and one rule in \mathcal{R}_- . $Y = 1$ represents the patient has heart disease.

To formulate the problem, we need the following notations and definitions.

Definition 2 A rule r covers an example \mathbf{x}_i if \mathbf{x}_i obeys the rule, denoted as $\text{covers}(r, \mathbf{x}_i) = 1$. A rule set R covers an example \mathbf{x}_i if \mathbf{x}_i obeys at least one rule in R , i.e.

Table 1. An example of a HyRS model

	Rules	Model
if	age < 35 and maximum heart rate ≥ 178 OR serum cholesterol ≥ 234 and thal $\neq 3$ and the number of vessels ≥ 1 $\rightarrow Y = 1$ (heart disease)	\mathcal{R}_+
else if	chest pain type $\neq 4$ and age > 40 $\rightarrow Y = 0$ (no heart disease)	\mathcal{R}_-
else	$\rightarrow Y = f_b(\mathbf{x})$	f_b

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

Definition 3 Given a data set \mathcal{D} , the support of a rule set R in \mathcal{D} is a set of observations covered by R , i.e., $\text{support}(R, \mathcal{D}) = \{i | \text{covers}(R, \mathbf{x}_i) = 1, \mathbf{x}_i \in \mathcal{D}\}$.

We formulate the objective function for HyRS following the objective described in the previous section. First, we measure the misclassification error to represent the predictive performance. Given rules \mathcal{R} , a black-box model f_b and data \mathcal{D} , the misclassification error is

$$\begin{aligned} \ell(\langle \mathcal{R}, f_b \rangle, \mathcal{D}) &= \sum_{i=1}^N \left((1 - y_i) \text{covers}(\mathcal{R}_+, \mathbf{x}_i) \triangleright \text{errors from } \mathcal{R}_+ \right. \\ &+ y_i \left(1 - \text{covers}(\mathcal{R}_+, \mathbf{x}_i) \right) \text{covers}(\mathcal{R}_-, \mathbf{x}_i) \triangleright \text{errors from } \mathcal{R}_- \\ &+ (1 - \text{covers}(\mathcal{R}_+, \mathbf{x}_i)) (1 - \text{covers}(\mathcal{R}_-, \mathbf{x}_i)) \\ &\times (y_i(1 - \hat{y}_{b_i}) + (1 - y_i)\hat{y}_{b_i}) \Big) / N. \quad \triangleright \text{errors from } f_b \end{aligned}$$

Definition 4 $\Omega(R)$ is the interpretability of R .

There are various definitions of interpretability for rule-based models, such as the number of rules (Dash et al., 2018a; Wang et al., 2017; Lakkaraju et al., 2016; Wang, 2018), the number of conditions, and those discussed in (Gacto et al., 2011). The definition of interpretability is domain specific and subject to applications. In this paper, we choose the number of rules in R to represent interpretability as an illustration of the framework but our formulation works with other interpretability measures in the literature.

We aim to minimize the objective function combining predictive accuracy, model interpretability and transparency.

$$\Lambda(\langle \mathcal{R}, f_b \rangle, \mathcal{D}) = \ell(\langle \mathcal{R}, f_b \rangle, \mathcal{D}) + \theta_1 \Omega(\mathcal{R}) - \theta_2 \frac{\text{support}(\mathcal{R}, \mathcal{D})}{N} \quad (1)$$

where the transparency of a HyRS model follows definition 1. Here, θ_1 and θ_2 are non-negative coefficients. Tuning the parameters will produce models at different operating points of accuracy, interpretability, and transparency. For example, in an extreme case when $\theta_2 \gg \theta_1$, the output will be a model that sends all data to f_b , producing a pure interpretable model. When $\theta_1 \gg \theta_2$ and $\theta_1 \gg 1$, then the model will force f_b to have complexity 0, i.e., producing a pure black-box model.

