



HAL
open science

Partial Order in Chaos: Consensus on Feature Attributions in the Rashomon Set

Gabriel Laberge, Yann Pequignot, Foutse Khomh, Mario Marchand,
Alexandre Mathieu

► **To cite this version:**

Gabriel Laberge, Yann Pequignot, Foutse Khomh, Mario Marchand, Alexandre Mathieu. Partial Order in Chaos: Consensus on Feature Attributions in the Rashomon Set. 2023. hal-04232199v1

HAL Id: hal-04232199

<https://hal.science/hal-04232199v1>

Preprint submitted on 7 Oct 2023 (v1), last revised 29 Dec 2023 (v3)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Partial Order in Chaos: Consensus on Feature Attributions in the Rashomon Set.

Gabriel Laberge¹

Yann Pequignot²

Alexandre Mathieu²

Foutse Khomh¹

Mario Marchand²

GABRIEL.LABERGE@POLYMTL.CA

YANN.PEQUIGNOT@IID.ULAAVAL.CA

ALEXANDRE.MATHIEU.7@ULAAVAL.CA

FOUTSE.KHOMH@POLYMTL.CA

MARIO.MARCHAND@IFT.ULAAVAL.CA

¹Génie Informatique et Génie Logiciel, Polytechnique Montréal

²Institut Intelligence et Données, Université de Laval à Québec

Abstract

Post-hoc feature attribution methods are progressively being employed to explain decisions of complex machine learning models. Yet, it is possible for practitioners to obtain a diversity of models that provide very different explanations to the same prediction, making it hard to derive insight from them. In this work, instead of aiming at reducing the under-specification of model explanations, we fully embrace it and extract logical statements about feature attributions that are consistent across multiple models with good performance. We show that a **partial** order of feature importance arises from this methodology enabling more nuanced explanations by allowing pairs of features to be incomparable when there is no consensus on their relative importance. We prove that every relation among features present in these partial order also holds in the rankings provided by existing approaches. Finally, we present use cases on three datasets where partial orders allow one to extract knowledge from models despite their under-specification.

Keywords: XAI, Feature Attribution, Under-Specification, Rashomon Set, Uncertainty

1. Introduction

The Machine Learning (ML) framework has proven to be an essential tool in many data-intensive domains such as software engineering, medicine, and cybersecurity (Esteves et al., 2020; Kaieski et al., 2020; Salih et al., 2021). However, the lack of interpretability of complex models is still an important limitation to their applicability. For this reason, various model-agnostic techniques such as LIME (Ribeiro et al., 2016), SHAP (Lundberg and Lee, 2017), and Integrated/Expected Gradient (IG/EG) (Sundararajan et al., 2017; Erion et al., 2021) have recently been developed to provide explanations of model decisions in the form of feature attributions. These attributions are meant to indicate the contribution (positive/negative/null) of individual features toward the model prediction, and their magnitudes (positive/null) can be used for feature selection or to rank features in order of importance. As researchers and practitioners have started to apply these model-agnostic explanations to real-world settings, it has become apparent that they are subject to variability. First, given a fixed model, re-running the explainer can yield different explanations (Visani et al., 2020; Slack et al., 2021; Zhou et al., 2021). Second, retraining the model can induce different explanations for the same decisions (Fel et al., 2021; Shaikhina et al., 2021; Schulz et al., 2021). This phenomenon, known as under-specification, arises when one

employs a rich hypothesis space containing various models that all fit the data while having very different behaviors (D’Amour et al., 2020).

In this work, we focus on uncertainty induced by retraining the model (the model under-specification), while controlling the variability arising from the explainer. Current literature addresses this uncertainty by aggregating explanations from an ensemble of independently trained models. The aggregation is either conducted by averaging the models (Shaikhina et al., 2021), or averaging the feature importance ranks (Schulz et al., 2021). We find that, although these methods provide a single feature attribution to explain all models, it is unclear what statements practitioners are allowed to make with confidence using said feature attribution.

Our characterization of explanations uncertainty departs from the current ones by focusing on *statements* about feature attribution. Our motto in this context of model under-specification is: *only consider statements on which all models with good performance agree*. Concretely, we are going to work with the set of all models with an empirical loss at most ϵ , or equivalently, with all models in the Rashomon Set (Fisher et al., 2019). At a fixed tolerance ϵ , all feature attribution statements on which there is a consensus in the Rashomon Set form **partial** orders, instead of the total orders typically used to rank features. The advantage of partial orders is that they enable safer interpretations by allowing two features to be incomparable. When two features are incomparable, they could be interpreted as “equivalently important” so a practitioner may leverage their background knowledge to select which of the two features is truly more important in the phenomenon being modeled. Here is a brief summary of the contributions of this work:

1. We identify feature attribution statements on which there is a perfect consensus across all models with an empirical loss at most ϵ (*i.e.*, all models in the Rashomon Set). These statements result in partial orders, which differ from the total orders commonly used to visualize feature attributions. Our methodology currently supports the Rashomon Sets of Linear/Additive Regression, Kernel Ridge Regression, and Random Forests Classification/Regression.
2. We prove that if feature i is more important than feature j in our partial orders, then the same relation holds in the total rankings proposed by (Shaikhina et al., 2021; Schulz et al., 2021). This property is crucial given the lack of ground truth in explainability, which restricts quantitative comparisons between competing approaches.
3. We finally present empirical evidence on three open-source datasets that our partial orders are indeed more cautious than total orders, while still conveying important information about the predictions. Each use-case employs a different class of models to better highlight the versatility of our framework.

The rest of the paper is structured as follow: **Section 2** introduces Machine Learning, feature attributions, and the problem of model under-specification, **Section 3** presents a toy-example that serves as the motivation behind our method, **Section 4** discusses our methodology for asserting consensus in the Rashomon Set, while **Sections 5, 6,** and **7** apply said methodology to Linear/Additive models, Kernel Ridge Regression, and Random Forests respectively. Finally, **Section 8** discusses the results and **Section 9** concludes the paper.

