
Sparse and Faithful Explanations Without Sparse Models

Yiyang Sun*
Duke University

Zhi Chen*
Duke University

Vittorio Orlandi
Duke University

Tong Wang
Yale University

Cynthia Rudin
Duke University

Abstract

Even if a model is not globally sparse, it is possible for decisions made from that model to be accurately and faithfully described by a small number of features. For instance, an application for a large loan might be denied to someone because they have no credit history, which overwhelms any evidence towards their creditworthiness. In this work, we introduce the *Sparse Explanation Value (SEV)*, a new way of measuring sparsity in machine learning models. In the loan denial example above, the SEV is 1 because only one factor is needed to explain why the loan was denied. SEV is a measure of *decision sparsity* rather than overall model sparsity, and we are able to show that many machine learning models – even if they are not sparse – actually have low decision sparsity, as measured by SEV. SEV is defined using movements over a hypercube, allowing SEV to be defined consistently over various model classes, with movement restrictions reflecting real-world constraints. We propose algorithms that reduce SEV without sacrificing accuracy, providing sparse and completely faithful explanations, even without globally sparse models.

1 Introduction

As machine learning is increasingly leveraged in business and societal contexts to make important decisions, there is an increasing need for insight into why these decisions were made. One classical important measure of interpretability is sparsity; famously, humans can handle only 7 ± 2 cognitive entities at once (Miller, 1956). Traditionally, sparsity is a property of the entire model, which we call global sparsity, for instance, the total number of terms in linear models,

*These authors contributed equally to this work

the total number of parameters in neural networks, or the number of leaves in decision trees (Murdoch et al., 2019). In this work, we argue that this notion of sparsity, while undeniably useful, may be overly restrictive. This is because *we do not need the model to be globally sparse in order to generate a sparse explanation for use in decision-making*. A far more relevant desideratum for users is that the explanation for each *individual prediction* is sparse. By this, we mean that the local explanation for a specific prediction depends only on a small number of features, which may differ across units. Since individuals care only about decisions made about themselves, we should consider aiming for sparse *explanations* – and not necessarily sparse *models*. That is, having an explanation for each prediction that does not necessarily contain all factors used by the model. Having sparse explanations, where we might faithfully explain a given prediction using only 1-3 features, allows for increased transparency in model-assisted decision-making processes, even if the model is not globally sparse.

To this end, we introduce a new metric for measuring the decision sparsity of classification models: the *Sparse Explanation Value (SEV)*. The SEV is defined by moving features from their values to the *population commons* (we call this our *reference*) or vice versa; the number of features that must be *aligned* (i.e., equate) to (or from) the reference in order to change the sign of the prediction is the SEV of the observation. SEV explains how many features change for predictions to switch from one class to another. SEV is easily explained and understood, and as we will see, the decision sparsity from SEV can be much sparser for each instance than the number of terms in a globally sparse model.

For example, consider a scenario where a loan applicant is denied a large loan due to their lack of credit history. This denial might occur even though the applicant has several positive factors, such as having a bank account, being employed, and not having a criminal record. In this case, the SEV is 1, which means that simply changing from having no credit history to having a credit history can flip the prediction, despite the other factors, and regardless of the total number of features in the model. If most of the feature values are 0 for a given individual (perhaps most people do not have complicated financial histories), or if most predictions depend only on a small set of factors (where important fac-

tors can differ across predictions), the global sparsity of the model is essentially irrelevant. Hence, SEV is an alternative way of measuring sparsity for these cases.

This work introduces the SEV and discusses its properties. SEV (as well as global sparsity) are most relevant for tabular datasets where each feature is meaningful. We show that decision sparsity, as measured by SEV, is naturally low in most common types of machine learning models already, without modification. **This means that despite the complexity of these models, their decisions are often based only on a few factors.** While SEV is useful as a post-hoc explanation method, it begs the question of whether models with better prediction-level sparsity can be created. In other words, if we care about sparse explanations, we should target them directly instead of obtaining globally sparse models and hoping they give the sparsest explanations. **Our second contribution is to create models that are optimized to yield sparse (i.e., low-SEV) explanations.** We do this by introducing SEV loss terms that can be used with a variety of models (linear, boosted trees, multi-layer perceptrons, etc.). Using real-world datasets, we show that SEV optimization is effective in improving decision sparsity.

2 Related Works

The concept of SEV revolves around finding models that are simple, in that the explanations for their predictions are sparse, while recognizing that different predictions can be simple in different ways (i.e., involving different features). In this way, it relates to (i) globally sparse models, (ii) local classification methods, which predict the outcomes of different units using local models, and (iii) black box explanation methods, which seek to explain predictions of complex models. We further comment on these below.

Global Model Sparsity. Loosely, globally sparse models (e.g., Rudin et al., 2022; Filzmoser et al., 2012; Tibshirani, 1996; Bach et al., 2012; Zhang et al., 2015) do not have many “components,” such as non-zero coefficients for additive models or leaves of a decision tree. Our focus in this paper is not overall model complexity, but rather the complexity of the individual predictions. These two notions of sparsity — global and local — are undoubtedly related. Even if every variable in a sparse model is used in every prediction, each explanation is necessarily sparse. However, non-sparse models can yield lower SEVs, as shown in Section 5.1.

Local Regression and Classification Methods. Such methods (e.g., Cleveland and Grosse, 1991; Bevilacqua and Marini, 2014; Dulac-Arnold et al., 2011) fit separate models for each data point, giving more weight to nearby points, and then possibly combining these models for final predictions. It is problematic to use numerous models instead of a single global model for applications like loan decisions since we do not want each person to be assigned their own credit scoring model; i.e., it is problematic that these models

are not globally consistent.

Instance-wise Explainability and Interpretability

There has been a large push to develop methods to explain the predictions of black boxes (e.g., Guidotti et al., 2018; Ribeiro et al., 2016, 2018; Lundberg and Lee, 2017; Baehrens et al., 2010). Those methods evaluate the contribution of features to the final outcome of the model. Some works are limited to particular domains and data types, such as images (e.g., Apicella et al., 2019, 2020; Boreiko et al., 2022), or text (e.g., Lei et al., 2016; Li et al., 2016; Treviso and Martins, 2020; Bastings et al., 2019; Yu et al., 2019, 2021). Chen et al. (2018) provides a subset features selection algorithm called *Learning to Explain (L2X)*, but with an entirely different premise — that sparsity values are fixed by the user and not adaptive. A more detailed comparison can be found in Appendix G.

Counterfactual explanations and *algorithmic recourse* study what could have happened if input to a model had been changed. Some methods are limited to text (Martens and Provost, 2014), or specific model classes, like linear models (Ustun et al., 2019), tree ensembles (Cui et al., 2015; Tolomei et al., 2017; Lucic et al., 2022), support vector machines (Gupta et al., 2019), or neural networks (Zhang et al., 2018). Most research in counterfactual explanations or algorithmic recourse focuses on minimizing ℓ_2 distance (Ross et al., 2021), ℓ_1 distance (Wachter et al., 2017; Zhang et al., 2018; McGrath et al., 2018; Russell, 2019; Joshi et al., 2019), cosine distance (Ramon et al., 2020), the overall cost of changing features by various amounts, a combination of distances (Laugel et al., 2018; Dhurandhar et al., 2018; Van Looveren and Klaise, 2021; Cheng et al., 2020) or model function outputs (Lash et al., 2017a,b) between the instance’s x value and the closest point across the decision boundary. An important distinction is that we use ℓ_0 distance between explanations to measure sparsity. The only paper we know of that uses ℓ_0 distance is that of Fernández-Loría et al. (2022), which is a special case of our SEV⁻; Fernández-Loría et al. (2022) does not discuss sparsity as an evaluation metric for explanations. As discussed, prediction sparsity is important: a classic counterfactual explanation might involve many changes to someone’s credit and lifestyle since the shortest path to crossing the decision boundary generally involves changes to many variables. Our explanations are sparse (e.g., “you have no credit history”), so they are easier to understand.

Additionally, there is ample evidence from human studies that the counterfactual explanations generated by most methods are not interpretable to humans because they aim only to cross the decision boundary, where the decision made by the model is not obvious to a human. Extensive experiments of Delaney et al. (2023) indicate that minimal changes to an observation that alter the predicted class are not natural for humans because the decision boundary is crossed in a

low-density part of the space where the observation is not clearly one class or the other. Our best competitor (counterfactual explanation method DiCE) has precisely this problem, shown in Section 5.4. By framing SEV with respect to a population attribute reference (i.e., typical values of features), it provides explanations that humans are more likely to understand. Importantly, Delaney et al. (2023) show that counterfactual explanations created by humans are not close to the decision boundary; people create counterfactuals that appear to be class-specific prototypes.

3 Sparse Explanation Values

Let us define SEV. We are in the setting of classification, where we have observations (units) with features $\mathbf{x} \in \mathcal{X} \subseteq \mathbb{R}^p$ and an outcome of interest. Here, we consider binary outcomes $\{0, 1\}$, though the extension to multi-class outcomes is straightforward. We predict the outcome using a classifier $f : \mathcal{X} \rightarrow \{0, 1\}$ that maps the input \mathbf{x} to a binary prediction \hat{y} . We consider explanations for only one class; to consider explanations for the other class, one would use a completely symmetric procedure (defining a reference for the other class). Without loss of generality, we assume class 1 is the class of interest (positive) and we focus on explaining the predictions of a unit (which we call a *query*, denoted $\mathbf{x}^{\text{query}}$) that is classified as positive ($f(\mathbf{x}^{\text{query}}) = 1$). As shown by human studies (Delaney et al., 2023), it is natural to compare a query with a *reference*; for example, a unit representative of the population as a whole or of negatively predicted (predicted to be class 0) units. Therefore, we define explanation sparsity with respect to the reference $\tilde{\mathbf{x}} \in \mathbb{R}^p$, where \tilde{x}_j is the reference value of feature j , e.g., the population mean or median of that feature in the dataset. We will consider aligning the observed values of the query $\mathbf{x}_j^{\text{query}}$ with the corresponding reference feature values \tilde{x}_j to see whether predictions change. Thinking of these adjustments as binary moves, it is convenient to represent all 2^p possible different feature sets as vertices of a Boolean hypercube. The hypercube is defined below:

Definition 1 (SEV hypercube) *An SEV hypercube $\mathcal{L}_{f, \text{query}}$ for a model f , reference $\tilde{\mathbf{x}}$, and query $\mathbf{x}^{\text{query}}$, such that $f(\mathbf{x}^{\text{query}}) = 1$, is a graph with 2^p vertices, where each vertex v is represented by a p -dimensional Boolean vector $\mathbf{b}_v \in \{0, 1\}^p$. Vertices u and v are adjacent when their Boolean vectors differ by one bit: $\|\mathbf{b}_u - \mathbf{b}_v\|_0 = 1$. 0's in \mathbf{b}_v indicate the reference value $\tilde{\mathbf{x}}$ of the corresponding feature, while 1's indicate the true feature value of the query. Thus, the actual feature vector represented by the vertex v is $\mathbf{x}_v := \mathbf{b}_v \odot \mathbf{x}^{\text{query}} + (\mathbf{1} - \mathbf{b}_v) \odot \tilde{\mathbf{x}}$, where \odot is the Hadamard product. The score of vertex v is $f(\mathbf{x}_v)$, also denoted as $\mathcal{L}_{f, \text{query}}(\mathbf{b}_v)$.*

The SEV hypercube definition can also be extended from a hypercube to a Boolean lattice as they have the same geometric structure. Based on this definition of SEV hypercubes,

we define two variants of the Sparse Explanation Value: one is the smallest number of features that, starting from the reference (corresponding to the origin of the hypercube $\mathbf{b} = \mathbf{0}$), need to be aligned with the query's feature values to make the prediction flip from negative to positive (called SEV^+). The other (called SEV^-) is the counterfactual case: it is the minimal number of features, starting from the query (corresponding to $\mathbf{b} = \mathbf{1}$) that need to be aligned with the reference's feature values to make the prediction flip from positive to negative. Formally:

Definition 2 (SEV Plus, denoted SEV^+) *For a query predicted as positive ($f(\mathbf{x}^{\text{query}}) = 1$), the Sparse Explanation Value Plus (SEV^+) is the length of the shortest path from the reference vertex to any positively predicted vertex, i.e.,*

$$\min_{\mathbf{b} \in \{0, 1\}^p} \|\mathbf{b}\|_0 \quad \text{s.t.} \quad \mathcal{L}_{f, \text{query}}(\mathbf{b}) = 1. \quad (1)$$

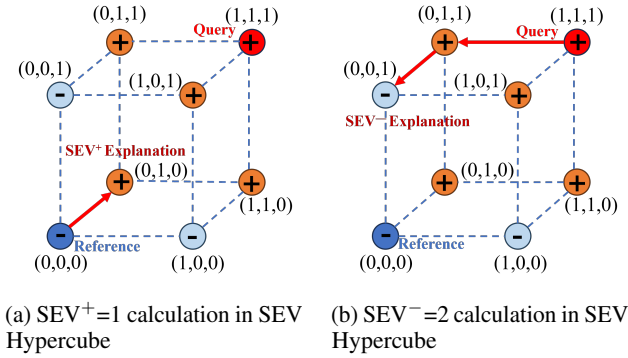
Definition 3 (SEV Minus, denoted SEV^-) *For a query predicted as positive ($f(\mathbf{x}^{\text{query}}) = 1$), the Sparse Explanation Value Minus (SEV^-) is the length of the shortest path from the query to any negatively predicted vertex,*

$$\min_{\mathbf{b} \in \{0, 1\}^p} \|\mathbf{1} - \mathbf{b}\|_0 \quad \text{s.t.} \quad \mathcal{L}_{f, \text{query}}(\mathbf{b}) = 0. \quad (2)$$

SEV^- is similar to a counterfactual explanation in which we use the length of the path on the SEV hypercube to measure the counterfactual change (see also Section 5.4). Figure 1a, 1b and Table 1c, 1d show examples of how $\text{SEV}^+ = 1$ and $\text{SEV}^- = 2$ are calculated in a loan credit evaluation decision involving three features: housing = renting, loan amount $\geq 10\text{k}$, and education level = high school level, where we need an SEV hypercube with $2^3 = 8$ vertices. The red vertex in Figure 1a is encoded as $(1, 1, 1)$, which represents the query; the dark blue vertex is encoded as $(0, 0, 0)$, representing the negatively-predicted reference value; the orange vertices are positively-predicted vertices; and the light blue ones are negatively predicted vertices. To compute SEV^+ , we start from the reference $(0, 0, 0)$ and find the shortest path to any positive (red/orange) nodes. Here, the two shortest paths end at $(0, 1, 0)$, which means that if we change the loan amount from 5-10k to more than 10k, then the model will evaluate it as high risk. Since we are changing only one feature (loan amount), SEV^+ equals 1, and the feature vector that $(0, 1, 0)$ represents is called an *SEV⁺ explanation*. For SEV^- , we start from the query $(1, 1, 1)$, and find the shortest path to any blue node. Here the shortest paths have length 2 (e.g., $(1, 1, 1) \rightarrow (0, 1, 1) \rightarrow (0, 0, 1)$), which means that if we changed the housing status of the renter from renting to managing and the loan amount from greater than 10k to 5-10k, then the credit risk will be evaluated as low instead of high. Moreover, there also exists other shortest paths like $(1, 1, 1) \rightarrow (1, 1, 0) \rightarrow (1, 0, 0)$ which means we can also align the loan amount and the education level with their reference values to obtain a low-risk prediction. This

alternative path to a negative vertex shows that we are able to provide multiple sparse explanations through the SEV hypercube. The feature vector that $(1, 0, 0)$ represents is called an SEV^- explanation.

Figures 1e and 1f show how SEV is distributed in 2D cases. The blue line is the decision boundary. The orange horizontal and vertical lines are where the decision boundary intersects one of the reference’s two feature values. If the query’s feature values are in the yellow area of the SEV^+ figure (Figure 1e), it means that for both features, the reference feature value must be changed to the query’s feature values in order to be predicted positive. In the yellow region of Figure 1f (SEV^- case), for both features, the query’s feature values must be aligned with the reference values in order to be predicted as negative.