5. Model Training

We describe a training algorithm to find an optimal solution f^* such that

$$f^* \in \arg \min_f \Lambda(f; \mathcal{D}) \quad (2)$$

Since $f^* = \langle \mathcal{R}^*, f_b \rangle$ and f_b is fixed, the problem reduces to finding an optimal rule set model $\mathcal{R}^* = \mathcal{R}_+^* \cup \mathcal{R}_-^*$ that covers the correct subset of data in the presence of f_b .

Learning rule-based models is challenging because the solution space (all possible rule sets) is a power set of the rule space. Fortunately, our objective has a nice structure that can be exploited for reducing computation.

Algorithm structure The algorithm is presented in Algorithm 1. Given training examples, \mathcal{D} , a black-box model f_b , parameters θ_1, θ_2 , base temperature C_0 and the total number of iterations T , the search procedure follows the main structure of a stochastic local search algorithm. Each state corresponds to a rule set model, indexed by the time stamp t , denoted as $\mathcal{R}_{[t]}$. 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 adding or removing a rule from the current set. At each iteration, the algorithm improves one of the three terms (accuracy, interpretability, and transparency) with approximately equal probabilities, by removing or adding a rule to the current model. The proposed neighbor is accepted with probability $\exp\left(\frac{\Lambda(\mathcal{R}_{[t]}) - \Lambda(\mathcal{R}_{[t+1]})}{C_0^{1 - \frac{t}{T}}}\right)$ which gradually decreases as the temperature cools down.

Rule Space Pruning We first use FP-growth¹ to generate a set of candidate rules, Υ^+ for positive rules and Υ^- for negative rules. The algorithm will search only within this rule space. Since the number of rules grows exponentially with the number of features, the complexity of the algorithm is directly determined by the size of the rule space. To facilitate faster computation, we derive a lower bound on the support of rules to prune the rule space. All proofs are in the supplementary material.

Theorem 1 (Lower Bound on Support) $\forall r \in \mathcal{R}_+^*, \text{support}(r) \geq N\theta_1; \forall r \in \mathcal{R}_-^*, \text{support}(r) \geq \frac{N\theta_1}{1 - \theta_2}$.

This means \mathcal{R}^* does not contain rules with too small a support. This theorem is used before the search begins to prune the rule space to only contain rules with large support, excluding unqualified rules from consideration, which greatly reduces computation without hurting the optimality. On the other hand, removing rules with low support naturally helps prevent overfitting. The bound increases as θ_1 increases, since θ_1 represents the penalty for adding a rule.

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

Algorithm 1 Stochastic Local Search algorithm

```

1: Input:  $f_b, \mathcal{D}, \theta_1, \theta_2, C_0$ 
2: Initialize:
3:  $\mathcal{R}^* = \mathcal{R}^{[0]} \leftarrow \emptyset$  ▷ start with a pure black-box model
4:  $\Upsilon^+ \leftarrow \text{FPGrowth}(\mathcal{D}, \text{minsupp} = N\theta_1)$ 
5:  $\Upsilon^- \leftarrow \text{FPGrowth}(\mathcal{D}, \text{minsupp} = \frac{N\theta_1}{1-\theta_2})$ 
   ▷ mine candidate rules from  $\mathcal{D}$  using minimum support from
   Theorem 1
6: for  $t = 0 \rightarrow T$  do
7:    $\delta = \text{random}()$ 
8:   if  $\delta \leq \frac{1}{3}$  or  $\Omega(\mathcal{R}_{[t]}) \geq \frac{\lambda_{[t]}^* + \theta_2}{\theta_1}$  then
9:      $\mathcal{R}_{[t+1]} \leftarrow$  remove a rule from  $\mathcal{R}_{[t]}$  ▷ decrease the size of
      $\mathcal{R}_{[t]}$  and use Theorem 2
10:  else if  $\delta \leq \frac{2}{3}$  or  $\text{support}(\mathcal{R}_{[t]}) \leq \frac{\theta_1 - \lambda_{[t]}^*}{\theta_2}$  then
11:     $\mathcal{R}_{[t+1]} \leftarrow$  add a rule to  $\mathcal{R}_{[t]}$  ▷ increase the transparency
    and use Theorem 3
12:  else
13:     $\epsilon \leftarrow \{k | f_{[t]}(\mathbf{x}_k) \neq y_k\}$  ▷ indices of misclassified exam-
    ples.
14:    if  $\epsilon \in \text{covrg}(\mathcal{R}_{[t]})$  then
15:      if  $\epsilon$  is negative then
16:         $\mathcal{R}_{[t+1]} \leftarrow$  remove a rule from  $\mathcal{R}_{+[t]}$  that covers  $\epsilon$ .
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:       $\mathcal{R}_{[t+1]} \leftarrow$  add a rule to  $\mathcal{R}_{[t]}$  that is consistent with the
      sign of  $\epsilon$ 
22:    end if
23:  end if
24:  accept  $\mathcal{R}_{[t+1]}$  with probability  $\exp(\frac{\Lambda(\mathcal{R}_{[t]}) - \Lambda(\mathcal{R}_{[t+1]})}{C_0^{\frac{1}{T}}})$ 
25:   $\mathcal{R}^* = \arg \min_{\mathcal{R}_{[t+1]}, \mathcal{R}^*} \Lambda(\mathcal{R})$  ▷ update the best solution
26: end for
27: Output:  $\mathcal{R}^*$ 