2. Background & Related Work

2.1 Machine Learning

Machine Learning (ML) is a programming paradigm where instead of hard-coding logic and rules, we let models adapt their internal logic based on data. We work with an input space $\mathcal{X} \subseteq \mathbb{R}^d$, an output space \mathcal{Y} , an hypothesis space $\mathcal{H} : \mathcal{X} \rightarrow \mathcal{Y}'$, and a loss function $\ell : \mathcal{Y}' \times \mathcal{Y} \rightarrow \mathbb{R}_+$. When the target variable is continuous ($\mathcal{Y} = \mathbb{R}$), the task is called Regression, and when the target is binary ($\mathcal{Y} = \{0, 1\}$), the task is called Classification. We suppose there exists a distribution \mathcal{D} over $\mathcal{X} \times \mathcal{Y}$ from which examples from the dataset $S = \{(\mathbf{x}^{(i)}, y^{(i)})\}_{i=1}^N \sim \mathcal{D}^N$ are sampled iid. The ultimate goal of the ML paradigm is to find a model $h^* \in \operatorname{argmin}_{h \in \mathcal{H}} \mathcal{L}_{\mathcal{D}}(h)$, with minimal population loss $\mathcal{L}_{\mathcal{D}}(h) := \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}}[\ell(h(\mathbf{x}), y)]$. However, since the data-generating distribution \mathcal{D} is unknown, we cannot compute the population loss $\mathcal{L}_{\mathcal{D}}(h)$ and must resort to studying the empirical loss on the dataset S

$$\widehat{\mathcal{L}}_S(h) := \frac{1}{N} \sum_{i=1}^N \ell(h(\mathbf{x}^{(i)}), y^{(i)}), \quad (1)$$

which can be minimized over \mathcal{H} to get an estimate $h_S \in \operatorname{argmin}_{h \in \mathcal{H}} \widehat{\mathcal{L}}_S(h)$ of h^* . In this work, we studied the hypothesis spaces \mathcal{H} of Additive Splines (Hastie et al., 2009, Chapter 5), Kernel Ridge Regression (Mohri et al., 2018, Chapters 6 & 11) and Random Forests (Breiman, 2001a). The two loss functions ℓ that were considered were the squared loss $\ell(y', y) = (y' - y)^2$ for regression and the 0-1 loss $\ell(y', y) = \mathbb{1}(y' \neq y)$ for classification.

2.2 Feature Attribution

The ML paradigm has been successful in tackling tasks where traditional programming methods fail. Still, the lack of transparency of some of the state-of-the-art models such as Random Forests and Multi-Layered Perceptrons prohibits their wide-spread application (Arrieta et al., 2020). To meet this novel challenge, the community of eXplainable Artificial Intelligence (XAI) has recently been growing with the ambition of *explaining* black box models. Formally, given a black box model h and a specific input of interest $\mathbf{x} \in \mathcal{X}$, there is increased interest in capturing the reasons/logic behind the decision $h(\mathbf{x})$. One family of techniques that aim at providing such information are feature attributions, which are vector-valued functionals $\phi : \mathcal{H} \times \mathcal{X} \rightarrow \mathbb{R}^d$ whose vector output represents the contribution of each feature towards the prediction $h(\mathbf{x})$. We are going to focus on feature attributions that are linear w.r.t to the model:

$$\phi(h_1 + \alpha h_2, \mathbf{x}) = \phi(h_1, \mathbf{x}) + \alpha \phi(h_2, \mathbf{x}), \quad (2)$$

for any models $h_1, h_2 \in \mathcal{H}$, and $\alpha \in \mathbb{R}$. Additionally, we will only study feature attributions that are “additive” in the sense that they sum up to a quantity $G(h, \mathbf{x})$ called the gap

$$G(h, \mathbf{x}) := h(\mathbf{x}) - \mathbb{E}_{z \sim \mathcal{B}}[h(z)] = \sum_{i=1}^d \phi_i(h, \mathbf{x}), \quad (3)$$

where \mathcal{B} is a background distribution which is usually a subset of the training data that arises naturally from asking the contrastive question: why is $h(\mathbf{x})$ so high/low compared

to the average prediction on \mathcal{B} ? The magnitude $|\phi_i(h)|$ is called the local importance of feature i .

We now present two linear additive feature attributions methods that have previously been used to understand black box predictions: SHAP (Lundberg and Lee, 2017), and Expected Gradient (EG) (Erion et al., 2021).

2.2.1 SHAPLEY VALUES

The Shapley values are a fundamental concept from cooperative game theory (Shapley, 1953). Letting $[d] = \{1, 2, \dots, d\}$ be the set of all d features, and given a subset $P \subseteq [d]$ of features, we define the replace function $r_P : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ as

$$r_P(\mathbf{z}, \mathbf{x})_i = \begin{cases} x_i & \text{if } i \in P \\ z_i & \text{otherwise.} \end{cases} \quad (4)$$

Moreover, let π be a permutation of $[d]$, $\pi(i)$ be the position of the feature i in π , and $\pi_{:i} = \{j \in [d] : \pi(j) < \pi(i)\}$. The Shapley values, as defined in the library SHAP (Lundberg and Lee, 2017), are the average marginal contributions of specifying the i th feature from the background distribution across all coalitions

$$\phi_i^{\text{SHAP}}(h, \mathbf{x}) = \mathbb{E}_{\substack{\pi \sim \Omega \\ \mathbf{z} \sim \mathcal{B}}} [h(\mathbf{r}_{\pi_{:i} \cup \{i\}}(\mathbf{z}, \mathbf{x})) - h(\mathbf{r}_{\pi_{:i}}(\mathbf{z}, \mathbf{x}))], \quad (5)$$

where Ω is the uniform distribution over all $d!$ permutations of the features. Since Ω encodes 2^d possible permutations, the computational cost of Shapley values is exponential in the number of features, although a method called TreeSHAP was recently developed to reduce the complexity to polynomial assuming the model being explained is an ensemble of decision trees (Lundberg et al., 2020).

2.2.2 INTEGRATED/EXPECTED GRADIENT

The Integrated/Expected Gradient (IG/EG) originates from a different background: cost-sharing in economics. It is also known as the Aumann-Shapley value and has been previously used to compute saliency maps of Convolutional Neural Networks (Sundararajan et al., 2017; Erion et al., 2021). The general definition of EG is

$$\phi_i^{\text{EG}}(h, \mathbf{x}) := \mathbb{E}_{\substack{\mathbf{z} \sim \mathcal{B}, \\ t \sim U(0,1)}} \left[(x_i - z_i) \frac{\partial h}{\partial x_i} \Big|_{t\mathbf{x} + (1-t)\mathbf{z}} \right]. \quad (6)$$

The main idea of this approach is to average the gradient along linear paths between reference inputs sampled from the background and the input \mathbf{x} of interest. When the background distribution degenerates to a single atom at input \mathbf{z} ($\mathcal{B} = \delta_{\mathbf{z}}$), the Expected Gradient falls back the so-called Integrated Gradient.

This work focuses on SHAP and EG feature attributions, but there exist many more post-hoc methods for feature attributions. For instance, LIME (Ribeiro et al., 2016) computes local feature attributions by training a linear model to mimic the behavior of h around \mathbf{x} . Moreover, some post-hoc techniques, such as Permutation Feature Importance (Breiman,

2001a) and SAGE (Covert et al., 2020), extract global feature importance that is not specific to any given input. These other methods are not studied in this work because they do not respect Equations 2 and 3.

2.3 Under-specification and Rashomon Set

The Rashomon Effect (Breiman, 2001b), also known as model under-specification (D’Amour et al., 2020) or model multiplicity (Marx et al., 2020) refers to the observation that there often exists a large diversity of models that fit the data well. This is especially true when one is employing a hypothesis space with a large capacity. Formally, model under-specification can be characterized via the Rashomon Set (Fisher et al., 2019)

Definition 1 (Rashomon Set) *Given a hypothesis space \mathcal{H} , a loss function ℓ , a data set S , and a tolerance threshold $\epsilon > 0$, the Rashomon set is defined as*

$$\mathcal{R}(\mathcal{H}, \epsilon) := \{h \in \mathcal{H} : \widehat{\mathcal{L}}_S(h) \leq \epsilon\}, \quad (7)$$

where we leave the dependence in S and ℓ implicit from the context.