(c) $SEV^+=1$ example

	Location	Housing	Loan	Education	Credit Risk
Reference	(0,0,0)	Mortgage	5-10k	Bachelor	Low
SEV^+ Explanation	(0,1,0)	Mortgage	>\$10k	Bachelor	High
Query	(1,1,1)	Rent	>\$10k	High School	High

(d) $SEV^-=2$ example

	Location	Housing	Loan	Education	Credit Risk
Query	(1,1,1)	Rent	>\$10k	High School	High
	(0,1,1)	Mortgage	>\$10k	High School	High
SEV^- Explanation	(0,0,1)	Mortgage	5-10k	High School	Low
Reference	(0,0,0)	Mortgage	5-10k	Bachelor	Low

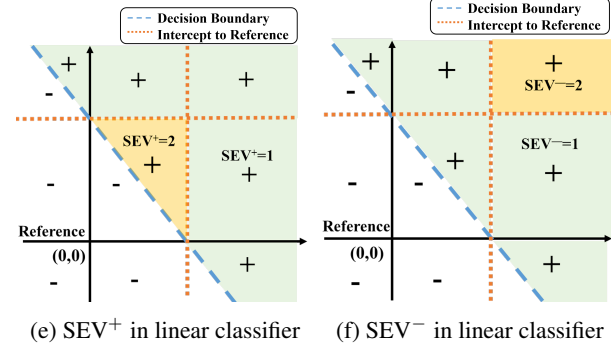


Figure 1: Visual illustrations of SEV definitions

In real-world applications, some features might be immutable or should not contribute to explanations for legal or ethical reasons. Suppose a bank uses a machine learning

model for making credit loan decisions. The bank may want to ensure that the explanation for a denied loan does not involve age since no one can adjust their age. We can accommodate this via a straightforward extension of SEV^- , which we call the Restricted Sparse Explanation Value (SEV^{\otimes}):

Definition 4 (Restricted SEV, denoted SEV^{\otimes}) For a query, predicted as positive ($f(\mathbf{x}^{query}) = 1$), and a set of restricted features $\mathcal{S}_r \subseteq \{1, 2, \dots, p\}$, the Restricted Sparse Explanation Value (SEV^{\otimes}) is the length of the shortest pathway from the sample’s feature values to a negatively predicted vertex, considering only non-restricted features, i.e.,

$$\min_{\mathbf{b} \in \{0,1\}^p} \|\mathbf{1} - \mathbf{b}\|_0 \quad (3)$$

$$s.t. \quad \mathcal{L}_{f,query}(\mathbf{b}) = 0, \quad \text{where } b_j = 1 \quad \forall j \in \mathcal{S}_r.$$

SEV^+ , SEV^- , and SEV^{\otimes} can all be computed by a breadth-first search along the SEV hypercube. Since we care about sparse and faithful explanations (i.e., no more than 5-7 features) we can set up search depth limits, which also have the effect of lowering the computational time. SEVs are usually very low in practice, and thus can generally be computed in milliseconds. Detailed results are shown in Section 5.1. Appendix D also shows that SEV requires equal or less time to generate explanations than other methods.

4 Optimizing Sparse Explanation Values

In this section, we describe how to obtain classifiers that optimize the average SEV^- or SEV^+ . As we mentioned, calculating SEVs is theoretically a combinatorial problem. Thus directly optimizing it is intractable. Therefore, we propose some easy-to-optimize surrogate objectives that obtain classifiers with very low SEVs in practice.

4.1 Volume-based SEV^+ Optimization for Linear Models

Consider a linear classifier, $f(\mathbf{x}) := \mathbf{1}[(\beta_0 + \sum_{j=1}^p \beta_j x_j) > 0]$. In this case, there exists an analytical surrogate objective for optimizing SEV^+ . Given this classifier, the positively classified region of the space $\{\mathbf{x} \in \mathbb{R}^p | f(\mathbf{x}) = 1\}$ can be partitioned into regions with $SEV^+ = 1$, $SEV^+ = 2$, ..., $SEV^+ = p$. Approximating the data to be uniformly distributed, we can minimize the average SEV^+ by minimizing the volume of the regions with $SEV^+ \geq 2$.

Theorem 5 Consider a linear classifier, $f(\mathbf{x}) := \mathbf{1}[(\beta_0 + \sum_{j=1}^p \beta_j x_j) > 0]$, where $\forall j, \beta_j \neq 0$, and for reference $\tilde{\mathbf{x}}$, we have $f(\tilde{\mathbf{x}}) = 0$ (i.e., reference predicted as negative). Let $g^{\text{ref}}(\beta) = \beta_0 + \sum_{j=1}^p \beta_j \tilde{x}_j$ be the raw score of the classifier f at reference $\tilde{\mathbf{x}}$. For all $k \in \{2, 3, \dots, p\}$, the volume of

region V_k in the input space \mathbb{R}^p with $SEV^+ = k$ is:

$$V_k = c_k \cdot \prod_{j=1}^p \left| \frac{g^{\text{ref}}(\beta)}{\beta_j} \right|, \quad (4)$$

where c_k is a finite constant unrelated to the β 's.

See the proof in Appendix F.

If some β_j 's are zero, we can either calculate volume on the subspace where the β_j 's are non-zero, or add an ϵ to the terms in the denominator. According to Theorem 5, the volume of the region with $SEV^+ \geq 2$ is proportional to the volume of a hyperrectangle, given by $\prod_{j=1}^p \left| \frac{g^{\text{ref}}(\beta)}{\beta_j} \right|$. Therefore, to minimize the expected SEV in the region, we minimize the volume of this hyperrectangle. For better optimization stability, we optimize the logarithm of the volume, i.e.,

$$\min_{\beta_0, \beta_1, \dots, \beta_p} \frac{1}{p} \sum_{j=1}^p \log \left(\min \left(\left| \frac{g^{\text{ref}}(\beta)}{\beta_j} \right|, \epsilon \right) \right) \quad (5)$$

where ϵ is used to ensure numerical stability.

This volume-based SEV loss, which we call **Vol-Opt**, is a simple and theoretically justified method that is effective for optimizing the SEV^+ of linear models. However, it is limited to linear models and cannot be applied to more complex models. In the next subsection, we introduce individual-based SEV loss, which can be applied to other types of models as well.

4.2 Individual-based SEV Optimization

The volume-based SEV loss above approximates queries as arising uniformly over \mathcal{X} , which might not be true. Making the less restrictive assumption that the distribution of a training dataset is close to that of test data, we can instead focus on minimizing the average training SEV. The difficulty here lies in the combinatorial number of steps required to calculate SEV. However, we have found that most trained classifiers yield low average SEVs (roughly < 3) on real datasets. With the majority of training data having low SEV, we can design the optimization objective to penalize queries with SEV^+ not equal to 1. On the other hand, we maximize the fraction of points where SEV^+ is exactly 1. Given a class \mathcal{F} of classification models that output the probability of belonging to the positive class, the loss term for minimizing SEV^+ on a training set is:

$$\begin{aligned} \ell_{SEV_All_Opt^+}(g) := & \\ & - \frac{1}{n^+} \sum_{i=1}^{n^+} \min \left(\max_{j=1, \dots, p} g(e_j \odot \mathbf{x}_i + (\mathbf{1} - e_j) \odot \tilde{\mathbf{x}}), \mathcal{T} \right) \end{aligned} \quad (6)$$

where \mathbf{x}_i is the i -th query in the training data, n^+ represents the number of positively predicted queries in the dataset, and

e_j denotes the vector with a 1 in the j^{th} coordinate and 0's elsewhere, and $g : \mathbb{R}^p \rightarrow [0, 1]$ is the continuous estimator associated with classifier f (e.g., in logistic regression, g is a sigmoid function), and \mathcal{T} is the threshold value for classification, which is usually 0.5 for balanced data. Here, we consider all vertices 1 step away from the reference and see whether at least one of them is predicted as positive (i.e., $SEV^+=1$). The min operation taken with the threshold \mathcal{T} is used to fix the value of the objective when the prediction is positive since we do not care where on the positive side of the decision boundary a prediction is, as long as it is positive. Similarly, we propose the following loss term for minimizing SEV^- :

$$\begin{aligned} \ell_{SEV_All_Opt^-}(g) := & \\ & \frac{1}{n^+} \sum_{i=1}^{n^+} \max \left(\min_{j=1, \dots, p} g((\mathbf{1} - e_j) \odot \mathbf{x}_i + e_j \odot \tilde{\mathbf{x}}), \mathcal{T} \right) \end{aligned} \quad (7)$$

We call these individual-based SEV losses **All-Opt⁺** and **All-Opt⁻**. The All-Opt⁻ loss can also be used to target SEV^{\otimes} by replacing the minimization over $j = 1, \dots, p$ with a minimization over $j \in \mathcal{S}_r$, which we denote as **All-Opt^{\otimes}**. The experiments section will show that these losses are not only effective in shrinking the average SEV but often result in attaining the minimum possible SEV value of 1 for most or all queries.

One thing to note is that both All-Opt and Vol-Opt assume that the reference is predicted negative; that is, $\mathcal{L}_{f, \text{query}}(\mathbf{0}) = 0$. Thus, if the reference is predicted as positive, then optimization for SEV^+ leads to a trivial solution where $SEV^+ = 0$ for all samples. Similarly SEV^- would not find a shortest path (or any path for that matter) from the observation to the negative class, since the reference is positive. Thus, to ensure the reference is predicted to be negative, we add a term that penalizes the reference receiving a positive prediction:

$$\ell_{\text{Pos_base}} := \max(g(\tilde{\mathbf{x}}), \mathcal{T} - \theta) \quad (8)$$

where $\theta > 0$ is a margin parameter, for example, $\theta = 0.05$. This term is $(\mathcal{T} - \theta)$ as long as the reference is predicted negative. As soon as it exceeds that amount, it is penalized (increasing linearly in $g(\tilde{\mathbf{x}})$).

Combining the loss terms above, we optimize a linear combination of them,

$$\min_{f \in \mathcal{F}} \ell_{\text{BCE}} + C_1 \ell_{SEV} + C_2 \ell_{\text{Pos_base}} \quad (9)$$

where ℓ_{BCE} is the Binary Cross Entropy Loss to control the accuracy of the training model. ℓ_{SEV} can be any of the SEV-based terms we introduced, and $\ell_{\text{Pos_base}}$ is for ensuring the reference is predicted as negative. C_1 and C_2 are the strengths of the terms and can be tuned through cross-validation. Minimizing (9) gives a low-SEV model.

5 Experiments

We present experiments on six real-world datasets: (i) the UCI Adult Income dataset for predicting income levels (Becker and Kohavi, 1996; Dua and Graff, 2017), (ii) the FICO Home Equity Line of Credit Dataset for assessing credit risk, used for the Explainable Machine Learning Challenge (FICO, 2018), (iii) the UCI German Credit dataset for determining creditworthiness (Hofmann, 1994; Dua and Graff, 2017), (iv) the MIMIC-III dataset for predicting patient outcomes in intensive care units (Johnson et al., 2016a,b), (v) the COMPAS dataset (Jeff Larson and Angwin, 2016; Wang et al., 2022) for predicting recidivism, and (vi) the Diabetes dataset (Strack et al., 2014; Dua and Graff, 2017) for predicting if patients will be re-admitted within two years. Additional details on data and preprocessing are in Appendix A.

We train four reference binary classifiers: (i, ii) Logistic Regression classifiers with ℓ_1 (**L1 LR**) and ℓ_2 (**L2 LR**) penalties, (iii) a Gradient Boosting Decision Tree classifier (**GBDT**), and (iv) a 2-layer Multi-Layer Perceptron (**MLP**). In addition, we train “optimized” variants of these models, in which the SEV penalties described in the previous sections are implemented. Details on training the methods are in Appendix B. Below, we report the average test SEV; i.e. averaged across queries in the test set.

5.1 Most models already have low SEVs.

We first compute SEV^+ and SEV^- for our four reference classifiers. We emphasize that these models are not optimized for SEV, they were created with standard machine-learning algorithms. All algorithms were run for 10 different train-test splits and the mean and standard deviation across these splits of the average test accuracy, AUC, SEV^+ and SEV^- on the German Credit data are shown in Table 1. Results for the other datasets are listed in Appendix C.


Table 1: SEV for reference classifiers in German Credit

Model	L1 LR	L2 LR	MLP	GBDT
Acc	0.74 ± 0.02	0.73 ± 0.03	0.76 ± 0.03	0.75 ± 0.03
AUC	0.77 ± 0.03	0.76 ± 0.03	0.77 ± 0.08	0.78 ± 0.02
SEV^+	1.00 ± 0.00	1.00 ± 0.00	1.04 ± 0.08	1.10 ± 0.24
SEV^-	1.29 ± 0.14	1.74 ± 0.16	1.58 ± 0.22	1.58 ± 0.22

Even though the German Credit models for L2LR, MLP and GBDT used up to all 20 features (they are not sparse), we observe from Table 1 that most of the queries need only to align one reference feature value with the query value to change to a positive prediction (SEV^+), and we need to align fewer than 2 feature values from the query to the reference for a negative prediction (SEV^-). Thus, *even without optimization for SEV at all, most models already have sparse explanations*. However, there is still room for improvement, as we now show.

5.2 SEV can be reduced with no loss in performance.

While existing models already have low SEVs on real datasets, our SEV optimization algorithm proposed in the previous section can further decrease SEVs – creating accurate models with extreme decision sparsity. To illustrate, Figure 2 shows the impact on SEV^+ of All-Opt and Vol-Opt on a reference linear classifier trained on the Adult dataset. Figure 2a shows the SEV distribution across queries; while L1 LR and L2 LR have reasonably low SEVs (per the discussion in the previous subsection), both of the proposed optimization algorithms are able to find models for which the SEV^+ of *all* queries is 1 (the minimum possible SEV). Table 7 in the Appendix shows that this incurs negligible performance loss. Figure 2b shows how the optimization affects the SEV^+ of each query. Its biggest impact on the mean is from reducing the explanations of many queries from involving 2 to 5 terms to a single one, with no query’s explanation becoming more complicated. Additionally, the predicted class distribution is comparable, and the impact of the optimization is not simply to predict a greater number of negative labels; it actually targets the explanation sparsity.



(a) SEV^+ counts across methods (b) SEV^+ Optimization Result

Figure 2: SEV^+ performance for linear classifiers in Adult

All-Opt⁺ and All-Opt⁻ can also be applied to nonlinear models. To illustrate, we apply them to baseline **MLP** and **GBDT** classifiers, again using the Adult data. For **GBDT**, we use our methods to fine-tune the weights for each tree generated by the original method. Both are optimized by Stochastic Gradient Descent. Figure 3 shows that across both model classes, and for both SEV^- and SEV^+ , All-Opt⁻ and All-Opt⁺ are able to reduce the SEV of all queries to 1, as in the linear case. Appendix C has other performance metrics, as well as results for other datasets, showing that we can consistently reduce SEV with negligible loss in accuracy, across various model classes and datasets.

Moreover, to evaluate the time consumption of our methods, we have tried different dataset sizes. The detailed information for each dataset is shown in Appendix A. To better compare the training time consumption, we have used **MLP** as the baseline model and compared its running time with and without All-Opt for different datasets. All the models were run 10 times and we report the mean and standard

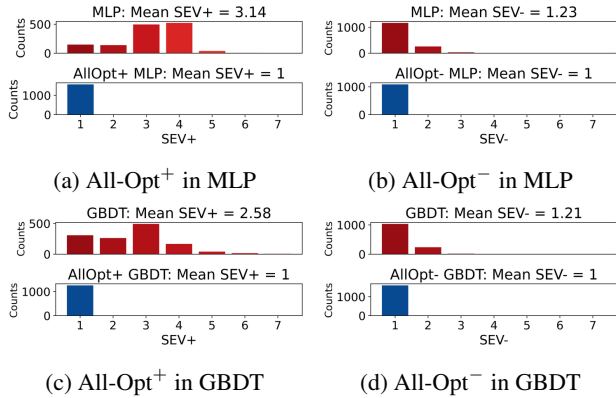


Figure 3: Performance of All-Opt⁺ and All-Opt⁻ in Adult

deviation in Table 2.

Dataset	Baseline	All-Opt ⁺	All-Opt ⁻
COMPAS	10.5 ± 0.69	26.9 ± 1.06	21.7 ± 0.03
Adult	61.1 ± 0.15	164 ± 1.56	136 ± 1.22
MIMIC III	78.2 ± 1.09	207 ± 2.71	167 ± 2.42
German Credit	1.81 ± 0.01	5.01 ± 0.01	4.08 ± 0.04
FICO	16.9 ± 0.78	47.6 ± 0.34	37.4 ± 1.30
Diabetes	188 ± 3.08	606 ± 3.65	517 ± 3.31

Table 2: The training time (seconds) comparison with and without All-Opt in MLP

Based on the run time results shown above, we can observe that even when optimizing a complex black box model, the increase in time consumption associated with optimizing SEV is quite low.

5.3 It is unnecessary to have global sparsity to have sparse explanations.

In this section, we discuss whether global sparsity is needed for sparse local explanations. We use the proportion of zero coefficients for linear models as a global sparsity measure. We will use the Adult income dataset as an example to see how the model distributes coefficients.