```

Search Chain Bounding We also derive two bounds that reduce the search space during the search. The bounds are applied in each iteration to confine the Markov Chain within promising solution space, preventing it from going too far and wasting too much search time. First, we derive a bound on the size of the model. 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 2 (Upper Bound on Size) $\Omega(\mathcal{R}^*) \leq \frac{\lambda_{[t]}^* + \theta_2}{\theta_1}$.

This theorem says that the size of an optimal model is upper bounded, which means the Markov Chain only needs focus on solution space of small models. Therefore, in the proposing step, if the current state violates the bound, the next state should be proposed by removing a rule from the current model.

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

Theorem 3 (Lower Bound on Transparency)

$$\text{support}(\mathcal{R}^*) \geq \frac{\theta_1 - \lambda_{[t]}^*}{\theta_2}.$$

The theorem says the transparency of an optimal model is lower bounded. Therefore whenever it is violated, the next proposal should be adding a rule to the current state.

Both bounds become smaller as $\lambda_{[t]}^*$ continuously gets smaller. Exploiting the theorem in the search algorithm, we check the intermediate solution at each iteration and pull the search chain back to promising an area (models of sizes smaller than $\frac{\lambda_{[t]}^* + \theta_1}{\theta_2}$ and models with transparency larger than $\frac{\theta_1 - \lambda_{[t]}^*}{\theta_2}$) whenever the bounds are violated (line 8 and line 10 in Algorithm 1).

Now we detail the proposing step below.

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 upper bound of the model in Theorem 2 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 is violated, we aim to increase 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 13-22). 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, we remove a rule from the positive that covers it. If the instance is positive, we add a rule to $\mathcal{R}_{+[t]}$ to cover it or remove a rule in the $\mathcal{R}_{-[t]}$ that covers it, re-routing it to f_b . If the example is not covered by $\mathcal{R}_{[t]}$, 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, re-routing it to f_i .

When choosing a rule to add or remove, we first evaluate the rules using precision, which is the percentage of correctly classified examples of a rule. Then we balance between exploitation, choosing the best rule, and exploration, choosing a random rule, to avoid getting in a local minimum.

6. Experiments

We perform a detailed experimental evaluation of HyRS using public datasets and a real-world application. We compare HyRS with state-of-the-art interpretable and black-box baselines.

6.1. Experiments on Public Datasets

The goal of the first set of experiments is to examine the accuracy and transparency of HyRS as well as their relationships. We use public datasets from domains interpretability is most pursued.