Note that if ϵ is chosen low enough, then the set could simply be empty. Although Rashomon Sets have an appealing and simple interpretation, their computation is intractable unless one is fitting linear models with squared loss. Hence, in general settings, the Rashomon Sets have to be estimated, which can be done by sampling models and keeping the ones with satisfactory performance (Dong and Rudin, 2019; Semenova et al., 2019). However, this method can be time-consuming and requires extensive memory to store hundreds/thousands of models.

Other approaches work implicitly with the Rashomon Set by solving optimization problems over \mathcal{H} under the constraint that $\widehat{\mathcal{L}}_S(h) \leq \epsilon$. By doing so, one can explore the different characteristics of models in the Rashomon Set without ever needing to represent the set explicitly. Such optimization problems have been studied to characterize the under-specification of model predictions (Marx et al., 2020; Coker et al., 2021), and Global Feature Importance (Fisher et al., 2019). However, there are currently no studies that explore the range of possible feature attributions $\phi(h, \mathbf{x})$ across all models from the Rashomon Set.

2.4 Under-Specification of Feature Attributions.

In the words of Leo Breiman (2001b) “*The multiplicity problem and its effect on conclusions drawn from models needs serious attention.*” Indeed, since models are under-specified, so are their interpretations via feature attributions. In practice, this translates to situations where a large set of independently trained models all yield different explanations to the same decision. If our goal is to understand not just one model, but also get insight from the data-generating mechanism, then contradicting explanations of the same decision are problematic. Previous work tackles this uncertainty by aggregating the feature attributions of multiple independently trained models. They both consider an ensemble $E = \{h_k\}_{k=1}^M$ of M models trained independently from a stochastic learning algorithm $h_k \sim \mathcal{A}(S)$. The feature attribution of each of these models are computed $\{\phi(h_k, \mathbf{x})\}_{k=1}^M$ and aggregated. Uncertainty scores are provided in tandem with the aggregated attributions as means to convey how “confident” the attributions are.

For instance, Shaikhina et al. (2021) aggregate feature attributions by explaining the average model and the uncertainty scores are the variances of feature attributions among models. That is, they define the average model $h_E = \frac{1}{M} \sum_{i=1}^M h_k$ and compute its corresponding feature attributions $\phi(h_E, \mathbf{x})$, which the authors show to be equivalent to averaging the feature attributions of each individual model when attribution is a linear functional. The uncertainty score for the attribution of feature i is the variance $\frac{1}{M} \sum_{k=1}^M (\phi_i(h_k, \mathbf{x}) - \phi_i(h_E, \mathbf{x}))^2$.

Additionally, Schulz et al. (2021) obtain an aggregated explanations by averaging the ranks of the feature importance across models $\frac{1}{M} \sum_{k=1}^M \mathbf{r}[|\phi(h_k, \mathbf{x})|]$, where $\mathbf{r} : \mathbb{R}_+^d \rightarrow [d]$ is the rank function that maps each component of a vector to its rank among the other components. The uncertainty score for feature i is the Ordinal consensus metric, which takes values between 0 and 1 and measures the consistency between the rankings. As we shall see in **Section 3**, both of these approaches share the same limitations: it is unclear what statements we can/cannot make with confidence when analyzing the resulting feature attributions. Indeed, they both end up providing a total order of feature importance which suggests that any feature i or more/less important than another feature j , irrespective of the explanation uncertainty. Moreover, the uncertainty scores shown in tandem with the explanations do not provide useful confidence scores to the statement: *feature i is less important than feature j* . Finally, they do not consider the whole Rashomon Set but rather employ ensembles of M independently trained models, which may under-estimate the true under-specification of the ML task.

3. Motivation

We show the limitation of current methods and motivate our own with a toy regression problem. We sampled 1000 4-dimensional points $\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \Sigma)$ where Σ is identity, except for $\Sigma_{1,2} = \Sigma_{2,1} = 0.75$, labelled them via $y := f(\mathbf{x}) + \epsilon$, with $f(\mathbf{x}) = -8 \cos(x_1 - x_2) \cos(x_1 + x_2) + 1.5x_3$ (x_4 is a dummy variable) and ϵ is Gaussian noise with standard deviation $\sigma = 0.1$. We then independently trained 5 Multi-Layered Perceptrons (MLP) with layerwidths=50,20,10 and ReLU activations. All models ended up having test set RMSE between 0.47 and 0.62, while the target had a standard deviation of 4.91. Given this satisfactory level of performance (and the lack of conviction for choosing one model over the other), we decided to study each individual model. We analyzed all their predictions at the input $\mathbf{x} = (\frac{\pi}{2}, \frac{\pi}{2}, \frac{\pi}{2}, \frac{\pi}{2})$ which ranged from 9.05 to 10.05 (the ground truth being $f(\mathbf{x}) = 0.75\pi + 8 \approx 10.36$). To provide insight into the model decisions, Figure 1 (Left) presents the SHAP feature attributions for all five models as blue lines. We see that the various MLPs lead to different interpretations. To make sense of these, we used the two state-of-the-art methods for feature attribution aggregation.

Following Shaikhina et al. (2021), we average the predictions of our 5 models, leading to a single predictor h_E with a test RMSE of 0.42. The resulting SHAP feature attribution is shown as an orange line in Figure 1 (Left). The total order of feature importance for this average model is represented in the first column of Table 1. In particular, this explanation suggests that x_2 is more important than x_1 , which given our knowledge of the symmetry of the ground truth seems somewhat spurious. Indeed, since the underlying data-generating distribution, the target function f , and the point \mathbf{x} to explain are all symmetric w.r.t x_1 and x_2 , an ideal explanation would certainly not support that x_2 is more important than




Figure 1: Left: Feature attributions for the average model h_E (orange line) and each individual model (blue lines). Right: Partial order of local feature importance. There is a directed path from feature x_i to feature x_j if **all good models** agree that feature x_i is more important than x_j .

Feature	Attribution h_E	Variance	Mean rank	Ordinal Consensus
$x_2 = 1.57$	4.65	0.46	2.8	0.87
$x_1 = 1.57$	3.90	0.34	2.2	0.87
$x_3 = 1.57$	2.23	0.10	1.0	1.00
$x_4 = 1.57$	-0.15	0.12	0.0	1.00

Table 1: Aggregated feature attributions and uncertainty scores following previous methods.

x_1 . The uncertainty of the feature attribution is characterized via the variance across the five models, see the second column of Table 1. We note that variance is higher for the attributions of features x_1 and x_2 , suggesting that their contribution toward the output is more uncertain. Still, it is unclear what variance values are low/high enough to label attributions as trustworthy/untrustworthy. Moreover, despite their higher variance, features x_1 and x_2 are consistently more important than features x_3 and x_4 . Thus, we argue that variance provides a pessimistic picture of the insights one can gather from feature attributions of multiple models.