Table 3: SEV and zero-coefficient % of LR in Adult

Methods	L1 LR	L2 LR	All-Opt ⁺ LR	Vol-Opt LR
Mean SEV ⁺	1.95±0.01	2.30±0.08	1.00±0.00	1.00±0.00
% of Zero Coefficients	87.9±0.00	0.09±0.00	0.00±0.00	0.00 ± 0.00

The L1 LR, L2 LR, All-Opt⁺ LR, and Vol-Opt LR methods are all compared in Table 3. L1 LR has already attained low SEVs, possibly because its models are globally sparse. Nevertheless, even without global sparsity, L2 penalized linear regression also generates a low SEV model. This suggests that if our objective is to construct models with low

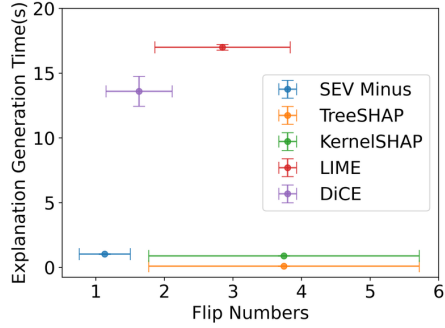
SEVs, we should not just focus on producing globally sparse models, as this may be unnecessarily restrictive. Moreover, after SEV⁺ optimization, neither the All-Opt⁺ LR nor the Vol-Opt LR optimization approaches produce sparse models and yet have perfect SEV’s. Thus, it is not necessary to have global sparsity to achieve sparse explanations.

5.4 Explaining a fixed model: comparing SEV to other local explanation methods


Here, we use GBDT to generate a model on the COMPAS dataset, and compare SEV⁻ with various post hoc explanation methods, specifically: SHAP-C (using treeSHAP and kernelSHAP) (Lundberg and Lee, 2017; Ramon et al., 2020), LIME-C (Ribeiro et al., 2016; Ramon et al., 2020), and Diverse Counterfactual Explanations (DiCE) (Mothilal et al., 2020). LIME and SHAP are post-hoc explainers that provide local feature importance scores to the prediction. The idea of SHAP-C and LIME-C is to order the features by importance, and based on that ordering, progressively (from the most important feature to less important ones) align the queries’ features with the reference values. We then calculate the number of alignments needed for the point to change from being positively predicted to negatively predicted (shown as Flip Number in Figure 4). We also compare our methods to DiCE, counting the number of features it modifies to generate a counterfactual explanation.

Figure 4 shows that, on average, LIME-C and SHAP-C flip about half of the features in the model (which uses 7 features in total) and do not induce sparse explanations. This conclusion aligns with Fernández-Loría et al. (2022)’s observation that high locally important features may not always influence model decisions.

Even though DiCE applies backward feature selection after generating explanations to control sparsity (Mothilal et al., 2020), Figure 4 shows that SEV⁻ provides sparser explanations than DiCE. Moreover, DiCE provides less interpretable explanations than SEV⁻. Recall that DiCE uses minimal edit distance to generate counterfactual explanations, whereas we find a minimal number of flips between original values and reference values. As we discussed in Related Works, the edit-distance-based counterfactual explanations may not be as interpretable to humans (Delaney et al., 2023) because the counterfactual often ends up in a low-density region where one cannot determine the class with certainty. The counterfactual explanations generated by DiCE can even lead away from the general direction of the other class. For instance, we found that on a criminal recidivism dataset, DiCE stated that a high-risk criminal offender who commits *more* crimes would be classified as *lower* risk, which is paradoxical. This unreasonable explanation arises because DiCE can reference *any* negatively-predicted point, including outliers. In contrast, the reference point for SEV is the center of the negative population, thus the direction of



(a) Flip Numbers and Mean Explanation Time Consumption



(b) Flip Number Distributions

Figure 4: Local Explanation Methods’ Performance Comparison in COMPAS

the SEV is generally meaningful: towards the negatives. For example, the reference value used in our approach for the number of prior misdemeanors is 0, which is more natural.

Table 4 shows counterfactual explanations for two different queries. For the first, DiCE suggests that the number of juvenile crimes should go from 1 to 5 to change the prediction from ‘will recidivate’ to ‘won’t recidivate.’ This is nonsensical, as the number of crimes should be positively correlated with the propensity to recidivate. In contrast, SEV⁻ suggests that the number of prior crimes should go from 2 to 1, which makes sense. Similarly, for the second query, DiCE suggests that the number of juvenile crimes should go from 3 to 7 and the number of juvenile misdemeanors from 2 to 11 in order to change the prediction from ‘will recidivate’ to ‘won’t recidivate.’ Again, this explanation is nonsensical. In contrast, SEV⁻ provides two possible explanations: (i) that the number of prior crimes should go from 3 to 1 or (ii) that the number of juvenile misdemeanors should go from 2 to 0. Again, this makes more sense. Given that DiCE uses a random initialization to find a counterfactual explanation whereas SEV⁻ provides the *minimum* number of feature alignments to cross the decision boundary, we would expect

SEV⁻ to be more stable and to generate useful explanations.


In addition, SEV has shown competitive speed in generating local explanations. Figure 4a shows that DiCE takes more time to generate stable sparse counterfactual explanations than SEV⁻. Table 5 shows the detailed time required to generate a local explanation for our method and baselines SHAP, LIME, and DiCE, for several datasets. We use **GBDT** as a baseline model and evaluate the median explanation runtime on queries in the test set. The table shows that SEV⁻ has a reasonable running time for black-box models on most datasets. We can observe that the explanation time is better than DiCE in all the datasets since it only flips from feature values to the reference. The reason why the FICO Dataset takes a long time is that its average SEV⁻ is about 4, which means that the search in the SEV hypercube starting from the query is too deep. However, if the time consumption for generating explanations is too long, it means that this model cannot be sparsely explained, but we can use All-Opt or Vol-Opt methods to optimize our model. Importantly, other types of explanations (SHAP, LIME, DiCE) do not have optimization methods such as All-Opt or Vol-Opt. The detailed comparison between KernelSHAP and SEV⁻ is shown in Appendix D.

5.5 Restricted SEV in practice

In real applications, some features are not changeable, such as race, gender, or age. To evaluate SEV[®] and our optimization algorithm under restricted feature sets, we apply our loss function to COMPAS and set `gender` and `age` as restricted features when training a linear classifier.

(a) SEV[®] Difference before and after optimization

	Before Optimization	After Optimization
Test Acc	0.68 ± 0.01	0.65 ± 0.01
Test AUC	0.73 ± 0.01	0.72 ± 0.01
Mean SEV [®]	3.20 ± 0.26	1.28 ± 0.38
Unexplained (%)	34.40 ± 4.40	3.35 ± 6.20



(b) Feature Prevalence in Explanations

Figure 5: The performance of All-Opt[®] on COMPAS data

In this restricted setting, there may exist queries whose predictions do not flip even if all nonrestricted features are

Explanations	Sex	Age	Juv fel counts	Juv misd count	Juvenile crimes	Priors counts	Current Charge Degree	Predicted Label
Query 1	Male	23	0	1	1	2	Not Felony	1
DiCE	-	-	-	-	5	-	-	0
SEV ⁻	-	-	-	-	-	1	-	0
Query 2	Male	19	1	2	3	3	Felony	1
DiCE	-	-	-	11	7	-	-	0
SEV ⁻	-	-	-	0	-	-	-	0
	-	-	-	-	-	1	-	0

Table 4: DiCE and SEV⁻ explanations for two different queries in COMPAS

Dataset	TreeSHAP	KernelSHAP	LIME	DiCE	SEV ⁻
Adult	1.08 ± 0.00	4.82 ± 0.01	20.27 ± 0.22	43.41 ± 0.16	1.39 ± 0.00
COMPAS	0.50 ± 0.01	0.22 ± 0.00	2.68 ± 0.10	20.86 ± 0.39	0.19 ± 0.02
MIMIC	1.13 ± 0.01	17.30 ± 0.10	6.10 ± 0.05	Crashed	0.45 ± 0.03
German Credit	0.53 ± 0.00	9.70 ± 0.17	7.11 ± 0.07	26.70 ± 0.91	1.37 ± 0.05
FICO	0.56 ± 0.02	8.25 ± 0.27	6.05 ± 0.14	27.82 ± 0.05	15.75 ± 1.01
Diabetes	0.55 ± 0.01	9.42 ± 0.10	22.49 ± 0.26	34.66 ± 1.33	2.24 ± 0.01

Table 5: The median explanation runtime (in 10⁻² s) for query units predicted as positive. The values in the table differ from those in Figure 4a since Figure 4a reports the mean runtime instead of median.

flipped. We call these queries “unexplainable.” Figure 5a shows the proportion of unexplainable queries in the restricted setting and their relevant SEV[®] before and after restriction and optimization of the All-Opt[®] objective. Figure 5b shows the proportion of queries using each feature for sparse explanations before (Red) and after restriction and optimization (Blue). After introducing restricted features, over one-third of the queries are unable to be explained by the baseline model. However, after All-Opt[®] is applied, the percentage of unexplainable queries decreases to 3% without much performance loss. Moreover, Figure 5b shows that after applying All-Opt[®] with `age` and `sex=female` as restricted features, `prior_count` becomes the most important factor for explanations for almost all queries, which is more actionable. In real cases, we might want to replace this feature with a dynamic feature, e.g., the number of priors over the last 5 years, to allow individuals to reduce that feature over time through their actions. In that case, one’s explanation would often be simply: *you are predicted as high risk because you committed a crime in the last 5 years*. More results on the performance of All-Opt[®] are shown in Appendix E. These results demonstrate that All-Opt[®] works well in restricted cases.

6 Discussion and Limitations

In this work, we introduce the Sparse Explanation Value (SEV) as a new way to measure and optimize interpretability in machine learning models. SEV measures decision sparsity, which focuses on how simply predictions can be explained, instead of global sparsity. Because users care about explanations for their individual predictions (rather than everyone else’s predictions), SEV better connects mod-

els to their real-world applications. Crucially, we have found (i) that many current models naturally have low SEV, (ii) that global sparsity is not necessary for decision sparsity, and (iii) that SEV can be directly optimized without sacrificing accuracy. The use of interpretable reference values and the development of optimization algorithms (Vol-Opt and All-Opt) further highlight SEV’s potential in applications.

There are several limitations to our approach. One is that causal relationships may exist among features, invalidating certain transitions across the SEV hypercube. This can be addressed by adapting the definition of vertices to change features in tandem, though it requires knowledge of the causal graph. Another practical consideration is choosing an appropriate reference and SEV variant. However, the flexibility in reference and variant choice actually helps tailor analyses to specific problems. For example, individuals or businesses might consider SEV⁺ to understand how they can make their profiles more appealing to lenders. In contrast, examining SEV⁻ can help lenders understand the fragility of their approvals and can provide useful information about important variables. Lastly, SEV[®], which excludes unchangeable characteristics, offers more actionable explanations to users. However, restricting explanations may obscure biases inherent in the model since the excluded features cannot contribute to explanations.

Ultimately, SEV shifts the burden for interpretability in machine learning models from prioritizing sparse models to prioritizing sparse decisions. SEV is distinct from existing black-box explanation methods, can be computed for pre-trained models, and can be optimized via efficient optimization algorithms. These qualities make it a candidate for real-world modeling pipelines.

References

- George Miller. The magical number seven, plus or minus two: Some limits on our capacity for processing information. *The Psychological Review*, 63:81–97, 1956.
- W James Murdoch, Chandan Singh, Karl Kumbier, Reza Abbasi-Asl, and Bin Yu. Definitions, methods, and applications in interpretable machine learning. *Proceedings of the National Academy of Sciences*, 116(44):22071–22080, 2019.
- Cynthia Rudin, Chaofan Chen, Zhi Chen, Haiyang Huang, Lesia Semenova, and Chudi Zhong. Interpretable machine learning: Fundamental principles and 10 grand challenges. *Statistics Surveys*, 16(none):1 – 85, 2022. doi: 10.1214/21-SS133. URL <https://doi.org/10.1214/21-SS133>.
- Peter Filzmoser, Moritz Gschwandtner, and Valentin Todorov. Review of sparse methods in regression and classification with application to chemometrics. *Journal of Chemometrics*, 26(3-4), 2012.
- Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society (Series B)*, 1996.
- Francis Bach, Rodolphe Jenatton, Julien Mairal, Guillaume Obozinski, et al. Optimization with sparsity-inducing penalties. *Foundations and Trends in Machine Learning*, 4(1):1–106, 2012.
- Zheng Zhang, Yong Xu, Jian Yang, Xuelong Li, and David Zhang. A survey of sparse representation: algorithms and applications. *IEEE Access*, 3:490–530, 2015.
- William S Cleveland and E. Grosse. Computational methods for local regression. *Statistics and Computing*, 1991.
- Marta Bevilacqua and Federico Marini. Local classification: Locally weighted–partial least squares–discriminant analysis (LW–PLS–DA). *Analytica Chimica Acta*, 2014.
- Gabriel Dulac-Arnold, Ludovic Denoyer, Philippe Preux, and Patrick Gallinari. Datum-wise classification: A sequential approach to sparsity. In *Machine Learning and Knowledge Discovery in Databases - European Conference, ECML PKDD 2011, Athens, Greece, September 5-9, 2011. Proceedings, Part I*, volume 6911 of *Lecture Notes in Computer Science*, pages 375–390. Springer, 2011.
- Riccardo Guidotti, Anna Monreale, Salvatore Ruggieri, Franco Turini, Fosca Giannotti, and Pedreschi Dino. A survey of methods for explaining black box models. *ACM Computing Surveys*, 51(5), 2018.
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. “Why should I trust you?” explaining the predictions of any classifier. In *Proceedings of the 22nd International Conference on Knowledge Discovery and Data Mining*, 2016.
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. Anchors: High-precision model-agnostic explanations. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence*, 2018.
- Scott M Lundberg and Su-In Lee. A Unified Approach to Interpreting Model Predictions. In *Advances in Neural Information Processing Systems*, pages 4765–4774, 2017.
- David Baehrens, Timon Schroeter, Stefan Harmeling, Motoaki Kawanabe, Katja Hansen, and Klaus-Robert Müller. How to explain individual classification decisions. *Journal of Machine Learning Research*, 11:1803–1831, aug 2010. ISSN 1532-4435.
- A. Apicella, F. Isgrò, R. Prevete, and G. Tamburrini. Contrastive explanations to classification systems using sparse dictionaries. In *Lecture Notes in Computer Science*, pages 207–218. Springer International Publishing, 2019.
- A. Apicella, F. Isgrò, R. Prevete, and G. Tamburrini. Middle-level features for the explanation of classification systems by sparse dictionary methods. *International Journal of Neural Systems*, 30(08):2050040, July 2020.
- Valentyn Boreiko, Maximilian Augustin, Francesco Croce, Philipp Berens, and Matthias Hein. Sparse visual counterfactual explanations in image space. In *Pattern Recognition: 44th DAGM German Conference, DAGM GCPD 2022, Konstanz, Germany, September 27–30, 2022, Proceedings*, pages 133–148. Springer, 2022.
- Tao Lei, Regina Barzilay, and Tommi Jaakkola. Rationalizing neural predictions. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, 2016.
- Jiwei Li, Will Monroe, and Dan Jurafsky. Understanding neural networks through representation erasure. *arXiv preprint arXiv:1612.08220*, 2016.
- Marcos V Treviso and André FT Martins. The explanation game: Towards prediction explainability through sparse communication. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*, 2020.
- Jasmijn Bastings, Wilker Aziz, and Ivan Titov. Interpretable neural predictions with differentiable binary variables. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2963–2977, 2019.
- Mo Yu, Shiyu Chang, Yang Zhang, and Tommi S Jaakkola. Rethinking cooperative rationalization: Introspective extraction and complement control. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4094–4103, 2019.
- Mo Yu, Yang Zhang, Shiyu Chang, and Tommi Jaakkola. Understanding interlocking dynamics of cooperative ra-