Datasets We use four structured datasets and a text dataset from domains where interpretability is highly desired, including healthcare, judiciaries and customer analysis. 1) *juvenile*(Osofsky, 1995) (4023 observations and 55 reduced features), to study the consequences of juvenile exposure to violence. The dataset was collected via a survey sent to juveniles. 2) *credit card* (30k observations and 23 features), to predict the default of credit card payment (Yeh & Lien, 2009) 3) *recidivism* (11,645 observations and 106 features) to predict if a criminal will re-offend after he is released from prison 4) *readmission* (100,000 observations and 34 features) to predict readmission of patients with diabetes. 5) *Yelp review* (Kotzias et al., 2015) (1,000 observations) that contains positive and negative reviews from Yelp. The goal is to do sentiment classification.

Implementation We process each dataset by binarizing categorical features and discretizing real-valued features with four cut-off points. For each structured dataset, we build three black-box models that are often the top performing models, Random Forests (Liaw et al., 2002), AdaBoost (Freund & Schapire, 1995) and extreme gradient boosting trees (XGBoost) (Chen & Guestrin, 2016). For the text classification, we build a Long Short-Term Memory (LSTM) neural network that consists of one embedding layer with an embedding vector of 32, one layer of 100 LSTM units, and two fully connected layers following it. We partition each dataset into 80% training and 20% testing. We do cross-validation for parameter tuning on the training set and evaluate the best model on the test set. The predictive performance of the black-box models are reported in Table 2.

Table 2. Test accuracy of black-box models

Models	Juvenile	Credit card	Recidivism	Diabetes	Yelp
RF	.91	.82	.73	.64	-
AdaBoost	.90	.82	.69	.64	-
XGBoost	.90	.82	.74	.64	-
LSTM	-	-	-	-	0.76

STUDY 1: FREE TRANSPARENCY

The goal of the analysis is to find the maximally achievable “free” transparency (no cost of the predictive performance), i.e., finding the area where a black-box is overkill.

We tune the parameters θ_1, θ_2 to get a set of models for each dataset. θ_1 controls the number of rules and is chosen from [0.001, 0.01]. θ_2 controls transparency and we choose

θ_2 from 0 to 1. We report in Table 3 the maximal free transparency. The performance is evaluated on test sets.

Table 3. Maximally achievable free transparency

	Juvenile	Credit Card	Recidivism	Diabetes	Yelp
$\langle \cdot, \text{RF} \rangle$.82	.91	.82	.10	-
$\langle \cdot, \text{AdaBoost} \rangle$.61	.89	.79	.26	-
$\langle \cdot, \text{XGBoost} \rangle$.78	.90	.65	.20	-
$\langle \cdot, \text{LSTM} \rangle$	-	-	-	-	.43

The rules are able to explain on average roughly 80% of the data for juvenile, credit card and recidivism, and about 20% for diabetes dataset. The results show that our fundamental assumption is true - there exists a subspace where the interpretable model is as accurate than the black-box model, even if the black-box model is better globally. Meanwhile, this free interpretability is obtained using only a few rules (see Table 4). All models use less than 5 rules in total. For the Yelp dataset, we select the top 2000 words with highest frequency as the input.

Table 4. The number of rules in HyRS models

	Juvenile	Credit Card	Recidivism	Diabetes	Yelp
$\langle \cdot, \text{RF} \rangle$	2	4	2	4	-
$\langle \cdot, \text{AdaBoost} \rangle$	1	3	2	5	-
$\langle \cdot, \text{XGBoost} \rangle$	1	4	1	5	-
$\langle \cdot, \text{LSTM} \rangle$	-	-	-	-	2

Examples of HyRS models We show two examples of HyRS models built from the datasets above. The first example is from juvenile dataset. The data was collected from a survey so the features are questions and feature values are answers to that question. The positive rule set is an empty set, and the negative rule set consists of one rule.

if Has any of your family members or friends ever attacked you with a weapon \neq Yes *and* Have your friends ever hit or threatened to hit someone without any reason? \neq Yes *and* Have your friends ever broken into a vehicle or building to steal something \neq Yes
then $Y = 0$
else $Y = f_b(\mathbf{x})$

Note that this one rule captures 78% of the instances and the overall predictive accuracy is just as accurate as of the black-box model.