Following Schulz et al. (2021), we averaged the ranks of the SHAP feature importance, see the third column of Table 1. This method again suggests that x_2 is more important than x_1 , which goes against with our knowledge of the data-generating mechanism. Using the Ordinal Consensus as an uncertainty metric (fourth column of Table 1) suggests that all feature importance ranks are confident. Indeed, both x_1 and x_2 have an Ordinal Consensus of 0.87 seeing that there is only a single model for which the ranks of these two features are switched. Nonetheless, looking at Figure 1 (Left), the model that contradicts all others has a test RMSE of 0.512, which is the second best of the whole ensemble. Simply put, this model offers a different but still valid perspective on the data. However, its opinion is “washed out” by the other four models in the computation of the Ordinal Consensus. Hence, we argue that the Ordinal Consensus offers a view of uncertainty that is too optimistic.

As we have just highlighted, the methods of Shaikhina et al. (2021) and Schulz et al. (2021) share the same limitations:

- It is unclear what statements one can/cannot make using these frameworks. For instance, is x_2 really more important than x_1 ? Both approaches return a total order of feature importance, which suggests one statement of relative importance for every pair of features *i.e.* feature i is less/more important than feature j . As we have seen, the uncertainty metrics provided in tandem with the total orders (Variance or Ordinal Consensus) do not help to decide what statements on relative importance are trustworthy.
- It is unclear what is the impact of model performances on the insights provided by these two methods. For instance, the second-best model in the ensemble contradicts all others regarding the relative importance of x_1 and x_2 . However, its opinions are diluted when aggregating all feature attributions.

In light of those takeaways, we decide to focus our method directly on statements about relative feature importance, and whether or not all good models agree on them. For instance, how can we decide if feature x_2 is more important than x_1 ? As noted earlier, one model considers, contrary to the other four, that x_1 is more important than x_2 . Given that this model is as good as any other, we can simply decide to **abstain** from claiming any relation of importance between x_1 and x_2 . In this case, abstention seems indeed a cautious position given the symmetry of the ground truth. Following this logic, for every other pair of features, we check if all five models agree on their relative importance. For instance, all five models agree that x_1 is more important than x_3 . We decide to record this consensus as trustworthy information and we represent it with an arrow from x_1 to x_3 in Figure 1 (Right). Furthermore, we observe that while all five models agree that x_1 , x_2 and x_3 have a positive attribution, this is not the case for x_4 (our dummy variable). Based on this observation, we decide to keep only the variables for which all models agree on the sign and exclude x_4 from our final explanation.

All relations of importance among pairs of features for which there is consensus among all five models actually form a *partial order*, a generalization of total orderings which can be conveniently represented using a Directed Acyclic Graph (DAG) called a Hasse diagram. The partial order of Figure 1 (Right) summarizes our explanation. Note that the partial order suggests that the only relative importance statements we can make are that features x_1 and x_2 are more important than x_3 . These two statements are also supported by the total orders of Shaikhina et al. (2021) and Schulz et al. (2021), a fact that always holds as discussed in **Section 4.5**.

4. Methodology

4.1 Statements on Feature Attributions

Having introduced a basic motivation for considering the consensus among diverse models with good performance, we now present a formal description of the approach. First and foremost, our theory focuses on statements $s : \mathcal{H} \times \mathcal{X} \rightarrow \{0, 1\}$ about feature attributions. Given a performance threshold $\epsilon > 0$, end-users will only be presented statements on which there is a perfect consensus for all models in the Rashomon Set

$$\forall h \in \mathcal{R}(\mathcal{H}, \epsilon) \quad s(h, \mathbf{x}) = 1. \quad (8)$$

We focus on the following statements about feature attributions.

Definition 2 (Positive (Negative) Gap) *We say that the gap $G(h, \mathbf{x})$ is positive (resp. negative) according to h if $G(h, \mathbf{x}) > 0$ (resp. $G(h, \mathbf{x}) < 0$). Formally, the statements take the form $s(h, \mathbf{x}) = \mathbb{1}[G(h, \mathbf{x}) > 0]$ and $s(h, \mathbf{x}) = \mathbb{1}[G(h, \mathbf{x}) < 0]$.*

Before running SHAP or EG, it is primordial to understand the sign of the gap as it is the basis behind the contrastive question we attempt to answer. There may exist instances $\mathbf{x}^{(i)}$ in the data where there is no consensus on the sign of the gap. Therefore, we let

$$\text{SG}(\epsilon) = \{i \in [N] : \forall h_1, h_2 \in \mathcal{R}(\mathcal{H}, \epsilon) \quad \text{sign}[G(h_1, \mathbf{x}^{(i)})] = \text{sign}[G(h_2, \mathbf{x}^{(i)})]\}, \quad (9)$$

be the sets of data instances on which a contrastive question makes sense. Once a contrastive question has been formulated, we can run SHAP or EG and analyze the feature attributions.

Definition 3 (Positive (Negative) Attribution) *We say that feature i has positive (resp. negative) attribution according to h if $\phi_i(h, \mathbf{x}) > 0$ (resp. $\phi_i(h, \mathbf{x}) < 0$). More formally, the statements are $s(h, \mathbf{x}) = \mathbb{1}[\phi_i(h, \mathbf{x}) > 0]$ and $s(h, \mathbf{x}) = \mathbb{1}[\phi_i(h, \mathbf{x}) < 0]$.*

We can now define the sets

$$\text{SA}(\epsilon, \mathbf{x}) = \{i \in [d] : \forall h_1, h_2 \in \mathcal{R}(\mathcal{H}, \epsilon) \quad \text{sign}[\phi_i(h_1, \mathbf{x})] = \text{sign}[\phi_i(h_2, \mathbf{x})]\}, \quad (10)$$

which store the features whose attribution has a consistent sign across all good models. After identifying the sign of the feature attributions, it makes sense to order them according to their magnitude.

Definition 4 (Relative Importance) *We say that feature i is less important than j (or equivalently j is more important than i) according to h if $|\phi_i(h, \mathbf{x})| \leq |\phi_j(h, \mathbf{x})|$. Formally, the statements take the form $s(h, \mathbf{x}) := \mathbb{1}[|\phi_i(h, \mathbf{x})| \leq |\phi_j(h, \mathbf{x})|]$.*

Note that model consensus on relative importance leads to a partial order $\preceq_{\epsilon, \mathbf{x}}$ such that: for all $i, j \in \text{SA}(\epsilon, \mathbf{x})$

$$i \preceq_{\epsilon, \mathbf{x}} j \iff \forall h \in \mathcal{R}(\mathcal{H}, \epsilon) \quad |\phi_i(h, \mathbf{x})| \leq |\phi_j(h, \mathbf{x})|. \quad (11)$$

By requiring a perfect consensus on the Rashomon Set, we guarantee that the order relations will be transitive. Partial orders differ from the common total orders by allowing some pairs of features to be incomparable when there exist two models with conflicting evidence on relative importance.