- tionalization. *Advances in Neural Information Processing Systems*, 34:12822–12835, 2021.
- Jianbo Chen, Le Song, Martin Wainwright, and Michael Jordan. Learning to explain: An information-theoretic perspective on model interpretation. In *International Conference on Machine Learning*, pages 883–892. PMLR, 2018.
- David Martens and Foster Provost. Explaining data-driven document classifications. *MIS quarterly*, 38(1):73–100, 2014.
- Berk Ustun, Alexander Spangher, and Yang Liu. Actionable recourse in linear classification. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, pages 10–19, 2019.
- Zhicheng Cui, Wenlin Chen, Yujie He, and Yixin Chen. Optimal action extraction for random forests and boosted trees. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 179–188, 2015.
- Gabriele Tolomei, Fabrizio Silvestri, Andrew Haines, and Mounia Lalmas. Interpretable predictions of tree-based ensembles via actionable feature tweaking. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 465–474, 2017.
- Ana Lucic, Harrie Oosterhuis, Hinda Haned, and Maarten de Rijke. Focus: Flexible optimizable counterfactual explanations for tree ensembles. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 5313–5322, 2022.
- Vivek Gupta, Pegah Nokhiz, Chitradeep Dutta Roy, and Suresh Venkatasubramanian. Equalizing recourse across groups. *arXiv preprint arXiv:1909.03166*, 2019.
- Xin Zhang, Armando Solar-Lezama, and Rishabh Singh. Interpreting neural network judgments via minimal, stable, and symbolic corrections. *Advances in Neural Information Processing Systems*, 31, 2018.
- Alexis Ross, Himabindu Lakkaraju, and Osbert Bastani. Learning models for actionable recourse. *Advances in Neural Information Processing Systems*, 34:18734–18746, 2021.
- Sandra Wachter, Brent Mittelstadt, and Chris Russell. Counterfactual explanations without opening the black box: Automated decisions and the GDPR. *Harv. JL & Tech.*, 31:841, 2017.
- Rory McGrath, Luca Costabello, Chan Le Van, Paul Sweeney, Farbod Kamiab, Zhao Shen, and Freddy Lecue. Interpretable credit application predictions with counterfactual explanations. In *NIPS 2018-Workshop on Challenges and Opportunities for AI in Financial Services: the Impact of Fairness, Explainability, Accuracy, and Privacy*, 2018.
- Chris Russell. Efficient search for diverse coherent explanations. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, pages 20–28, 2019.
- Shalmali Joshi, Oluwasanmi Koyejo, Warut Vijitbenjaronk, Been Kim, and Joydeep Ghosh. Towards realistic individual recourse and actionable explanations in black-box decision making systems. *arXiv preprint arXiv:1907.09615*, 2019.
- Yanou Ramon, David Martens, Foster Provost, and Theodoros Evgeniou. A comparison of instance-level counterfactual explanation algorithms for behavioral and textual data: SEDC, LIME-C and SHAP-C. *Advances in Data Analysis and Classification*, 14:801–819, 2020.
- Thibault Laugel, Marie-Jeanne Lesot, Christophe Marsala, Xavier Renard, and Marcin Detyniecki. Comparison-based inverse classification for interpretability in machine learning. In *Information Processing and Management of Uncertainty in Knowledge-Based Systems. Theory and Foundations: 17th International Conference, IPMU 2018, Cádiz, Spain, June 11-15, 2018, Proceedings, Part I 17*, pages 100–111. Springer, 2018.
- Amit Dhurandhar, Pin-Yu Chen, Ronny Luss, Chun-Chen Tu, Paishun Ting, Karthikeyan Shanmugam, and Payel Das. Explanations based on the missing: Towards contrastive explanations with pertinent negatives. *Advances in Neural Information Processing Systems*, 31, 2018.
- Arnaud Van Looveren and Janis Klaise. Interpretable counterfactual explanations guided by prototypes. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 650–665. Springer, 2021.
- Furui Cheng, Yao Ming, and Huamin Qu. Dece: Decision explorer with counterfactual explanations for machine learning models. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):1438–1447, 2020.
- Michael T Lash, Qihang Lin, Nick Street, Jennifer G Robinson, and Jeffrey Ohlmann. Generalized inverse classification. In *Proceedings of the 2017 SIAM International Conference on Data Mining*, pages 162–170. SIAM, 2017a.
- Michael T Lash, Qihang Lin, W Nick Street, and Jennifer G Robinson. A budget-constrained inverse classification framework for smooth classifiers. In *2017 IEEE International Conference on Data Mining Workshops (ICDMW)*, pages 1184–1193. IEEE, 2017b.
- Carlos Fernández-Loría, Foster Provost, and Xintian Han. Explaining data-driven decisions made by ai systems: The counterfactual approach. *MIS Quarterly*, 46(3):1635–1660, 2022.
- Eoin Delaney, Arjun Pakrashi, Derek Greene, and Mark T Keane. Counterfactual explanations for misclassified images: How human and machine explanations differ. *Artificial Intelligence*, 324:103995, 2023.

Barry Becker and Ronny Kohavi. Adult. UCI Machine Learning Repository, 1996. DOI: <https://doi.org/10.24432/C5XW20>.

Dheeru Dua and Casey Graff. UCI machine learning repository, 2017. URL <http://archive.ics.uci.edu/ml>.

FICO. Explainable Machine Learning Challenge, 2018. URL <https://community.fico.com/s/explainable-machine-learning-challenge>. Accessed: 2018-11-02.

Hans Hofmann. Statlog (German Credit Data). UCI Machine Learning Repository, 1994. DOI: <https://doi.org/10.24432/C5NC77>.

A. Johnson, T. Pollard, and R. Mark. MIMIC-III clinical database. <https://physionet.org/content/mimiciii/>, 2016a.

Alistair E.W. Johnson, Tom J. Pollard, Lu Shen, Li wei H. Lehman, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G. Mark. MIMIC-III, a freely accessible critical care database. *Scientific Data*, 3(1), May 2016b. doi: 10.1038/sdata.2016.35. URL <https://doi.org/10.1038/sdata.2016.35>.

Lauren Kirchner Jeff Larson, Surya Mattu and Julia Angwin. How we analyzed the COMPAS recidivism algorithm. <https://www.propublica.org/article/how-we-analyzed-the-compas-recidivism-algorithm>, 2016. Accessed: 2023-02-01.

Caroline Wang, Bin Han, Bhrij Patel, and Cynthia Rudin. In Pursuit of Interpretable, Fair and Accurate Machine Learning for Criminal Recidivism Prediction. *Journal of Quantitative Criminology*, pages 1–63, 2022. ISSN 0748-4518. doi: 10.1007/s10940-022-09545-w.

Beata Strack, Jonathan P. DeShazo, Chris Gennings, Juan L. Olmo, Sebastian Ventura, Krzysztof J. Cios, and John N. Clore. Impact of HbA1c measurement on hospital readmission rates: Analysis of 70, 000 clinical database patient records. *BioMed Research International*, 2014:1–11, 2014.

Ramaravind K Mothilal, Amit Sharma, and Chenhao Tan. Explaining machine learning classifiers through diverse counterfactual explanations. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pages 607–617, 2020.

F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.

Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary

DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. PyTorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019. URL <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>.

Checklist

1. For all models and algorithms presented, check if you include:
 - (a) A clear description of the mathematical setting, assumptions, algorithm, and/or model. [Yes]
 - (b) An analysis of the properties and complexity (time, space, sample size) of any algorithm. [Yes]
 - (c) (Optional) Anonymized source code, with specification of all dependencies, including external libraries. [Yes]
2. For any theoretical claim, check if you include:
 - (a) Statements of the full set of assumptions of all theoretical results. [Yes]
 - (b) Complete proofs of all theoretical results. [Yes]
 - (c) Clear explanations of any assumptions. [Yes]
3. For all figures and tables that present empirical results, check if you include:
 - (a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL). [Yes]
 - (b) All the training details (e.g., data splits, hyperparameters, how they were chosen). [Yes]
 - (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). [Yes]
 - (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). [Yes]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:
 - (a) Citations of the creator If your work uses existing assets. [Yes]
 - (b) The license information of the assets, if applicable. [Yes]
 - (c) New assets either in the supplemental material or as a URL, if applicable. [Yes]
 - (d) Information about consent from data providers/curators. [Not Applicable]

- (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. [Not Applicable]
5. If you used crowdsourcing or conducted research with human subjects, check if you include:
- (a) The full text of instructions given to participants and screenshots. [Not Applicable]
 - (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. [Not Applicable]
 - (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. [Not Applicable]

A Datasets Description

All datasets were partitioned into training and test subsets according to an 80-20 stratified split. Numeric features were linearly rescaled to have mean zero and variance one. Categorical features with k unique levels were one-hot encoded into k dummy variables. Binary features were one-hot encoded into a single dummy variable. Dataset sizes pre- and post-encoding are displayed in Table 6. The reference for features is given by the feature mean for continuous features and the mode for binary and categorical features.

	Observations	Pre-Encoded Features	Post-Encoded Features
COMPAS	6,907	7	7
Adult	32,561	14	107
MIMIC-III	48,786	14	14
Diabetes	101,766	33	101
German	1,000	20	59
Credit			
FICO	10,459	23	23

Table 6: Training Dataset Sizes

COMPAS

The COMPAS dataset consists of data on criminal recidivism in Broward County, Florida (Jeff Larson and Angwin, 2016). The objective is to predict two-year recidivism using the following features: sex, age, number of priors, number of juvenile felonies/misdemeanors/crimes, and whether the current charge is a felony.

Adult

The Adult data is based on U.S. census data including demographics, education and job information, marital status, and capital gain/loss (Becker and Kohavi, 1996; Dua and Graff, 2017). The target outcome is whether an individual’s salary is above \$50,000.

MIMIC-III

MIMIC-III is a database containing a variety of medical information for patients during ICU stays in the Beth Israel Deaconess Medical Center (Johnson et al., 2016a,b). We take the ‘hospital expires flag’, a binary indicator of whether or not a patient died in the given hospitalization, to be the outcome of interest and select the following subset of variables as features: age, preiculos (pre-ICU length of stay), gcs (Glasgow Coma Scale), heartrate_min, heartrate_max, meanbp_min (min blood pressure), meanbp_max (max blood pressure), resprate_min, resprate_max, tempc_min, tempc_max, urineoutput, mechvent (whether the patient is on mechanical ventilation), and electivesurgery (whether the patient had elective surgery).

Diabetes

The Diabetes dataset is derived from 10 years (1999-2008) of clinical care at 130 US hospitals and integrated delivery networks (Strack et al., 2014; Dua and Graff, 2017). It comprises over 50 features representing patient and hospital outcomes. The dataset includes attributes such as race, gender, age, admission type, time in hospital, the medical specialty of the admitting physician, number of lab tests performed, number of medications, etc. We take whether the patient will revisit the hospital in 2 years as a binary indicator.

German Credit

The German Credit data (Hofmann, 1994; Dua and Graff, 2017) uses financial and demographic markers (checking account status, credit history, employment/marital status, etc.) to predict whether an individual is at risk for defaulting on a loan.

FICO

The FICO Home Equity Line of Credit (HELOC) dataset used for the Explainable Machine Learning Challenge contains a variety of financial markers (number of inquiries into a user’s account, max delinquency, number of satisfactory trades, etc.) for various loan applicants (FICO, 2018). The target outcome is whether a consumer was ever 90+ days past due in any 2-year period since they opened their account.

B Model Training

Baseline models were fit using `sklearn` (Pedregosa et al., 2011) implementations in Python. The logistic regression models L1 LR and L2 LR were fit using regularization parameter $C = 0.01$. The 2-layer MLP used ReLU activation and consisted of two fully-connected layers with 128 nodes each. It was trained with early stopping. The gradient-boosted classifier used 200 trees with a max depth of 3.

The optimized models were trained by adding the SEV losses from Section 4 to the standard loss term for the models. For GBDT, the training goal is to reweight the trees from the reference GBDT model. The resulting loss was minimized via gradient descent in `PyTorch` (Paszke et al., 2019), with a batch size of 128, a learning rate of 0.1, and the Adam optimizer. To maintain high accuracy, the first 70 training epochs are warm-up epochs optimizing just Binary Cross Entropy Loss for classification (`BCELoss`). The next 30 epochs add the All-Opt and Vol-Opt terms and the reference positive penalty term to encourage low SEV values.

C Model Results

In this section, we compare our optimized classifiers to the reference classifiers in terms of accuracy/AUC and SEV^+/SEV^- .

C.1 Results Overview

Here we report training and testing accuracy, AUC, mean SEV^+ , and mean SEV^- . Table entries are means and standard deviations across 10 random train-test splits for each dataset. The optimized mean SEV^+ and SEV^- are marked in red from Table 7 to Table 24. More detailed SEV counts for each dataset before and after optimization are shown in Appendix C.2. Blank entries are caused by incompatible loss functions and objectives (e.g., All-Opt⁺, designed to target SEV^+ , and SEV^-). From these tables, we observe that with the application of All-Opt⁺ and Vol-Opt methods for SEV^+ and All-Opt⁻ for SEV^- , we can obtain models with lower SEV^+ or SEV^- without significant performance loss. This means that we are able to find models with sparse explanations for more queries, without reducing the predictive power of the model.

	Train		Test			
	Accuracy	AUC	Accuracy	AUC	Mean SEV^+	Mean SEV^-
L2	0.85 ± 0.00	0.90 ± 0.00	0.85 ± 0.02	0.90 ± 0.00	2.27 ± 0.11	1.18 ± 0.15
L1	0.85 ± 0.00	0.90 ± 0.00	0.85 ± 0.01	0.90 ± 0.00	1.96 ± 0.07	1.13 ± 0.01
All-Opt ⁺	0.85 ± 0.00	0.90 ± 0.00	0.85 ± 0.01	0.90 ± 0.01	1.10 ± 0.26	
All-Opt ⁻	0.85 ± 0.00	0.90 ± 0.00	0.84 ± 0.01	0.90 ± 0.00		1.03 ± 0.02
Vol-Opt	0.84 ± 0.01	0.90 ± 0.02	0.84 ± 0.01	0.89 ± 0.02	1.01 ± 0.00	

Table 7: SEV Optimization Performance of linear classifiers on Adult

	Train		Test			
	Accuracy	AUC	Accuracy	AUC	Mean SEV^+	Mean SEV^-
MLP	0.87 ± 0.00	0.93 ± 0.00	0.86 ± 0.00	0.92 ± 0.00	2.97 ± 0.01	1.21 ± 0.04
All-Opt ⁺	0.86 ± 0.01	0.92 ± 0.01	0.84 ± 0.00	0.90 ± 0.00	1.01 ± 0.01	
All-Opt ⁻	0.86 ± 0.00	0.92 ± 0.00	0.85 ± 0.00	0.93 ± 0.00		1.03 ± 0.03

Table 8: SEV Optimization Performance of multi-layer perceptrons on Adult

	Train		Test			
	Accuracy	AUC	Accuracy	AUC	Mean SEV^+	Mean SEV^-
GBDT	0.88 ± 0.00	0.93 ± 0.00	0.87 ± 0.00	0.93 ± 0.00	2.66 ± 0.17	1.21 ± 0.02
All-Opt ⁺	0.88 ± 0.00	0.93 ± 0.01	0.87 ± 0.00	0.90 ± 0.00	1.03 ± 0.01	
All-Opt ⁻	0.87 ± 0.01	0.93 ± 0.01	0.85 ± 0.02	0.91 ± 0.01		1.00 ± 0.03

Table 9: SEV Optimization Performance of gradient boosting decision trees on Adult

	Train		Test			
	Accuracy	AUC	Accuracy	AUC	Mean SEV ⁺	Mean SEV ⁻
L2	0.68 ± 0.00	0.73 ± 0.01	0.68 ± 0.01	0.73 ± 0.02	1.04 ± 0.01	1.25 ± 0.03
L1	0.68 ± 0.00	0.73 ± 0.00	0.68 ± 0.01	0.73 ± 0.02	1.06 ± 0.01	1.25 ± 0.03
All-Opt ⁺	0.67 ± 0.01	0.72 ± 0.01	0.66 ± 0.01	0.71 ± 0.02	1.03 ± 0.03	
All-Opt ⁻	0.65 ± 0.01	0.72 ± 0.01	0.65 ± 0.02	0.72 ± 0.01		1.05 ± 0.07
Vol-Opt	0.65 ± 0.01	0.72 ± 0.02	0.65 ± 0.01	0.71 ± 0.02	1.03 ± 0.05	

Table 10: SEV Optimization Performance of linear classifiers on COMPAS

	Train		Test			
	Accuracy	AUC	Accuracy	AUC	Mean SEV ⁺	Mean SEV ⁻
MLP	0.68 ± 0.01	0.74 ± 0.00	0.68 ± 0.01	0.73 ± 0.02	1.01 ± 0.01	1.53 ± 0.27
All-Opt ⁺	0.68 ± 0.01	0.73 ± 0.00	0.68 ± 0.01	0.73 ± 0.01	1.00 ± 0.00	
All-Opt ⁻	0.68 ± 0.00	0.73 ± 0.00	0.68 ± 0.00	0.74 ± 0.01		1.24 ± 0.14

Table 11: SEV Optimization Performance of multi-layer perceptrons on COMPAS

	Train		Test			
	Accuracy	AUC	Accuracy	AUC	Mean SEV ⁺	Mean SEV ⁻
GBDT	0.70 ± 0.00	0.76 ± 0.00	0.68 ± 0.00	0.73 ± 0.00	1.12 ± 0.03	1.13 ± 0.03
All-Opt ⁺	0.70 ± 0.01	0.76 ± 0.01	0.66 ± 0.01	0.70 ± 0.01	1.01 ± 0.01	
All-Opt ⁻	0.70 ± 0.01	0.77 ± 0.01	0.66 ± 0.01	0.70 ± 0.01		1.01 ± 0.01

Table 12: SEV Optimization Performance of gradient boosting decision trees on COMPAS