The second model we show is built from the text classification using the Yelp review data. When we build HyRS on the text data, each unique word is a feature, and a rule is a phrase (conjunction of features), i.e., word 1 and word 2 and ... all appear in the text. Therefore, a rule set contains a set of words and phrases. On this dataset, the rules are all words. We show the HyRS model and the words for the positive set and the negative set in Figure 3.

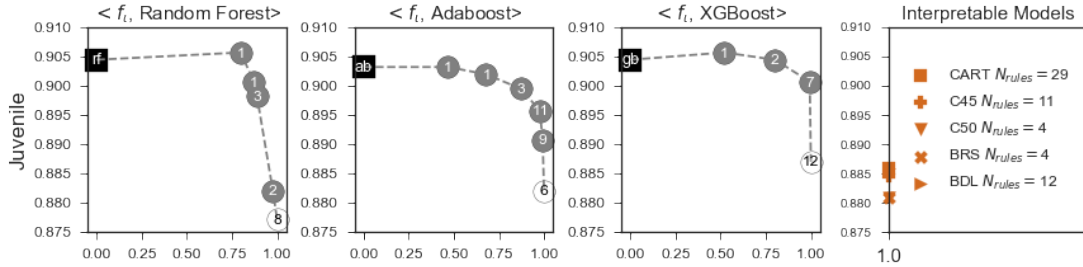


Figure 2. The trade-off between transparency and accuracy for HyRS on Juvenile dataset. The black-squared represent black-box models. Grey circles represent HyRS models and the transparent circle represent HyRS models with transparency equal to one (reduced to interpretable models).

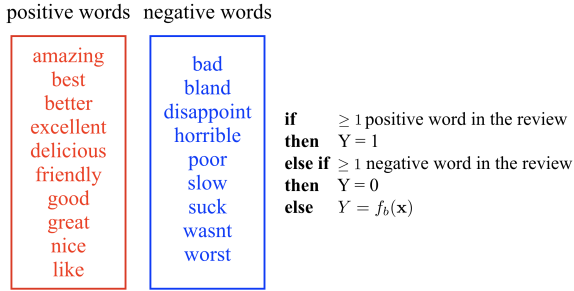


Figure 3. Positive words and negative words mined from Yelp review dataset. The HyRS model built from these words captures 43% of the reviews without losing predictive accuracy compared to the LSTM model.

STUDY 2: ACCURACY-TRANSPARENCY TRADE-OFF

In addition to studying the maximally achievable “free” transparency, we study if some loss in predictive performance is tolerable, how can that be efficiently utilized and transformed into transparency. To visualize this trade-off, we plot models’ transparency as the X-axis and accuracy as the Y-axis, both evaluated on test sets. Here we choose juvenile dataset for demonstration. We tune θ_1 from 0.001 to 0.01 and θ_2 from 0 to 1 and only show the models on the frontier of the curves. See Figure 2.

Baselines We benchmark the performance of HyRS against other rule-based interpretable models, C4.5 (Quinlan, 2014) and C5.0 (Kuhn et al., 2014), Scalable Bayesian Rule Lists (SBRL) (Yang et al., 2017) and Bayesian Rule Sets (BRS) (Wang 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. See a description of parameter tuning in the supplementary material. We find the models with the highest cross-validated accuracy and show their test accuracy in Figure 2. All models are at the transparency of 1 since they are interpretable.

The curves remain almost flat when the transparency is smaller than roughly 70% and start to drop quickly after

that. It means there exists a large subset of data that can be captured by simple rules (only one or two as shown in Figure 2). The rules perform as well as black-box models. But the remaining 30% is much harder to be characterized by rules. Thus the number of rules increases quickly as transparency approaches 1, and the accuracy drops to values comparable to the interpretable baselines.