4.2 Asserting Consensus on Statements

We now show the equivalence between asserting a consensus on the Rashomon Set (*i.e.*, verifying that $\forall h_1, h_2 \in \mathcal{R}(\mathcal{H}, \epsilon)$, $s(h_1, \mathbf{x}) = s(h_2, \mathbf{x}) = 1$) and solving optimization problems.

Definition 5 (Computing Consensus on the Rashomon Set) *Given a tolerance level $\epsilon > 0$, a Rashomon Set $\mathcal{R}(\mathcal{H}, \epsilon)$, and a feature attribution $\phi : \mathcal{H} \times \mathcal{X} \rightarrow \mathbb{R}^d$, consensus on statements are asserted via the following optimization problems.*

1. **Positive (Negative) Gap :** *There is a consensus that the gap $G(h, \mathbf{x})$ is positive (resp. negative) if $\inf_{h \in \mathcal{R}(\mathcal{H}, \epsilon)} G(h, \mathbf{x}) > 0$ (resp. $\sup_{h \in \mathcal{R}(\mathcal{H}, \epsilon)} G(h, \mathbf{x}) < 0$).*
2. **Positive (Negative) Attribution :** *There is a consensus that feature i has a positive (resp. negative) attribution if $\inf_{h \in \mathcal{R}(\mathcal{H}, \epsilon)} \phi_i(h, \mathbf{x}) > 0$ (resp. $\sup_{h \in \mathcal{R}(\mathcal{H}, \epsilon)} \phi_i(h, \mathbf{x}) < 0$).*
3. **Relative Importance :** *Let there be a consensus that the attribution of features i and j have signs s_i and s_j . Under this assumption, the feature importance becomes $|\phi_i(h, \mathbf{x})| = s_i \phi_i(h, \mathbf{x})$ for any $h \in \mathcal{R}(\mathcal{H}, \epsilon)$, and similarly for feature j . Consequently, there is a consensus that i is less important than j if*

$$\sup_{h \in \mathcal{R}(\mathcal{H}, \epsilon)} s_i \phi_i(h, \mathbf{x}) - s_j \phi_j(h, \mathbf{x}) \leq 0.$$

These optimization problems may potentially be intractable depending on the hypothesis set \mathcal{H} and loss functions ℓ . Nonetheless, we will see that they can be solved exactly and efficiently for Linear/Additive Regression, Kernel Ridge Regression, and Random Forests.

4.3 Tune the Error Tolerance

Once we can assert model consensus on a feature attribution statement, we are left with specifying the error tolerance ϵ . By increasing the tolerance, we can explore the under-specification of our ML task via the Rashomon Set. Still, if our tolerance becomes too high, we might reach a point where we always abstain from making statements and our explainability system loses all practical utility. Therefore, tuning the parameter ϵ involved a trade-off: it should be high enough so that we can characterize the under-specification of our problem, but also low enough so that we don't always abstain from making statements. The way we propose to measure the utility of the system is to compute the ratio of statements we make to the ratio of statements we could have made had we known the best-in-class model h^* . Had we known h^* , for each instance in the dataset, we could run SHAP or EG and get d statements on the sign of the attribution and $d(d-1)/2$ statements of relative importance, which leads to $Nd(d+1)/2$ statements total. The utility of our framework is thus defined as

$$u(\epsilon) = \left(N \frac{d(d+1)}{2} \right)^{-1} \sum_{i \in \text{SG}(\epsilon)} |\{(j, k) \in \text{SA}(\epsilon, \mathbf{x}^{(i)})^2 : j \preceq_{\epsilon, \mathbf{x}^{(i)}} k\}|. \quad (12)$$

Simply put, it is the sum of the cardinality of all partial orders that we can define for data instances where the contrastive question is well-defined. We note that $u(\epsilon)$ is a monotonically decreasing function of ϵ . The trade-off between ϵ and $u(\epsilon)$ should be visualized graphically in order to select the tolerance ϵ^* used in practice.

4.4 Full Methodology

We enumerate all the required steps to apply our methodology to a given problem.

1. Program optimization routines to assess feature attribution consensus for any error tolerance level ϵ (cf. **Definition 5**).
2. Apply a linesearch over \mathbb{R}^+ for ϵ and compute $u(\epsilon)$. Choose a preferred value ϵ^* .
3. Only present users with statements on which all models with an empirical loss at most ϵ^* agree. This will implies visualizing feature attributions as partial orders.

4.5 Relation To Prior Work

Prior methods for characterizing the effect of model uncertainty on feature attributions have mainly focused on explaining an ensemble of models $E = \{h_k\}_{k=1}^M$ trained with the same stochastic learning algorithm $h_k \sim \mathcal{A}(S)$ (Shaikhina et al., 2021; Schulz et al., 2021). We go a step further by studying the feature attributions of all models in the Rashomon Set. For this reason, it may not be immediately clear how our method compares to prior work. The following proposition shows that what we propose is a more conservative alternative to both existing methods.

Proposition 6 *Let $\phi(\cdot, \mathbf{x})$ be a linear feature attribution functional, and $E = \{h_k\}_{k=1}^M$ be an ensemble of M models from \mathcal{H} trained with the same stochastic learning algorithm $h_k \sim \mathcal{A}(S)$. Said feature attribution and ensemble will be employed in the methods of (Shaikhina et al., 2021; Schulz et al., 2021). Moreover let $\epsilon \geq \max\{\widehat{\mathcal{L}}_S(h_k)\}_{k=1}^M$ be an error tolerance, and let $\preceq_{\epsilon, \mathbf{x}}$ be the consensus order relation on $SA(\epsilon, \mathbf{x})$ (cf. Equation 11). If the relation $i \preceq_{\epsilon, \mathbf{x}} j$ holds, we have that i is less important than j in the two total orders of prior work (Shaikhina et al., 2021; Schulz et al., 2021).*

This proposition is key as it implies that our framework will not provide users with statements that are not supported by existing approaches. In a way, all we do is abstain from making statements whose uncertainty is highest. We think this is an important property to have because, unlike model predictions, there are no ground truths for feature attributions. For example, a practitioner can apply multiple aggregation mechanisms to model predictions (Arithmetic Mean, Geometric Mean, Majority Vote etc.) and compare the resulting test set performances using the target y as ground truth. However, when aggregating feature attributions using different schemes, there is no metric for what feature importance ranking is the best, or closest to ground truth. This is one of the major challenges currently faced by the explainability community. Still, since our framework only highlights statements supported by existing approaches, we eliminate the need for quantitative comparisons.