	Train		Test			
	Accuracy	AUC	Accuracy	AUC	Mean SEV ⁺	Mean SEV ⁻
L2	0.89 ± 0.00	0.80 ± 0.00	0.89 ± 0.01	0.80 ± 0.02	4.27 ± 0.04	1.16 ± 0.02
L1	0.89 ± 0.00	0.80 ± 0.00	0.89 ± 0.00	0.80 ± 0.01	4.36 ± 0.01	1.14 ± 0.03
All-Opt ⁺	0.89 ± 0.00	0.78 ± 0.01	0.89 ± 0.00	0.78 ± 0.01	3.58 ± 0.33	
All-Opt ⁻	0.89 ± 0.00	0.78 ± 0.01	0.89 ± 0.00	0.78 ± 0.01		1.06 ± 0.06
Vol-Opt	0.88 ± 0.02	0.76 ± 0.01	0.88 ± 0.03	0.76 ± 0.01	3.70 ± 0.71	

Table 13: SEV Optimization Performance of linear classifiers on MIMIC III

	Train		Test			
	Accuracy	AUC	Accuracy	AUC	Mean SEV ⁺	Mean SEV ⁻
MLP	0.90 ± 0.00	0.86 ± 0.01	0.90 ± 0.00	0.85 ± 0.01	4.27 ± 0.10	1.21 ± 0.03
All-Opt ⁺	0.89 ± 0.00	0.81 ± 0.01	0.89 ± 0.00	0.81 ± 0.01	2.34 ± 0.77	
All-Opt ⁻	0.89 ± 0.00	0.82 ± 0.01	0.89 ± 0.00	0.82 ± 0.01		1.02 ± 0.01

Table 14: SEV Optimization Performance of multi-layer perceptrons on the MIMIC III

	Train		Test			
	Accuracy	AUC	Accuracy	AUC	Mean SEV ⁺	Mean SEV ⁻
GBDT	0.91 ± 0.00	0.87 ± 0.00	0.90 ± 0.00	0.85 ± 0.01	5.26 ± 0.08	1.19 ± 0.03
All-Opt ⁺	0.90 ± 0.00	0.87 ± 0.00	0.89 ± 0.01	0.83 ± 0.01	3.27 ± 0.02	
All-Opt ⁻	0.90 ± 0.00	0.87 ± 0.00	0.89 ± 0.01	0.83 ± 0.01		1.02 ± 0.02

Table 15: SEV Optimization Performance of gradient boosting decision trees on MIMIC III

	Train		Test			
	Accuracy	AUC	Accuracy	AUC	Mean SEV ⁺	Mean SEV ⁻
L2	0.79 ± 0.01	0.85 ± 0.00	0.73 ± 0.03	0.76 ± 0.03	1.00 ± 0.00	1.75 ± 0.16
L1	0.76 ± 0.01	0.80 ± 0.01	0.73 ± 0.02	0.77 ± 0.03	1.00 ± 0.00	1.13 ± 0.01
All-Opt ⁺	0.79 ± 0.01	0.83 ± 0.01	0.75 ± 0.05	0.77 ± 0.05	1.00 ± 0.00	
All-Opt ⁻	0.78 ± 0.01	0.83 ± 0.01	0.75 ± 0.05	0.77 ± 0.05		1.04 ± 0.05
Vol-Opt	0.77 ± 0.01	0.83 ± 0.01	0.70 ± 0.05	0.75 ± 0.05		1.00 ± 0.00

Table 16: SEV Optimization Performance of linear classifiers on German Credit

	Train		Test			
	Accuracy	AUC	Accuracy	AUC	Mean SEV ⁺	Mean SEV ⁻
MLP	0.80 ± 0.00	0.86 ± 0.02	0.76 ± 0.03	0.79 ± 0.03	1.04 ± 0.08	1.58 ± 0.22
All-Opt ⁺	1.00 ± 0.03	1.00 ± 0.03	0.73 ± 0.02	0.80 ± 0.04	1.00 ± 0.00	
All-Opt ⁻	1.00 ± 0.01	1.00 ± 0.02	0.76 ± 0.04	0.81 ± 0.05		1.18 ± 0.12

Table 17: SEV Optimization Performance of multi-layer perceptrons on German Credit

	Train		Test			
	Accuracy	AUC	Accuracy	AUC	Mean SEV ⁺	Mean SEV ⁻
GBDT	0.96 ± 0.00	0.99 ± 0.00	0.74 ± 0.03	0.78 ± 0.02	1.10 ± 0.24	1.50 ± 0.06
All-Opt ⁺	1.00 ± 0.00	1.00 ± 0.00	0.73 ± 0.02	0.76 ± 0.02	1.01 ± 0.01	
All-Opt ⁻	1.00 ± 0.00	1.00 ± 0.00	0.74 ± 0.02	0.76 ± 0.03		1.17 ± 0.07

Table 18: SEV Optimization Performance of gradient boosting decision trees on German Credit

	Train		Test			
	Accuracy	AUC	Accuracy	AUC	Mean SEV ⁺	Mean SEV ⁻
L2	0.72 ± 0.00	0.78 ± 0.00	0.71 ± 0.01	0.78 ± 0.01	1.28 ± 0.02	2.76 ± 0.08
L1	0.71 ± 0.00	0.78 ± 0.00	0.71 ± 0.01	0.78 ± 0.01	1.12 ± 0.01	2.46 ± 0.07
All-Opt ⁺	0.71 ± 0.00	0.77 ± 0.00	0.70 ± 0.01	0.77 ± 0.01	1.03 ± 0.02	
All-Opt ⁻	0.68 ± 0.01	0.75 ± 0.01	0.68 ± 0.01	0.75 ± 0.01		1.14 ± 0.04
Vol-Opt	0.70 ± 0.02	0.77 ± 0.01	0.69 ± 0.01	0.76 ± 0.02	1.04 ± 0.04	

Table 19: SEV Optimization Performance of linear classifiers on FICO

	Train		Test			
	Accuracy	AUC	Accuracy	AUC	Mean SEV ⁺	Mean SEV ⁻
MLP	0.73 ± 0.01	0.81 ± 0.01	0.72 ± 0.01	0.79 ± 0.01	1.01 ± 0.01	2.83 ± 0.22
All-Opt ⁺	0.68 ± 0.01	0.78 ± 0.01	0.67 ± 0.01	0.78 ± 0.02	1.01 ± 0.04	
All-Opt ⁻	0.77 ± 0.00	0.85 ± 0.00	0.67 ± 0.01	0.75 ± 0.01		1.52 ± 0.44

Table 20: SEV Optimization Performance of multi-layer perceptrons on FICO

	Train		Test			
	Accuracy	AUC	Accuracy	AUC	Mean SEV ⁺	Mean SEV ⁻
GBDT	0.77 ± 0.00	0.85 ± 0.00	0.73 ± 0.03	0.80 ± 0.01	1.00 ± 0.00	3.58 ± 0.13
All-Opt ⁺	0.80 ± 0.00	0.98 ± 0.01	0.70 ± 0.01	0.75 ± 0.01	1.00 ± 0.00	
All-Opt ⁻	0.80 ± 0.01	0.89 ± 0.01	0.69 ± 0.01	0.75 ± 0.01		1.84 ± 0.10

Table 21: SEV Optimization Performance of gradient boosting decision trees on FICO

	Train		Test			
	Accuracy	AUC	Accuracy	AUC	Mean SEV ⁺	Mean SEV ⁻
L2	0.62 ± 0.00	0.67 ± 0.00	0.62 ± 0.00	0.67 ± 0.00	1.00 ± 0.00	1.53 ± 0.01
L1	0.62 ± 0.00	0.67 ± 0.00	0.62 ± 0.00	0.67 ± 0.00	1.00 ± 0.00	1.49 ± 0.01
All-Opt ⁺	0.56 ± 0.00	0.63 ± 0.00	0.56 ± 0.00	0.63 ± 0.00	1.00 ± 0.00	
All-Opt ⁻	0.56 ± 0.00	0.63 ± 0.00	0.56 ± 0.00	0.63 ± 0.00		1.00 ± 0.00
Vol-Opt	0.57 ± 0.04	0.62 ± 0.03	0.57 ± 0.03	0.61 ± 0.03	1.00 ± 0.00	

Table 22: SEV Optimization Performance of linear classifiers on Diabetes

	Train		Test			
	Accuracy	AUC	Accuracy	AUC	Mean SEV ⁺	Mean SEV ⁻
MLP	0.73 ± 0.01	0.81 ± 0.01	0.72 ± 0.01	0.79 ± 0.01	1.01 ± 0.01	2.83 ± 0.22
All-Opt ⁺	0.68 ± 0.01	0.78 ± 0.01	0.67 ± 0.01	0.78 ± 0.02	1.01 ± 0.04	
All-Opt ⁻	0.77 ± 0.00	0.85 ± 0.00	0.67 ± 0.01	0.75 ± 0.01		1.52 ± 0.44


Table 23: SEV Optimization Performance of multi-layer perceptrons on Diabetes

	Train		Test			
	Accuracy	AUC	Accuracy	AUC	Mean SEV ⁺	Mean SEV ⁻
GBDT	0.77 ± 0.00	0.85 ± 0.00	0.73 ± 0.03	0.80 ± 0.01	1.00 ± 0.00	3.58 ± 0.13
All-Opt ⁺	0.80 ± 0.00	0.98 ± 0.01	0.70 ± 0.01	0.75 ± 0.01	1.00 ± 0.00	
All-Opt ⁻	0.80 ± 0.01	0.89 ± 0.01	0.69 ± 0.01	0.75 ± 0.01		1.84 ± 0.10


Table 24: SEV Optimization Performance of gradient boosting decision trees on Diabetes

C.2 Detailed SEV Distributions Plots


In this section, we will display the SEV^+ and SEV^- counts distributions for optimized and unoptimized models for a single train-test split of each dataset. The distribution of the SEV^+ and SEV^- counts are shown in subfigures (a), while the accompanying Sankey plot shows how counts change due to optimization. The red on the left side of the Sankey Plots shows the distribution of the SEV before optimization, while the red on the right shows the SEV distribution after optimization.




(a) SEV^+ counts across methods



(b) L2 LR SEV^+ to All-Opt+ LR



(c) L1 LR SEV^+ to All-Opt+ LR



(d) L2 LR SEV^+ to Vol-Opt LR

Figure 6: SEV^+ Optimization performance for linear classifiers on Adult




Figure 7: SEV⁻ optimization performance for linear classifiers on Adult




Figure 8: SEV⁺ optimization performance for multi-layer perceptions on Adult

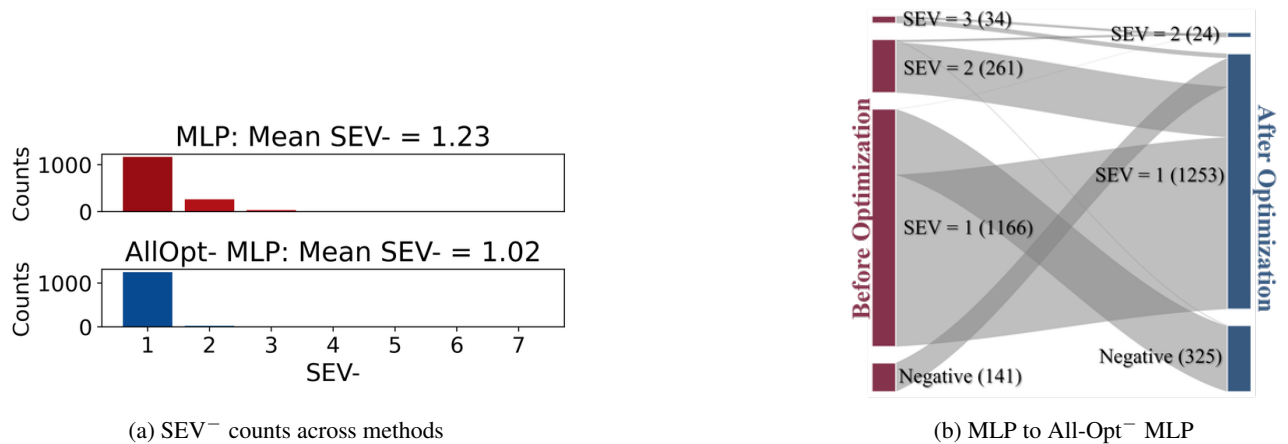




Figure 9: SEV⁻ optimization performance for multi-layer perceptions on Adult




(a) SEV⁺ counts across methods




(b) GBDT to All-Opt⁺ GBDT

Figure 10: SEV⁺ optimization performance for gradient boosting trees on Adult




(a) SEV⁻ counts across methods




(b) GBDT to All-Opt⁻ GBDT

Figure 11: SEV⁻ optimization performance for gradient boosting trees on Adult



(a) SEV⁺ counts across methods



(b) L2 LR SEV⁺ to All-Opt⁺ LR

Figure 12: SEV⁺ Optimization performance for linear classifiers on COMPAS

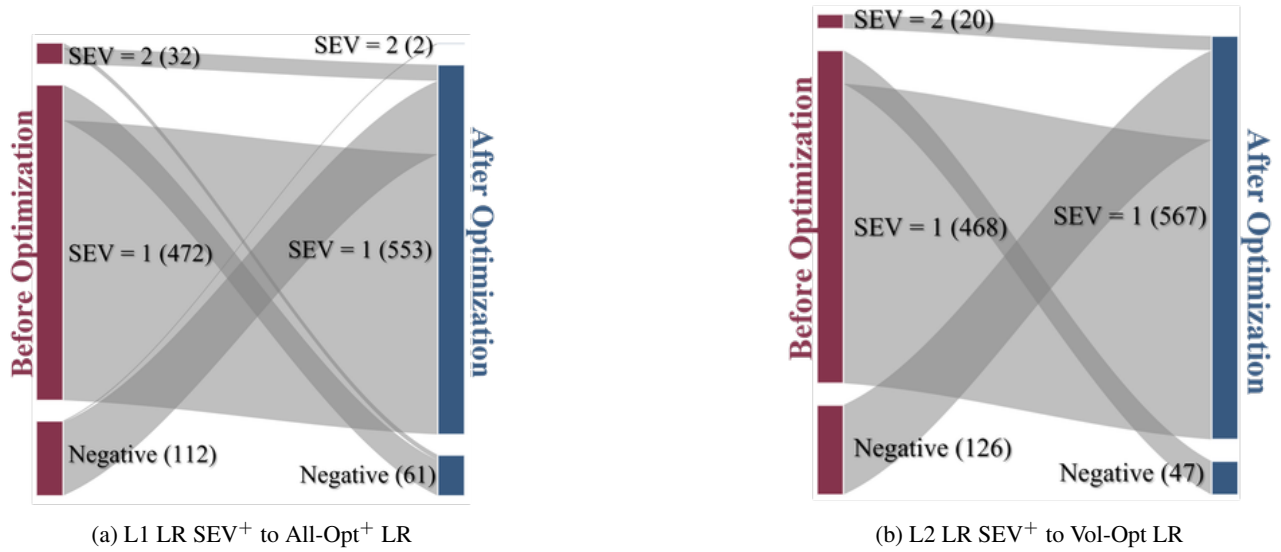


Figure 13: SEV⁺ Optimization performance for linear classifiers on COMPAS(cont.)




Figure 14: SEV⁻ optimization performance for linear classifiers on COMPAS






Figure 15: SEV⁺ optimization performance for multi-layer perceptions on COMPAS




(a) SEV⁻ counts across methods




(b) MLP to All-Opt⁻ MLP

Figure 16: SEV⁻ optimization performance for multi-layer perceptions on COMPAS




(a) SEV⁺ counts across methods




(b) GBDT to All-Opt⁺ GBDT

Figure 17: SEV⁺ optimization performance for gradient boosting trees on COMPAS




(a) SEV⁻ counts across methods

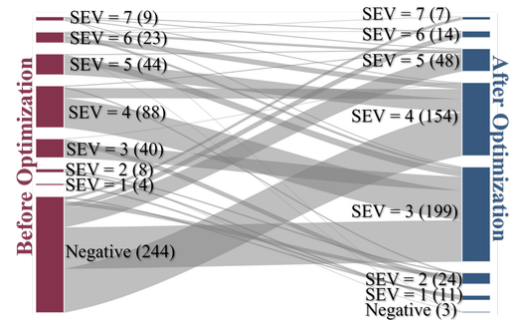


(b) GBDT to All-Opt⁻ GBDT

Figure 18: SEV⁻ optimization performance for gradient boosting trees on COMPAS




(a) SEV+ counts across methods




(b) L2 LR SEV+ to All-Opt+ LR

Figure 19: SEV+ Optimization performance for linear classifiers on MIMIC




(a) L1 LR SEV+ to All-Opt+ LR




(b) L2 LR SEV+ to Vol-Opt LR


Figure 20: SEV+ Optimization performance for linear classifiers on MIMIC(cont.)