6.2. Application to Medical Crowdfunding Prediction

We apply HyRS to a real-world application, medical crowdfunding prediction, using real-world data. Medical crowdfunding is a type of donation-based crowdfunding, helping users raise funds to pay medical bills by collecting small donations from many people. A crowdfunding could last for several weeks. But since medical crowdfunding is time-critical in its nature, an early prediction of its outcome, especially future failure is very valuable since users can look for alternatives channels of funding early if they could not raise enough money from the current platform. We want to predict the failure of fundraising, which is defined as raising less than 10% of the target amount. The data is provided by medical crowdfunding company and consists of 51,228 cases from October 2016 to June 2018.

There are two types of features, time-invariant and time series. When a fundraiser creates a new case, he submits patient information including demographics (age, gender, city, etc.), insurance status (commercial or basic medical insurance), the target amount, verification from the patient’s hospital, along with a short “call for donation” post. These are time-invariant features. Then, after a case is published, a set of features are collected daily. The daily features include the number of times a case is shared on social networks, the number of times a fundraiser responds to users’ questions and requests for more information (or proof) on the case pages, the number of views for the case, the number of users who verified the case, the number of users who donate, and etc. These are time series features that are collected daily. We use the observations for the first week after a case is published.

Table 5. A HyRS model for predicting outcomes for medical crowdfunding. The accuracy is 0.77 and transparency is 0.82.

	Rules	Model
if	day of launch (of a month) < 6 <i>and</i> approval status \neq approved <i>and</i> fundraiser’s gender = missing OR month of launch = June OR day of launch (of a month) < 18 OR content length < 519 words <i>and</i> target amount \geq 44,000	\mathcal{R}_+
then	$Y = 1$ (failure)	
else if	content length < 334 words <i>and</i> patient age < 44 <i>and</i> target amount < 44,000 OR approval status = approved <i>and</i> target amount \leq 29,000 <i>and</i> month \neq June OR content length \geq 519 words <i>and</i> patient age < 44 <i>and</i> month \neq June	\mathcal{R}_-
then	$Y = 0$ (success)	
else	$Y = f_b(\mathbf{x})$	f_b

We design a deep neural network that first uses two layers of LSTM (20 units and 40 units) to process the time series features, and then the output is merged with the time-invariant features at two fully connected layers. The last layer uses a sigmoid activation function to produce a prediction for failure. We partition the dataset into 80% training and 20% testing. We train the network for 200 epochs, and the test accuracy is 0.96.

Then we build a HyRS model. To ensure interpretable and early prediction, here we restrict our model only to use time-invariant features that are immediately available after a user launches a case. We would like to know how accurately can HyRS provide an early prediction only using a few features.

To understand the trade-off between transparency and accuracy, we set θ_1 to 0.001 and vary θ_2 from 0 to 1 to obtain a curve in Figure 4. The numbers of rules are represented by the size of the markers and are also annotated in the figure.

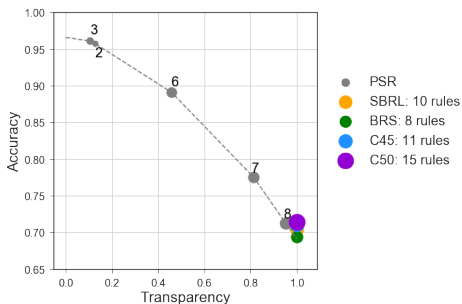


Figure 4. The accuracy vs transparency curve for HyRS models for predicting fundraising failure for medical crowdfunding.

For the baseline models, we use the same interpretable models as described in Section 6.1. We follow the same steps tuning the parameters. Their accuracy and the number of rules are shown in Figure 4.

Discussions This curve characterizes the trade-off between transparency and accuracy. As transparency increases, HyRS covers more and more data with rules, at the cost of predictive accuracy, but always higher than pure interpretable models alone. The users can decide the best oper-

ating point based on their specific requirement of accuracy using the curve. Compared to using either purely black-box models or purely interpretable models, HyRS provides more options for model selections.