5. Application to Linear/Additive Regression

5.1 Rashomon Set

One of the simplest hypothesis spaces \mathcal{H} is the set of Additive models of the form $h(\mathbf{x}) = \omega_0 + \sum_{j=1}^d h_j(x_j)$ where each function h_j only depends on the feature x_j . We note that linear models are a subset of additive models with $h_j(x_j) = \omega_j x_j$. The particularity of Linear/Additive models is that the output is the sum of the contributions of d functions h_j which each only depend on one input feature. When employing such models, the contribution of each individual feature toward the output is readily available, which is why additive models are advertised as being transparent. To fit an additive model, one must find a way to represent the univariate functions h_j . A first method is to represent each of the functions non-parametrically via a sum of univariate decision trees. This scheme is what is currently done in the `ExplainableBoostingRegressor` of the `InterpreML` Python library (Nori et al., 2019) for instance. The parametric alternative is to define a basis $\{h_{jk}\}_{k=1}^{M_j}$ along each dimension j (for example using Splines) and represent the additive model using linear combinations of these basis functions (Hastie et al., 2009, Chapter 5)

$$h_{\omega}(\mathbf{x}) = \omega_0 + \underbrace{\sum_{j=1}^d \sum_{k=1}^{M_j} \omega_{jk} h_{jk}(x_j)}_{h_j(x_j)}, \quad (13)$$

where

$$\boldsymbol{\omega} := [\omega_0, \underbrace{\omega_{11}, \omega_{12}, \dots, \omega_{1, M_1}}_{\text{feature 1}}, \underbrace{\omega_{21}, \omega_{22}, \dots, \omega_{2, M_2}}_{\text{feature 2}}, \dots, \underbrace{\omega_{d1}, \omega_{d2}, \dots, \omega_{d, M_d}}_{\text{feature d}}]^T.$$

In that case, be letting \mathbf{H} be the $N \times (1 + \sum_{j=1}^d M_j)$ matrix whose i th row is

$$[1, \underbrace{h_{11}(\mathbf{x}^{(i)}), h_{12}(\mathbf{x}^{(i)}), \dots, h_{1M_1}(\mathbf{x}^{(i)})}_{\text{feature 1}}, \dots, \underbrace{h_{d1}(\mathbf{x}^{(i)}), h_{d2}(\mathbf{x}^{(i)}), \dots, h_{dM_d}(\mathbf{x}^{(i)})}_{\text{feature d}}],$$

the empirical loss minimizer takes the familiar form

$$\boldsymbol{\omega}_S = (\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T \mathbf{y}. \quad (14)$$

Definition 7 (Rashomon Set for Parametric Additive Regression) *Let \mathcal{H} be the set of Parametric Additive Regression models (cf Equation 13), ℓ be the squared loss, S be a dataset of size N , and $\boldsymbol{\omega}_S = \operatorname{argmin}_{h \in \mathcal{H}} \widehat{\mathcal{L}}_S(h)$ be the least-square estimate. If one uses the performance threshold $\epsilon \geq \widehat{\mathcal{L}}_S(\boldsymbol{\omega}_S)$, then the Rashomon set $\mathcal{R}(\mathcal{H}, \epsilon)$ contains all models h_{ω} s.t.*

$$(\boldsymbol{\omega} - \boldsymbol{\omega}_S)^T \frac{\mathbf{H}^T \mathbf{H}}{N} (\boldsymbol{\omega} - \boldsymbol{\omega}_S) \leq \epsilon - \widehat{\mathcal{L}}_S(\boldsymbol{\omega}_S), \quad (15)$$

where \mathbf{H} is the $N \times (1 + \sum_{j=1}^d M_j)$ augmented feature matrix. We see that the Rashomon Set is isomorphic to an ellipsoid. Moreover, if we let $\epsilon < \widehat{\mathcal{L}}_S(\boldsymbol{\omega}_S)$, then the Rashomon Set is empty.

This result is a simple generalization of the Rashomon of Linear Regression models derived in Semenova et al. (2019) to Parametric Additive Regression models.

5.2 Asserting Model Consensus

Beyond the fact that the Rashomon Set has an analytical expression, additive models are the only types of models where the notion of feature attribution is not ambiguous. For instance, running SHAP and EG on a Parametric Additive model while taking the whole dataset S as the background yields the same answer

$$\begin{aligned} \phi_j^{\text{SHAP}}(h, \mathbf{x}) &= \phi_j^{\text{EG}}(h, \mathbf{x}) = h_j(x_j) - \frac{1}{N} \sum_{i=1}^N h_j(x_j^{(i)}) \\ &= \sum_{k=1}^{M_j} \omega_{jk} \left(h_{jk}(x_j) - \frac{1}{N} \sum_{i=1}^N h_{jk}(x_j^{(i)}) \right), \end{aligned} \quad (16)$$

which is a linear function of the weights ω_{jk} . We have seen previously in Definition 5 that asserting the consensus on feature attribution statements amounts to optimization problems that are linear with respect to the attributions. Therefore, asserting a consensus on the Rashomon Set of Parametric Additive models requires maximizing/minimizing a linear function on an ellipsoid

$$\begin{aligned} \min/\max_{\boldsymbol{\omega}} \quad & \mathbf{a}^T \boldsymbol{\omega} \\ \text{with} \quad & (\boldsymbol{\omega} - \boldsymbol{\omega}_S)^T \mathbf{A} (\boldsymbol{\omega} - \boldsymbol{\omega}_S) \leq \epsilon - \widehat{\mathcal{L}}_S(\boldsymbol{\omega}_S), \end{aligned} \quad (17)$$

with $\mathbf{A} := \frac{\mathbf{H}^T \mathbf{H}}{N}$ and assuming $\epsilon \geq \widehat{\mathcal{L}}_S(\boldsymbol{\omega}_S)$. This optimization problem has an analytical solution that can be computed rapidly using the Cholesky decomposition $\mathbf{A} = \mathbf{A}^{\frac{1}{2}} \mathbf{A}^{\frac{1}{2}T}$. The optimal values of Equation 17 are

$$\pm \sqrt{\epsilon - \widehat{\mathcal{L}}_S(\boldsymbol{\omega}_S)} \|\mathbf{a}'\| + \mathbf{a}^T \boldsymbol{\omega}_S, \quad (18)$$

where $\mathbf{a}' = \mathbf{A}^{-\frac{1}{2}} \mathbf{a}$ and $\mathbf{A}^{-\frac{1}{2}} := (\mathbf{A}^{\frac{1}{2}})^{-1}$, see **Appendix A.2** for more details. This result is a generalization of Theorem 4 from Coker et al. (2021) to Parametric Additive models and arbitrary linear functionals of the weights $\mathbf{a}^T \boldsymbol{\omega}$. We deduce from Equation 18 that the minimum and maximum values of any linear functional evaluated on the Rashomon Set are a deviation of $\sqrt{\epsilon - \widehat{\mathcal{L}}_S(\boldsymbol{\omega}_S)} \|\mathbf{a}'\|$ from $\mathbf{a}^T \boldsymbol{\omega}_S$ the value of the functional evaluated on the least-square. Since the deviation is an explicit function of the tolerance ϵ , a consensus on feature attribution statements can be asserted at any tolerance level.