(a) SEV- counts across methods




(b) L2 LR SEV- to All-Opt- LR




(c) L1 LR SEV- to All-Opt- LR

Figure 21: SEV- optimization performance for linear classifiers on MIMIC




(a) SEV⁺ counts across methods




(b) MLP to All-Opt⁺ MLP

Figure 22: SEV⁺ optimization performance for multi-layer perceptions on MIMIC




(a) SEV⁻ counts across methods




(b) MLP to All-Opt⁻ MLP

Figure 23: SEV⁻ optimization performance for multi-layer perceptions on MIMIC




(a) SEV⁺ counts across methods




(b) GBDT to All-Opt⁺ GBDT

Figure 24: SEV⁺ optimization performance for gradient boosting trees on MIMIC




(a) SEV⁻ counts across methods




(b) GBDT to All-Opt⁻ GBDT

Figure 25: SEV⁻ optimization performance for gradient boosting trees on MIMIC




(a) SEV⁺ counts across methods




(b) L2 LR SEV⁺ to All-Opt⁺ LR

Figure 26: SEV⁺ Optimization performance for linear classifiers on German Credit



(a) L1 LR SEV⁺ to All-Opt⁺ LR



(b) L2 LR SEV⁺ to Vol-Opt LR

Figure 27: SEV⁺ Optimization performance for linear classifiers on German Credit (cont.)




Figure 28: SEV^- optimization performance for linear classifiers on German Credit




Figure 29: SEV^+ optimization performance for multi-layer perceptions on German Credit

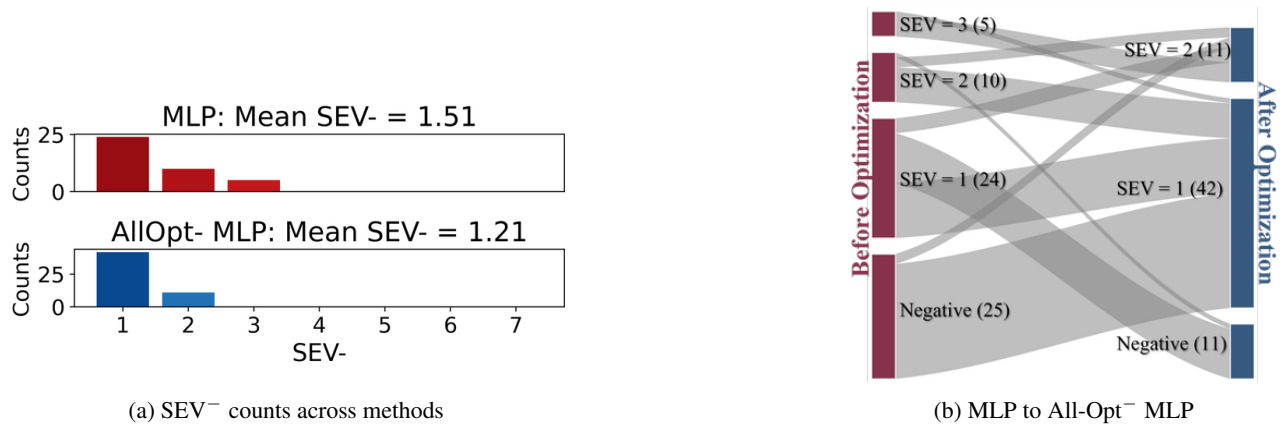




Figure 30: SEV^- optimization performance for multi-layer perceptions on German Credit




(a) SEV⁺ counts across methods




(b) GBDT to All-Opt⁺ GBDT

Figure 31: SEV⁺ optimization performance for gradient boosting trees on German Credit




(a) SEV⁻ counts across methods




(b) GBDT to All-Opt⁻ GBDT

Figure 32: SEV⁻ optimization performance for gradient boosting trees on German Credit



(a) SEV⁺ counts across methods



(b) L2 LR SEV⁺ to All-Opt⁺ LR

Figure 33: SEV⁺ Optimization performance for linear classifiers on FICO

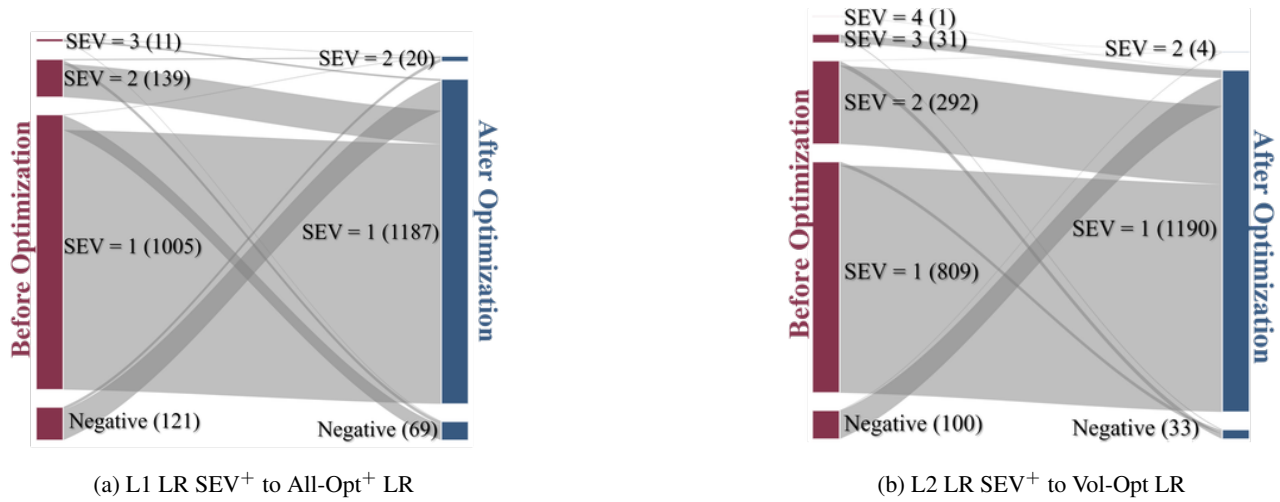


Figure 34: SEV⁺ Optimization performance for linear classifiers on FICO (cont.)




Figure 35: SEV⁻ optimization performance for linear classifiers on FICO

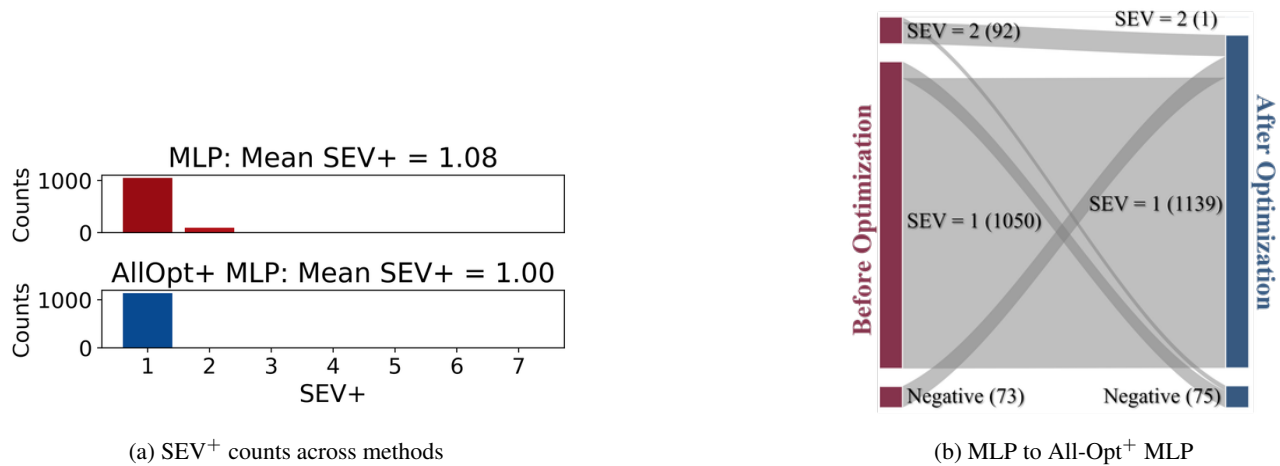




Figure 36: SEV⁺ optimization performance for multi-layer perceptions on FICO




(a) SEV⁻ counts across methods




(b) MLP to All-Opt⁻ MLP

Figure 37: SEV⁻ optimization performance for multi-layer perceptions on FICO




(a) SEV⁺ counts across methods




(b) GBDT to All-Opt⁺ GBDT

Figure 38: SEV⁺ optimization performance for gradient boosting trees on FICO




(a) SEV⁻ counts across methods




(b) GBDT to All-Opt⁻ GBDT

Figure 39: SEV⁻ optimization performance for gradient boosting trees on FICO




(a) SEV⁺ counts across methods




(b) L2 LR SEV⁺ to All-Opt⁺ LR

Figure 40: SEV⁺ Optimization performance for linear classifiers on Diabetes




(a) L1 LR SEV⁺ to All-Opt⁺ LR




(b) L2 LR SEV⁺ to Vol-Opt LR


Figure 41: SEV⁺ Optimization performance for linear classifiers on Diabetes(cont.)



(a) SEV⁻ counts across methods




(b) L2 LR SEV⁻ to All-Opt⁻ LR

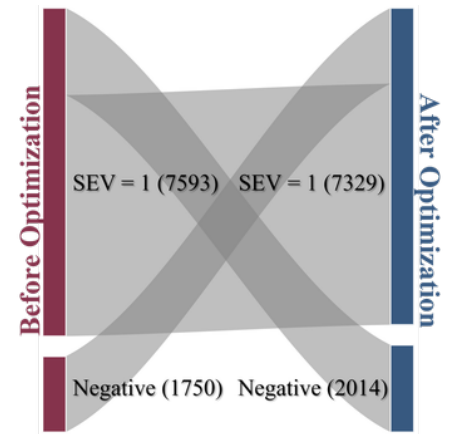


(c) L1 LR SEV⁻ to All-Opt⁻ LR

Figure 42: SEV⁻ optimization performance for linear classifiers on Diabetes




(a) SEV⁺ counts across methods




(b) MLP to All-Opt⁺ MLP

Figure 43: SEV⁺ optimization performance for multi-layer perceptions on Diabetes




(a) SEV⁻ counts across methods




(b) MLP to All-Opt⁻ MLP

Figure 44: SEV⁻ optimization performance for multi-layer perceptions on Diabetes




(a) SEV⁺ counts across methods




(b) GBDT to All-Opt⁺ GBDT

Figure 45: SEV⁺ optimization performance for gradient boosting trees on Diabetes



(a) SEV⁻ counts across methods



(b) GBDT to All-Opt⁻ GBDT

Figure 46: SEV⁻ optimization performance for gradient boosting trees on Diabetes

C.3 Local Explanations Comparison

In this section, we show how LIME-C, SHAP-C, DiCE, and SEV^- each explain a fixed GBDT reference classifier across different datasets. These plots are analogous to those in Figure 4b. These figures do not include models optimized for SEV^- .

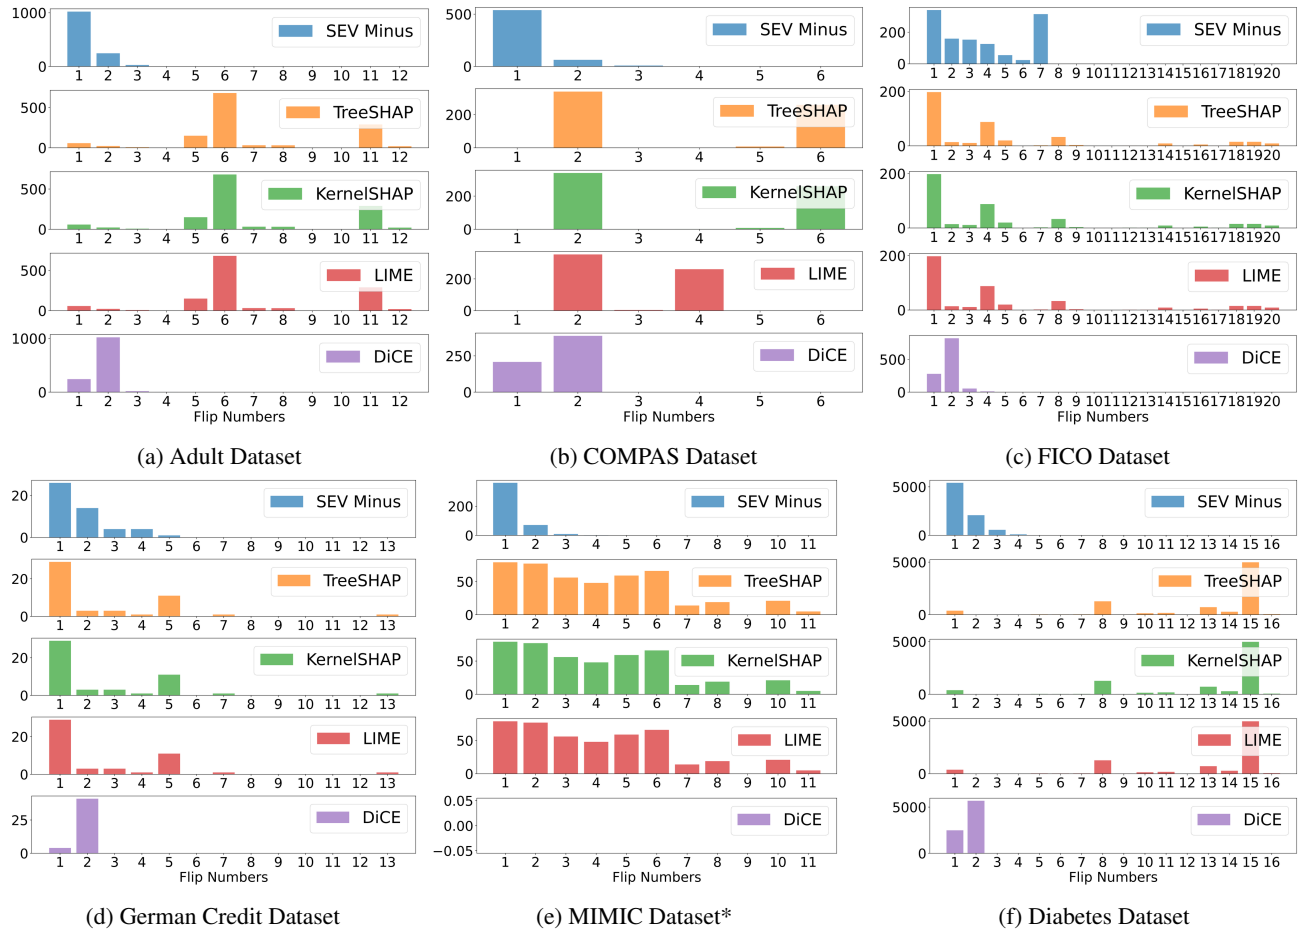


Figure 47: The flip counts for different local explanations in different datasets (continued)

* Values for DiCE are not plotted for MIMIC, as after one week of running it still had not generated counterfactual explanations. Additional details are shown in Table 5. Figure 47 shows that the local feature importance scores created by LIME-C and SHAP-C do not lead to sparse explanations. Some of the explanations need more than 5 features to explain one prediction. SEV^- and DiCE, on the other hand, do generate sparse explanations.


D Local Explanation Time Comparison between Kernel SHAP and SEV⁻

Here, we compare the time required to generate a local explanation for our method and KernelSHAP. We use **GBDT** as the baseline model and evaluate the median explanation runtime on queries in the test set.


Figure 48 and Table 25 give additional information on the runtime distribution across queries for KernelSHAP and SEV⁻. We see that for 3 out of the 6 datasets in which KernelSHAP is faster than SEV⁻ for explaining all queries, the total SEV⁻ runtime is skewed by a small number of queries. In all cases, with SEV⁻, a greater proportion of queries can be explained in less than 0.1 seconds than with KernelSHAP.

Dataset	Methods	max single query explanation time	% of query explanation time less than 0.1s	% of query explanation time less than 0.5s	% of query explanation time less than 1s
Adult	SEV ⁻	0.18	99.12	100.00	100.00
	KernelSHAP	0.48	83.44	100.00	100.00
German	SEV ⁻	0.83	88.46	98.08	100.00
	KernelSHAP	0.20	5.66	100.00	100.00
MIMIC	SEV ⁻	0.58	97.84	99.78	100.00
	KernelSHAP	0.18	1.94	100.00	100.00
FICO*	SEV ⁻	69.05	45.34	61.78	63.05
	KernelSHAP	0.19	0.00	100.00	100.00
COMPAS	SEV ⁻	0.01	100.00	100.00	100.00
	KernelSHAP	0.00	100.00	100.00	100.00
Diabetes	SEV ⁻	103.03	74.70	94.25	94.93
	KernelSHAP	0.20	11.26	100.00	100.00


Table 25: Time comparison between KernelSHAP and SEV⁻. The red parts emphasize that KernelSHAP has a smaller proportion of queries explained less than 0.1s. (*) For the FICO dataset, we set the max depth for SEV⁻ search to be 6, and we observe that more than half (61.78%) of the queries take less than 0.5s to generate the SEV⁻ explanations.