We show a model that achieves transparency of 0.82 and an accuracy of 0.77 in Table 5. This model consists of four positive rules and three negative rules. Only 18% instances are not captured and therefore sent to LSTM for decision.

7. Conclusions

We proposed a general framework for learning a hybrid model that integrates an interpretable partial substitute with any black-box model to introduce transparency into the predictive process at no or low cost of the predictive performance. We instantiated this framework with Hybrid Rule Sets Hybrid (HyRS) model using rules as the interpretable component. Experiments demonstrated partial transparency is possible in the presence of black-box models. It suggests that in some cases, always using a black-box is overkill, and replacing with a simpler and interpretable model will save resources and provide partial transparency.

The HyRS model is one example of the proposed hybrid model framework. An important contribution of this work is that we proposed a general framework for combining an interpretable substitute model with a black-box model. Our framework can support the exploration of a variety of interpretable models, such as linear models (see our latest paper (Wang & Lin, 2019)), decision trees and prototype-based models.

The proposed framework provides a new solution when one wishes not to give up the high predictive accuracy of black-box models. A hybrid model can serve as a pre-step for a black-box explainer: we first find a region that can be captured and explained by an interpretable model and then sends the rest of the data to a black-box to predict and an explainer to explain.

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

References

- Adler, P., Falk, C., Friedler, S. A., Rybeck, G., Scheidegger, C., Smith, B., and Venkatasubramanian, S. Auditing black-box models for indirect influence. In *ICDM*, pp. 1–10. IEEE, 2016.
- Alonso, J. M., Castiello, C., and Mencar, C. A bibliometric analysis of the explainable artificial intelligence research field. In *International Conference on Information Processing and Management of Uncertainty in Knowledge-Based Systems*, pp. 3–15. Springer, 2018a.
- Alonso, J. M., Soto, A. R., Castiello, C., and Mencar, C. Hybrid data-expert explainable beer style classifier. 2018b.
- Angelino, E., Larus-Stone, N., Alabi, D., Seltzer, M., and Rudin, C. Learning certifiably optimal rule lists. In *SIGKDD*, pp. 35–44. ACM, 2017.
- Chen, T. and Guestrin, C. Xgboost: A scalable tree boosting system. In *SIGKDD*, pp. 785–794. ACM, 2016.
- Choi, E., Bahadori, M. T., Sun, J., Kulas, J., Schuetz, A., and Stewart, W. Retain: An interpretable predictive model for healthcare using reverse time attention mechanism. In *Advances in Neural Information Processing Systems*, pp. 3504–3512, 2016.
- Dash, S., Gunluk, O., and Wei, D. Boolean decision rules via column generation. In *Advances in Neural Information Processing Systems 31*, pp. 4660–4670. Curran Associates, Inc., 2018a.
- Dash, S., Gunluk, O., and Wei, D. Boolean decision rules via column generation. In *Advances in Neural Information Processing Systems*, pp. 4655–4665, 2018b.
- Doshi-Velez, F. and Kim, B. A roadmap for a rigorous science of interpretability. *arXiv preprint arXiv:1702.08608*, 2017.
- Freund, Y. and Schapire, R. E. A decision-theoretic generalization of on-line learning and an application to boosting. In *European conference on computational learning theory*, pp. 23–37. Springer, 1995.
- Gacto, M. J., Alcalá, R., and Herrera, F. Interpretability of linguistic fuzzy rule-based systems: An overview of interpretability measures. *Information Sciences*, 181(20): 4340–4360, 2011.
- Hua, Z. and Zhang, B. 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.
- Koh, H. C., Tan, W. C., and Goh, C. P. A two-step method to construct credit scoring models with data mining techniques. *International Journal of Business and Information*, 1(1), 2015.
- Kohavi, R. Scaling up the accuracy of naive-bayes classifiers: A decision-tree hybrid. In *KDD*, volume 96, pp. 202–207, 1996.
- Kotzias, D., Denil, M., De Freitas, N., and Smyth, P. From group to individual labels using deep features. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 597–606. ACM, 2015.
- Kuhn, M., Weston, S., Coulter, N., and Quinlan, R. C50: C5.0 decision trees and rule-based models. *R package version 0.1.0-21*, 50, 2014.
- Lakkaraju, H., Bach, S. H., and Leskovec, J. Interpretable decision sets: A joint framework for description and prediction. In *SIGKDD*. ACM, 2016.
- Lakkaraju, H., Kamar, E., Caruana, R., and Leskovec, J. Interpretable & explorable approximations of black box models. *arXiv preprint arXiv:1707.01154*, 2017.
- Liaw, A., Wiener, M., et al. Classification and regression by randomforest. *R news*, 2(3):18–22, 2002.
- Lichman, M. Uci machine learning repository, 2013. URL <http://archive.ics.uci.edu/ml>.
- Lissack, M. Dealing with ambiguity—the black box as a design choice. *SheJi (forthcoming)*, 2016.
- McCormick, T., Rudin, C., and Madigan, D. A hierarchical model for association rule mining of sequential events: An approach to automated medical symptom prediction. 2011.
- Osofsky, J. D. The effect of exposure to violence on young children. *American Psychologist*, 50(9):782, 1995.
- Parliament and of the European Union, C. General data protection regulation, 2016.
- Quinlan, J. R. *C4.5: programs for machine learning*. Elsevier, 2014.
- Ribeiro, M. T., Singh, S., and Guestrin, C. Why should i trust you?: Explaining the predictions of any classifier. In *SIGKDD*, pp. 1135–1144. ACM, 2016.
- Rijnbeek, P. R. and Kors, J. A. Finding a short and accurate decision rule in disjunctive normal form by exhaustive search. *Machine learning*, 80(1):33–62, 2010.