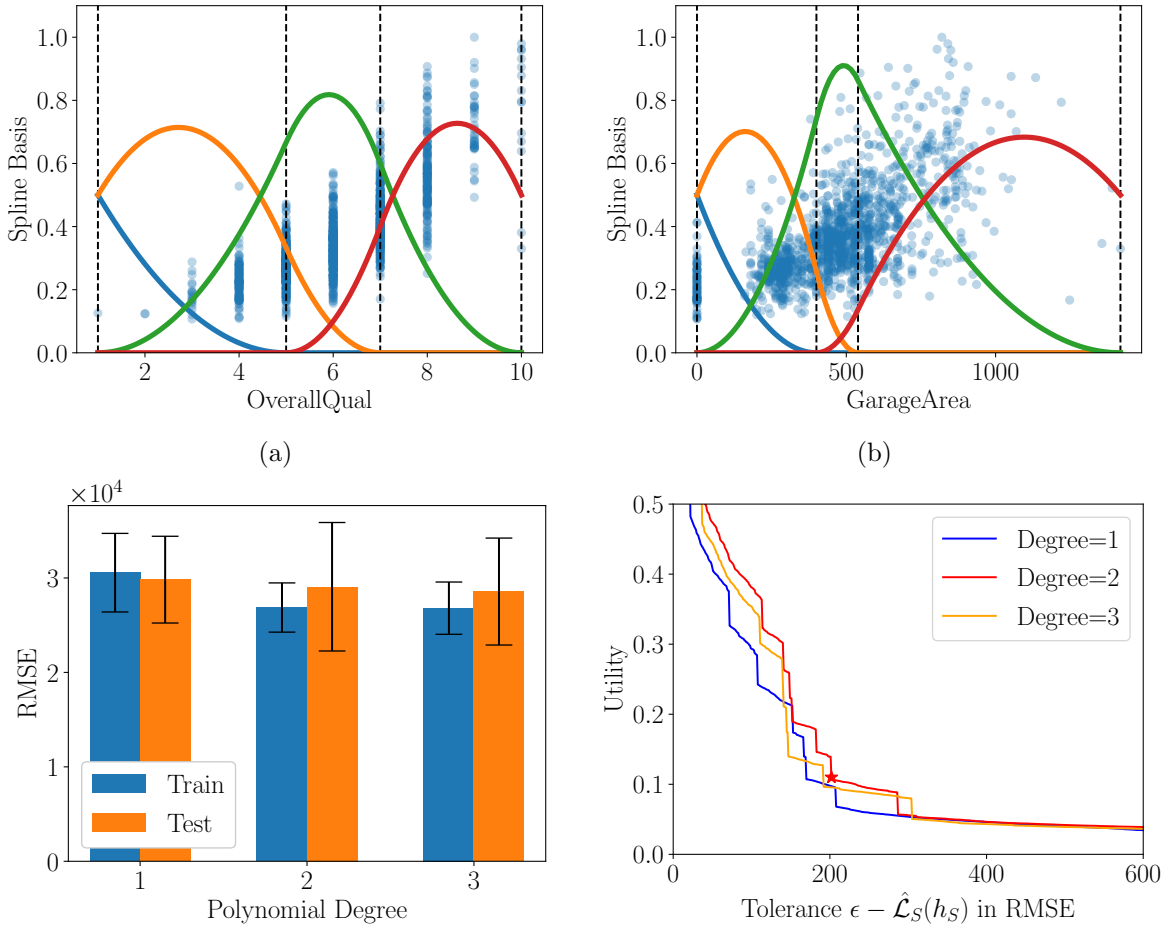
5.3 House Price Prediction

The Houses regression dataset, available on Kaggle¹, consists of predicting the selling price in USD of 1446 houses based on 79 numerical and categorical features. For simplicity, only numerical features were selected and, when feature pairs had a high Spearman correlation (>0.55), only one feature was kept. Correlated features were removed because we suspect they would aggravate under-specification of feature attributions since competing models would rely on different subsets of correlated features. Moreover, we removed the two time-related features `YearSold` and `YearRemodAdd`, since we are only interested in the physical properties of the houses. We were finally left with 15 features total.

We decided to train an additive regression model with uni-variate spline bases. The first modeling step was deciding which features to represent with spline bases and which to keep linear (*i.e.* $h_j(x_j) = \omega_j x_j$). We regressed the target on each feature individually using a depth-3 decision tree and selected the four features with the highest R^2 score as candidates for spline modeling. The reasoning behind this approach is that additive models only consider the main effects of the features and hence, following Occam’s razor, our modeling efforts focus on features that are highly predictive on their own. We ended up using spline bases with 4 knots on the features `LotArea`, `OverallQual`, `1stFlrSF`, and `GarageArea`, see Figure 2(a) and (b) for examples. The function `SplineTransformer` of the Scikit-Learn Python library (Pedregosa et al., 2011) was used to generate the bases. We left the polynomial degree of splines as a hyper-parameter to tune between values 1, 2, and 3. Figure 2(c) presents the test set empirical loss of ω_S for the three hypothesis spaces, along with Asymptotic-Gaussian confidence intervals of the population loss $\mathcal{L}_{\mathcal{D}}(\omega_S)$. We see that none of these polynomial degrees results in a test error that is significantly better than the others. This is an opportunity to compute the error-utility trade-off for all three hypothesis spaces and see the effects of model complexity on the under-specification of feature attributions.

The result of such an analysis is presented in Figure 2(d). We show the excess tolerance $\epsilon - \widehat{\mathcal{L}}_S(\omega_S)$ instead of the absolute tolerance ϵ to ease comparisons between the three hypothesis spaces. Contrary to our intuition, we see that the smallest hypothesis space (degree 1) exhibits the steepest decrease in utility w.r.t increases in tolerance. We suspect this is because the simpler models perform worst on the training set and for the same excess of error relative to the least-square, the Rashomon Set contains worst models than other hypothesis spaces. We also note that polynomials of degree 3 have a steeper decrease of utility compared to degree 2, which could be attributed to their higher capacity, leading to more contradictions between competing models. As splines of degree 2 show the most gradual decreases in utility, we select this hypothesis space and fix $\epsilon^* = 200 + \widehat{\mathcal{L}}_S(\omega_S)$ to coincide with an “elbow” in Figure 2(d). We illustrate the selected value with a red star. The reason we search for an “elbow” is that we do not want our statements to cease to hold if we were to slightly increase our tolerance to error. Bear in mind that an extra tolerance of 200 USD in error is negligible considering the models are predicting prices of houses that range from 50k to 500k USD.

1. <https://www.kaggle.com/c/house-prices-advanced-regression-techniques>



(c) Performance of the Spline Additive models for three different polynomial degrees. Error bars are due to the finite test set size.

(d) Decrease of the utility as the excess tolerance increases. A red star indicates the value of tolerance selected.