(a) Adult Dataset




(b) COMPAS Dataset




(c) FICO Dataset



(d) German Dataset



(e) MIMIC Dataset



(f) Diabetes Dataset

Figure 48: Time consumption distribution between KernelSHAP and SEV⁻ less than 1 second

E The performance of All-Opt[®]

This section will show that All-Opt[®] can be applied in different datasets. We use LR as the baseline model and compare the SEV[®] distribution, and the proportion of unreachable queries between L2 LR and All-Opt[®] LR for the Adult, COMPAS, MIMIC, and German Credit datasets. For Adult Dataset, age, marital-status, relationship, race, sex, native-country and occupation are considered as restricted features. For MIMIC Dataset, age and preiculos are considered restricted features. For COMPAS Dataset, gender=female and age are restricted features. For the German Credit Dataset, Age, Personal-status-sex, and Job are considered restricted. The results for test Accuracy, AUC, mean SEV[®], and the proportion of unreachable cases in explanations are shown in Tables 49a, 50a, 51a, and 52a. For unexplainable queries, we consider the sparse explanations as the number of features used in the model. Figures 49b, 50b, 51b, and 52b show the SEV distribution changes before and after optimization. Figures 49c, 50c, 51c, and 52c show the proportion of features used in each query before optimization and the restricted usage of features after optimization (restricted features are marked as red). All of the results show that All-Opt[®] performs well in avoiding unexplainable queries and can still effectively lower SEV.

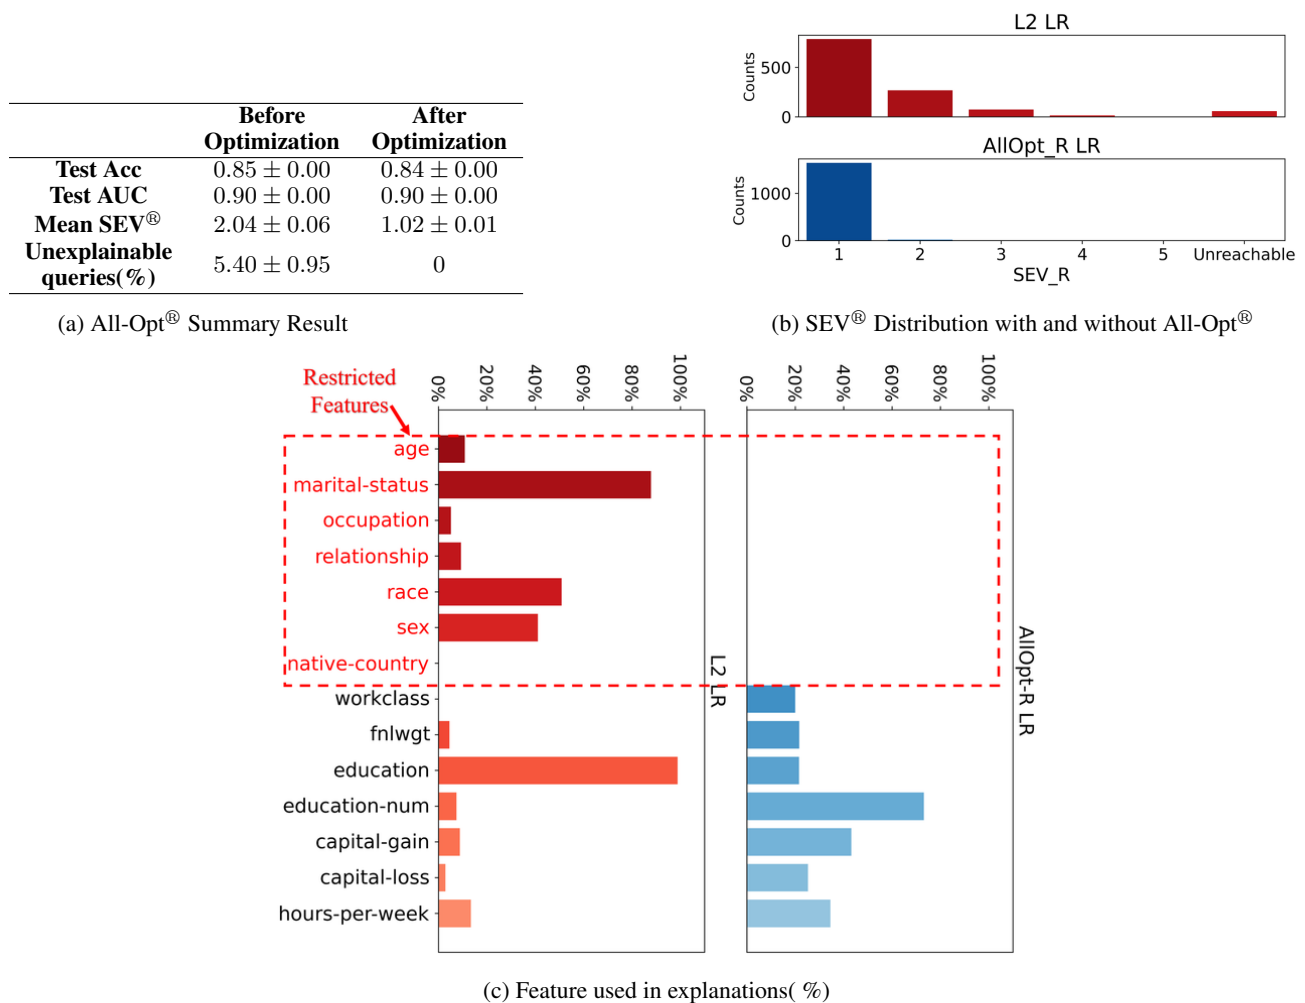


Figure 49: All-Opt[®] performance on Adult Dataset

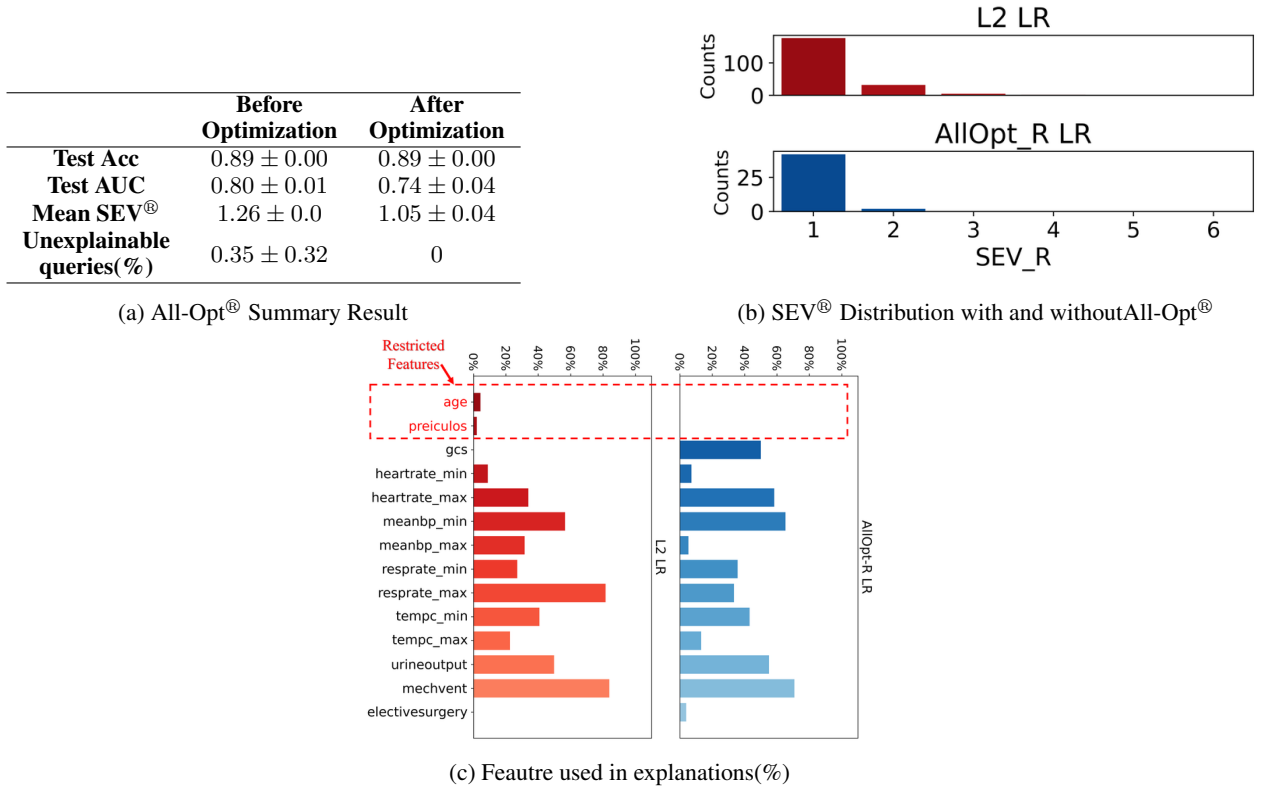


Figure 50: All-Opt[®] performance in MIMIC Dataset

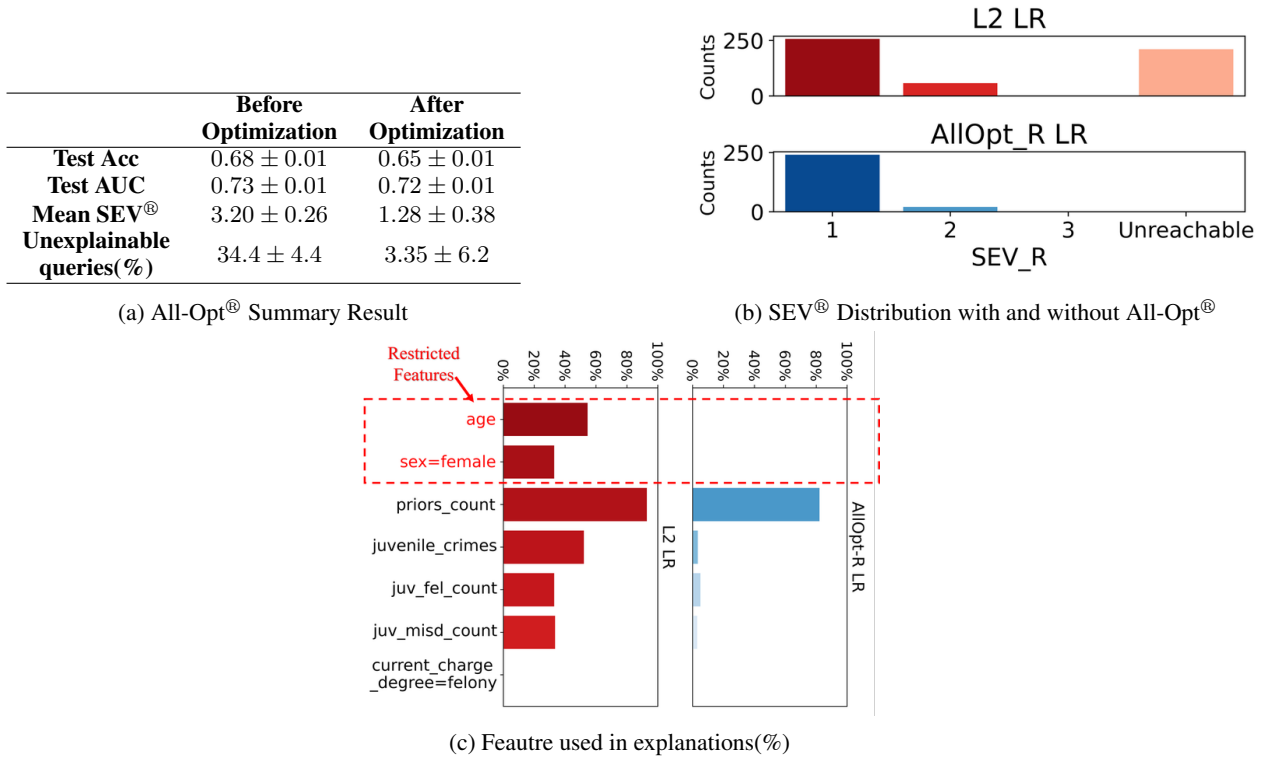




Figure 51: All-Opt[®] performance in COMPAS Dataset

	Before Optimization	After Optimization
Test Acc	0.74 ± 0.03	0.76 ± 0.03
Test AUC	0.79 ± 0.05	0.78 ± 0.04
Mean SEV [®]	1.27 ± 0.05	1.05 ± 0.06
Unexplainable Queries(%)	0	0

(a) All-Opt[®] Summary Result



(b) SEV[®] Distribution under All-Opt[®]



(c) Feautre used in explanations(%)

Figure 52: All-Opt[®] performance in German Credit Dataset

F Proof of Theorem 5

Theorem 5

Consider a linear classifier, $f(\mathbf{x}) := \mathbf{1}[(\beta_0 + \sum_{j=1}^p \beta_j x_j) > 0]$, where $\forall j, \beta_j \neq 0$, and for reference $\tilde{\mathbf{x}}$, we have $f(\tilde{\mathbf{x}}) = 0$ (i.e., reference predicted as negative). Let $g^{\text{ref}}(\boldsymbol{\beta}) = \beta_0 + \sum_{j=1}^p \beta_j \tilde{x}_j$ be the raw score of the classifier f at reference $\tilde{\mathbf{x}}$. For all $k \in \{2, 3, \dots, p\}$, the volume of region V_k in the input space \mathbb{R}^p with $\text{SEV}^+ = k$ is:

$$V_k = c_k \cdot \prod_{j=1}^p \left| \frac{g^{\text{ref}}(\boldsymbol{\beta})}{\beta_j} \right|, \quad (10)$$

where c_k is a finite constant unrelated to the β 's.

Proof:

Consider a simple linear classifier $h(\mathbf{x}) := \mathbf{1}[-1 + \sum_{j=1}^p x_j > 0]$. Let $\tilde{\mathbf{x}}$ be a reference where the reference feature values are all 0, i.e., $\tilde{\mathbf{x}} = \mathbf{0}$. Consider a point \mathbf{x} predicted as positive by the classifier h . The SEV^+ of \mathbf{x} is greater or equal to 2 if and only if $x_j \leq 1, \forall j \in \{1, \dots, p\}$. Let us show both directions.

If: If $x_j \leq 1, \forall j \in \{1, \dots, p\}$, for any vertex on the SEV hypercube one step away from the origin, which can be represented as $\mathbf{x}_v^{(j)} = \mathbf{e}_j \odot \mathbf{x}^{\text{query}} + (\mathbf{1} - \mathbf{e}_j) \odot \tilde{\mathbf{x}}$, we have $h(\mathbf{x}_v^{(j)}) = \mathbf{1}[x_j - 1 > 0] = 0$. This suggests that any vertex one step away from the origin fails to flip the prediction, which means $\text{SEV}^+ \geq 2$.

Only if: Suppose there exists k such that $x_k > 1$, then for the vertex $\mathbf{x}_v^{(k)} = \mathbf{e}_k \odot \mathbf{x}^{\text{query}} + (\mathbf{1} - \mathbf{e}_k) \odot \tilde{\mathbf{x}}$ in the SEV hypercube, $h(\mathbf{x}_v^{(k)}) = \mathbf{1}[-1 + x_k > 0] = 1$, which flips the sign of prediction but with $\text{SEV}^+ = 1$.

Therefore, the region with $\text{SEV}^+ \geq 2$ is the polytope comprised by the hyperplanes $(-1 + \sum_{j=1}^p x_j) > 0$, and $x_j \leq 1, \forall j \in \{1, \dots, p\}$. The volume of this polytope is $1 - \frac{1}{p!} = \frac{p!-1}{p!}$ (a unit hypercube whose volume is 1 with one corner cut off, and the volume of the corner is $\frac{1}{p!}$).

As the corner represents a unit simplex, its volume is well known to be $1/p!$. For completeness, however, we supply a brief proof. Define the corner as a set $S := \{(x_1, \dots, x_p)^T | 0 \leq x_i \leq 1, \sum_{m=1}^p x_m \leq 1\}$. Consider another set $T := \{(y_1, \dots, y_p)^T | 0 \leq y_1 \leq y_2 \leq \dots \leq y_p \leq 1\}$. There exists an invertible linear transformation, represented by a matrix $A := (a_{mj})$, where $a_{mm} = 1, a_{mj} = -1$ for $m = j + 1$, and $a_{mj} = 0$ otherwise, that defines a one-to-one mapping between points in S and points in T . To see this, consider a point $\mathbf{y} = (y_1, \dots, y_p)^T \in T$ and its mapping $A\mathbf{y} = (y_1, y_2 - y_1, \dots, y_p - y_{p-1})^T$. The monotonicity of the y_j ensures that each of these entries lies in $[0, 1]$. And the sum of the entries in $A\mathbf{y}$ is simply $\sum_{j=1}^p y_j = y_p \leq 1$. Therefore $A\mathbf{y} \in S$ and $\text{Vol}(S) = \text{Det}(A)\text{Vol}(T)$. $\text{Vol}(T) = \frac{1}{p!}$ because the volume is the same as the probability that a random permutation is in sorted order, and $\text{Det}(A) = 1$ since A is lower triangular with all diagonal elements equal to 1. Thus, $\text{Vol}(S) = \frac{1}{p!}$.