- Ross, A. S., Hughes, M. C., and Doshi-Velez, F. Right for the right reasons: Training differentiable models by constraining their explanations. *arXiv preprint arXiv:1703.03717*, 2017.
- Shin, C.-K., Yun, U. T., Kim, H. K., and Park, S. C. A hybrid approach of neural network and memory-based learning to data mining. *IEEE Transactions on Neural Networks*, 11(3):637–646, 2000.
- Tsang, M., Cheng, D., and Liu, Y. Detecting statistical interactions from neural network weights. *arXiv preprint arXiv:1705.04977*, 2017.
- Ustun, B. and Rudin, C. Supersparse linear integer models for optimized medical scoring systems. *Machine Learning*, 102(3):349–391, 2016.
- Wang, J., Fujimaki, R., and Motohashi, Y. Trading interpretability for accuracy: Oblique treed sparse additive models. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 1245–1254. ACM, 2015.
- Wang, T. Multi-value rule sets for interpretable classification with feature-efficient representations. In *Advances in Neural Information Processing Systems 31*, pp. 10858–10868. 2018.
- Wang, T. and Lin, Q. Hybrid predictive model: When an interpretable model collaborates with a black-box model. 2019.
- Wang, T., Rudin, C., Doshi, F., Liu, Y., Klampfl, E., and MacNeille, P. A bayesian framework for learning rule set for interpretable classification. *Journal of Machine Learning Research*, 2017.
- Yang, H., Rudin, C., and Seltzer, M. Scalable bayesian rule lists. *ICML*, 2017.
- Yeh, I.-C. and Lien, C.-h. The comparisons of data mining techniques for the predictive accuracy of probability of default of credit card clients. *Expert Systems with Applications*, 36(2):2473–2480, 2009.
- Yi, K., Wu, J., Gan, C., Torralba, A., Kohli, P., and Tenenbaum, J. Neural-symbolic vqa: Disentangling reasoning from vision and language understanding. In *Advances in Neural Information Processing Systems*, pp. 1039–1050, 2018.
- Zeng, J., Ustun, B., and Rudin, C. Interpretable classification models for recidivism prediction. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 180(3):689–722, 2017.