Within this polytope (unit cube with one corner cut off), suppose the proportion of the region with $\text{SEV}^+ = k$ is q_k , then the size of this region is $\frac{q_k(p!-1)}{p!}$.

Define $c_k := \frac{q_k(p!-1)}{p!}$. For a linear classifier, $f(\mathbf{x}) := \mathbf{1}[(\beta_0 + \sum_{j=1}^p \beta_j x_j) > 0]$, where for all j , and $f(\tilde{\mathbf{x}}) = 0$, let $g^{\text{ref}}(\boldsymbol{\beta})$ as the raw output of the linear classifier f , we can apply change of variables: $x'_j = -\frac{\beta_j(x_j - \tilde{x}_j)}{g^{\text{ref}}(\boldsymbol{\beta})}$ for all $j \in \{1, \dots, p\}$. Then, the classifier becomes $\mathbf{1}[-1 + \sum_{j=1}^p x'_j > 0]$ (the greater than sign remains because we assume $f(\tilde{\mathbf{x}}) = 0$, i.e., $g^{\text{ref}}(\boldsymbol{\beta}) = \beta_0 + \sum_{j=1}^p \beta_j \tilde{x}_j < 0$), and the reference value \tilde{x}'_j becomes 0. This is the same as the setting of g mentioned above. Hence, the volume of the region with $\text{SEV}^+ = k$ after the change of variables is c_k . Since the change of variable is a linear transformation, the volume of the same region before transformation can be obtained by dividing the new volume with the absolute value of the determinant of the transformation matrix, i.e.,

$$\begin{aligned} V_k &= c_k / \left| \prod_{j=1}^p -\frac{\beta_j}{g^{\text{ref}}(\boldsymbol{\beta})} \right| \\ &= c_k \cdot \prod_{j=1}^p \left| \frac{g^{\text{ref}}(\boldsymbol{\beta})}{\beta_j} \right|, \end{aligned}$$


where, as mentioned above, $c_k = \frac{q_k(p!-1)}{p!}$ is a finite constant unrelated to the β 's.

□


G The difference between explanations of Chen et al. (2018) and sparse explanation value

The work of Chen et al. (2018) is fundamentally different from ours in several ways.


- First, they impose that each query has an explanation of size k (manually determined), whereas SEV varies across queries and is often 1.
- Second, their explanations are designed (i.e., optimized) to set features *outside* the explanation to 0. Note that setting many features to 0 is likely to be out of distribution. Thus, predictions made at these points are not meaningful. (For instance, they would set all dummy categorical variables of a data point to 0, which means the point is no longer in any category, which is nonsensical.) Our SEV^- instead considers what would happen if we set positive features *inside* the explanation to a reference value, leaving all other features the same; these points are more likely to be in distribution. SEV^+ starts from a reference (which, when chosen properly, is in distribution), and some features are set to the feature value of the data point.
- Third, SEV can be used for many different kinds of models, whereas their approach must be applied only to differentiable models, such as neural networks. (It cannot be used for decision trees, for example.) Because they use a complex neural network as their model, the computation within the explanation is a black box – we do not know how the features are combined to form a prediction.
- Fourth, their approach requires optimization of their sparsity metric, which hurts accuracy (see Figure 53a to 54f). As we show in Appendix C.1, SEV is already small for black box models as shown in Section 5.1, so optimizing it might not even be necessary.



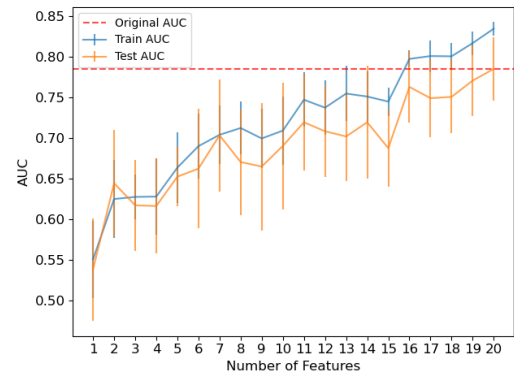
(a) The relationship between accuracy and k for L2X Model in Adult



(b) The relationship between AUC and k for L2X Model in Adult




(c) The relationship between accuracy and k for L2X Model in German Credit




(d) The relationship between AUC and k for L2X Model in German Credit

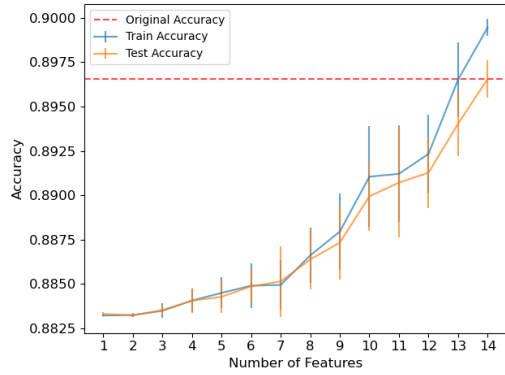
Figure 53: The relationship between model performance and k in L2X Model




(a) The relationship between accuracy and k for L2X Model in COMPAS




(b) The relationship between AUC and k for L2X Model in COMPAS




(c) The relationship between accuracy and k for L2X Model in MIMIC



(d) The relationship between AUC and k for L2X Model in MIMIC



(e) The relationship between accuracy and k for L2X Model in FICO



(f) The relationship between AUC and k for L2X Model in FICO

Figure 54: The relationship between model performance and k in L2X Model (cont.)

H Examples for SEV⁺ and SEV⁻ explanations

In this section, we show some examples of SEV⁺ and SEV⁻ explanations. Tables 26 to 29 show a few SEV⁺ and SEV⁻ calculations for various queries. For SEV⁺ explanations, we show the query, the reference, and values that are aligned to the query (from the reference). For SEV⁻, we show the values aligned to the reference (from the query). For instance, taking the first query of Table 26, the first row represents the query, and the second row is the reference. The following two rows show different SEV⁺ explanations: if we start from the reference and change its 9.6 year education to the query’s 13 years, and change the occupation from administrative clerical to executive manager, the income will be predicted as over 50K. Or, if, starting from the reference, education year is changed to 13, and the marital status is changed to civilian spouse, then income will be predicted as over 50K. The next three lines are the SEV⁻ explanations: starting from the query, if we change either the year of education, the hours worked per week, or the marital status from the query values to the reference values, the income will be predicted as less than 50K.

Table 26: Examples of SEV⁺/SEV⁻ Explanations for Adult Income Dataset

Type	age	education year	capital gain	capital loss	hour per week	workclass	marital status	occupation	relationship	race
Query	29	13	0	0	55	Private	Married-civ-spouse	Exec-managerial	Husband	Black
Reference	36.78	9.6	148.75	53.14	38.84	Private	Never-married	Adm-clerical	Not-in-family	White
SEV ⁺ Explanations	—	13	—	—	—	—	—	Exec-managerial	—	—
	—	13	—	—	—	—	Married-civ-spouse	—	—	—
SEV ⁻ Explanations	—	9.6	—	—	—	—	—	—	—	—
	—	—	—	—	—	—	Never-married	—	—	—
	—	—	—	—	38.84	—	—	—	—	—
Query	50	9	3103	0	40	Private	Married-civ-spouse	Craft-repair	Husband	White
Reference	36.78	9.6	148.75	53.14	38.84	Private	Never-married	Adm-clerical	Not-in-family	White
SEV ⁺ Explanation	—	—	3103	—	—	—	Married-civ-spouse	—	—	—
SEV ⁻ Explanation	36.78	—	—	—	—	—	—	—	—	—
	—	—	—	—	—	—	Never-married	—	—	—
	—	—	—	—	—	—	—	—	—	—
	—	—	148.75	—	—	—	—	—	—	—
Query	30	13	0	1902	40	Private	Married-civ-spouse	Adm-clerical	Husband	White
Reference	36.78	9.6	148.75	53.14	38.84	Private	Never-married	Adm-clerical	Not-in-family	White
SEV ⁺ Explanation	—	—	—	1902	—	—	Married-civ-spouse	—	—	—
SEV ⁻ Explanation	—	—	—	—	—	—	Never-married	—	—	—
	—	—	—	53.14	—	—	—	—	—	—
Query	45	10	0	0	65	Federal-gov	Married-civ-spouse	Transport-moving	Husband	White
Reference	36.78	9.6	148.75	53.14	38.84	Private	Never-married	Adm-clerical	Not-in-family	White
SEV ⁺ Explanation	—	—	—	—	65	Federal-gov	Married-civ-spouse	—	—	—
	—	—	0	—	65	—	Married-civ-spouse	—	—	—
SEV ⁻ Explanation	—	—	—	—	—	—	Never-married	—	—	—
Query	41	13	0	0	12	Private	Married-civ-spouse	Sales	Wife	White
Reference	36.78	9.6	148.75	53.14	38.84	Private	Never-married	Adm-clerical	Not-in-family	White
SEV ⁺ Explanation	—	13	—	—	—	—	Married-civ-spouse	—	Wife	—
SEV ⁻ Explanation	—	9.6	—	—	—	—	—	—	—	—
	—	—	—	—	—	—	Never-married	—	—	—
	—	—	—	—	—	—	—	—	Not-in-family	—
	—	—	148.75	—	—	—	—	—	—	—

Sparse and Faithful Explanations Without Sparse Models

Table 27: Examples of SEV⁺/SEV⁻ Explanations for MIMIC Dataset

Type	age	preiculus	gcs	heartrate min	meanbp min	resprate min	resprate max	tempc min	tempc max	urine output	mechvent
Query	59.2	0.4	3.0	49.0	54.0	9.0	16.0	34.3	36.3	240.0	1.0
Reference	72.8	2215.9	13.9	70.9	59.4	12.2	26.9	36.1	37.5	2024.1	0.0
SEV ⁺ Explanation	—	—	3.0	—	—	—	—	34.3	36.3	240.0	—
SEV ⁻ Explanation	—	—	13.9	—	—	—	—	—	—	—	—
	—	—	—	—	—	—	—	36.1	—	—	—
	—	—	—	—	—	—	—	—	37.5	—	—
	—	—	—	—	—	—	—	—	—	2024.1	—
	—	—	—	—	—	—	—	—	—	—	0.0
Query	85.8	1.6	15.0	87.0	61.0	18.0	26.0	35.9	36.4	183.0	1.0
Reference	72.8	2215.9	13.9	70.9	59.4	12.2	26.9	36.1	37.5	2024.1	0.0
SEV ⁺ Explanation	85.8	1.6	—	—	—	18.0	—	—	36.4	183.0	1.0
SEV ⁻ Explanation	72.8	—	—	—	—	—	—	—	—	—	—
	—	—	—	—	—	12.2	—	—	—	—	—
	—	—	—	—	—	—	—	—	—	2024.1	—
	—	—	—	—	—	—	—	—	—	—	0.0
Query	87.4	6953.8	3.0	38.0	47.0	11.0	45.0	35.9	37.4	852.0	1.0
Reference	72.8	2215.9	13.9	70.9	59.4	12.2	26.9	36.1	37.5	2024.1	0.0
SEV ⁺ Explanation	87.4	6953.8	3.0	38.0	—	—	45.0	—	—	852.0	—
	87.4	—	3.0	38.0	—	—	45.0	—	—	852.0	1.0
SEV ⁻ Explanation	72.8	—	—	—	—	—	—	—	—	—	—
	—	—	13.9	—	—	—	—	—	—	—	—
	—	—	—	70.9	—	—	—	—	—	—	—
	—	—	—	—	—	—	26.9	—	—	—	—
	—	—	—	—	—	—	—	—	—	2024.1	—
	—	—	—	—	—	—	—	—	—	—	0.0
Query	78.7	0.5	15.0	17.0	21.0	20.0	35.0	32.3	35.2	8.0	1.0
Reference	72.8	2215.9	13.9	70.9	59.4	12.2	26.9	36.1	37.5	2024.1	0.0
SEV ⁺ Explanation	—	—	—	17.0	21.0	—	—	—	—	8.0	—
	—	—	—	17.0	—	—	—	32.3	35.2	—	—
	—	—	—	17.0	—	—	—	32.3	—	8.0	—
	—	—	—	17.0	—	—	—	—	35.2	8.0	—
	—	—	—	—	21.0	—	—	32.3	—	8.0	—
	—	—	—	—	21.0	—	—	—	35.2	8.0	—
	—	—	—	—	—	20.0	—	32.3	35.2	—	—
	—	—	—	—	—	—	—	32.3	35.2	8.0	—
SEV ⁻ Explanation	—	—	—	—	—	—	—	36.1	—	2024.1	0.0

Table 28: Examples SEV⁺/SEV⁻ Explanations for FICO Dataset

Type	External RiskEstimate	MSinceMost RecentTradeOpen	NumSatisfactoryTrades	PercentTrades NeverDelq	MaxDelq2Public RecLast12M	PercentInstallTrades	NumInq Last6M	NumInqLast6 Mexcl7days	NetFraction RevolvingBurden	NumRevolving TradesWBalance
Query	79	138	1	100	7	100	0	0	Missing	Missing
Reference	72.21	9.2	21.1	89.98	5.36	29.82	0.6	0.56	22.26	2.95
SEV ⁺ Explanation	—	138	—	—	—	—	—	—	—	—
SEV ⁻ Explanation	—	—	—	—	—	100	—	—	—	—
Query	60	8	55	95	4	34	1	1	54	6
Reference	72.21	9.2	21.1	89.98	5.36	29.82	0.6	0.56	22.26	2.95
SEV ⁺ Explanation	—	—	—	—	—	—	—	—	54	—
SEV ⁻ Explanation	72.21	—	—	—	—	—	—	—	—	—
Query	59	12	18	85	2	30	10	10	95	5
Reference	72.21	9.2	21.1	89.98	5.36	29.82	0.6	0.56	22.26	2.95
SEV ⁺ Explanation	—	—	—	—	—	—	10	—	—	—
SEV ⁻ Explanation	72.21	—	—	—	5.36	—	0.6	0.56	22.26	—
Query	57	4	9	43	0	14	6	6	43	3
Reference	72.21	9.2	21.1	89.98	5.36	29.82	0.6	0.56	22.26	2.95
SEV ⁺ Explanation	—	—	—	43	—	—	—	—	—	—
SEV ⁻ Explanation	72.21	—	—	—	—	—	6	—	—	—
Query	82	61	3	100	7	33	0	0	50	1
Reference	72.21	9.2	21.1	89.98	5.36	29.82	0.6	0.56	22.26	2.95
SEV ⁺ Explanation	—	61	—	—	—	—	—	—	—	—
SEV ⁻ Explanation	—	9.2	21.1	—	—	—	—	—	—	—
SEV ⁻ Explanation	—	9.2	—	—	—	—	—	—	22.26	—

Sparse and Faithful Explanations Without Sparse Models

Table 29: Examples SEV⁺/SEV⁻ Explanations for German Credit Dataset. 0DM in the Checking account status column means that checking amount has no Deutsche marks (DM) in it, and 0-200DM means that the checking amount has 0-200 Deutsche marks (DM).

Type	Duration	Credit Amount	Age	Checking account status	Purpose	Other debtors	Housing
Query	18	4380	35	0DM	car(new)	none	own
Reference	19.21	2985.46	36.22	no checking account	radio/television	none	own
SEV ⁺ Explanation	—	—	—	0DM	—	—	—
	—	—	—	—	car(new)	—	—
	—	2985.46	—	—	—	—	—
SEV ⁻ Explanation	—	—	—	no checking account	—	—	—
	—	—	—	—	radio/television	—	—
Query	24	2325	32	0DM	car(new)	none	own
Reference	19.21	2985.46	36.22	no checking account	radio/television	none	own
SEV ⁺ Explanation	—	—	—	0DM	—	—	—
	24	—	—	—	—	—	—
	—	—	—	—	car(new)	—	—
SEV ⁻ Explanation	—	—	—	no checking account	radio/television	—	—
Query	12	585	20	0-200DM	radio/television	co-applicant	rent
Reference	19.21	2985.46	36.22	no checking account	radio/television	none	own
SEV ⁺ Explanation	—	585	—	—	—	—	—
SEV ⁻ Explanation	—	—	—	—	—	none	—
	—	—	36.22	—	—	—	—
	—	—	—	—	—	—	own
Query	36	2323	24	0-200DM	radio/television	none	rent
Reference	19.21	2985.46	36.22	no checking account	radio/television	none	own
SEV ⁺ Explanation	36	—	—	—	—	—	—
SEV ⁻ Explanation	19.21	—	—	—	—	—	—
	—	—	36.22	—	—	—	—