Figure 2

Having selected the tolerance ϵ^* and hypothesis space \mathcal{H}_2 , we can extract statements on feature attributions for which there is a consensus on $\mathcal{R}(\mathcal{H}_2, \epsilon^*)$. The first instance we wish to explain is the house with the smallest price in the dataset: 52K USD. The predictions for this instance ranged from 40k to 83k across the Rashomon Set. Figure 3 (Top) shows the feature attribution on this instance and the partial order that summarizes all the statements we make with confidence. We observe that features **OverallQual=4** (quality of materials and finish of the house from a scale of 1 to 10) and **1stFlrSF=very small** have maximal importance when explaining the drop in price relative to the mean. Also, we note that **GarageArea=0** is not part of the partial order despite clearly having a negative attribution according to the bar chart. This is because, as indicated by the error bars, the models in the Rashomon Set do not agree on the sign of this feature’s attribution. This highlights the danger of explaining a single model and the nuance introduced by explaining a set of good models.

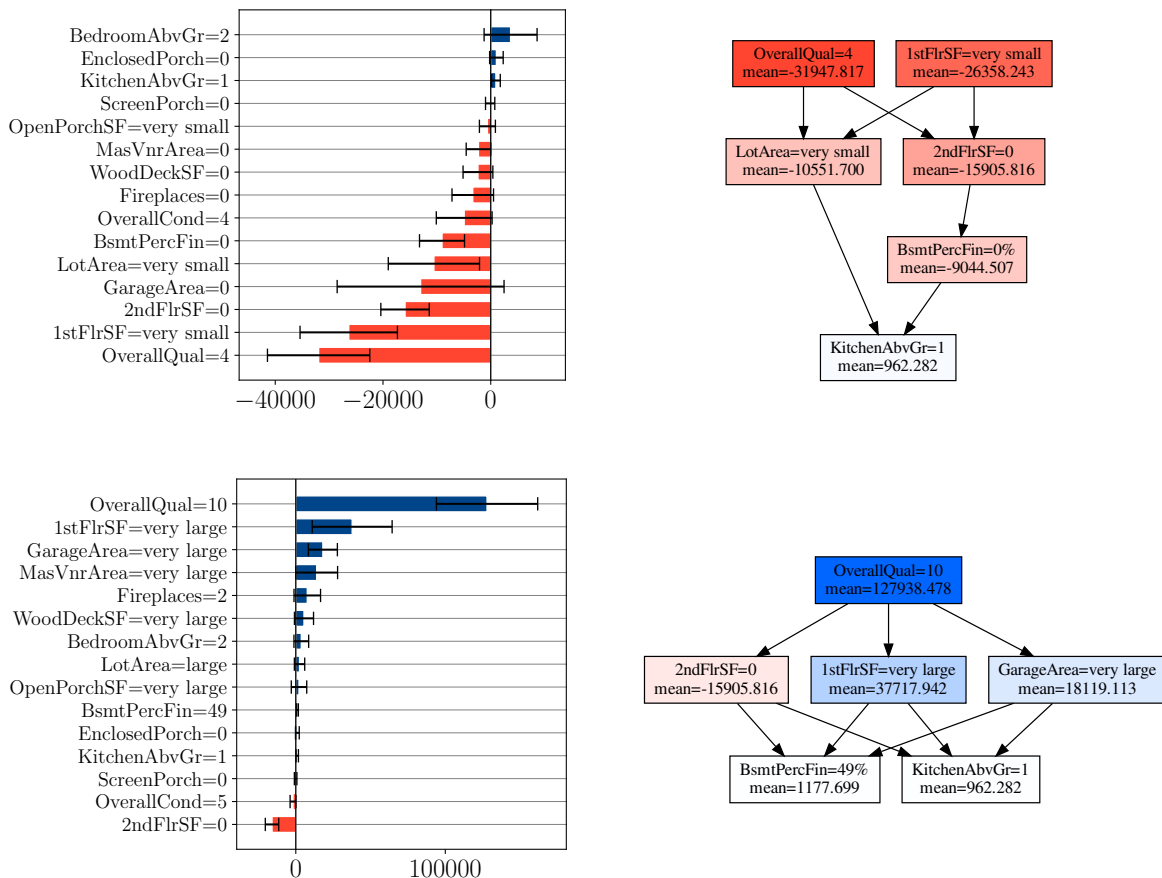


Figure 3: Feature attributions on two houses instances. (Top) House with the smallest price. (Bottom) House with the third largest price.

The second instance we explain is the house with the third highest price: 466k USD. The predictions on this instance ranged from 350k to 415k across the Rashomon Set. Figure 3 (Bottom) shows the feature attribution on this instance and the partial order that summarizes all the statements on which the models agree. We see that feature **OverallQual=10** has maximum importance to explain the increase in price predictions. This is a strong statement because it means that all models with satisfactory performance put this feature at the top 1. We also observe that three other features **2ndFlrSF=0**, **1stFlrSF=very large**, and **GarageArea=very large** are all important for the prediction but are all mutually incomparable. Indeed, we abstain from claiming that any one of these features is more important than the others. By allowing features to be incomparable and putting them on the same vertical position in the partial order, we make the explanation more nuanced and allow end-users to focus their attention on whichever feature they prefer. Note that, unlike Figure 3 (Top), the feature **GarageArea** does appear in the partial order for this instance since all models agree that this feature has a positive attribution. This means that whether or not there is a consensus on the sign of a feature’s attribution can depend on the instance being explained.

6. Application to Kernel Ridge Regression

6.1 Rashomon Set

Let $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}_+$ be a symmetric positive definite kernel. Then such a kernel induces a functional space called a Reproducing-Kernel-Hilbert-Space (RKHS), which is actually the completion of the Pre-Hilbert space (Mohri et al., 2018)

$$\mathcal{H}_k := \left\{ h(\mathbf{x}) = \sum_{j=1}^R \alpha_j k(\mathbf{x}, \mathbf{r}^{(j)}) \text{ for } R \in \mathbb{N}, \boldsymbol{\alpha} \in \mathbb{R}^R, \mathbf{r}^{(j)} \in \mathcal{X} \right\} \quad (19)$$

endowed with the scalar product

$$\langle k(\cdot, \mathbf{r}^{(i)}), k(\cdot, \mathbf{r}^{(j)}) \rangle_{\mathcal{H}_k} = k(\mathbf{r}^{(i)}, \mathbf{r}^{(j)}), \quad (20)$$

from which the terminology ‘‘Reproducing-Kernel’’ arises. The space \mathcal{H}_k is impossibly large since it requires specifying any integer R and any R reference inputs $\mathbf{r}^{(j)}$. For simplicity, we will fix the R reference inputs in advance and store them in a dictionary $D := \{\mathbf{r}^{(j)}\}_{j=1}^R$. We will then employ the much simpler space

$$\mathcal{H}_k^D := \left\{ h_{\boldsymbol{\alpha}}(\mathbf{x}) = \sum_{j=1}^R \alpha_j k(\mathbf{x}, \mathbf{r}^{(j)}) \text{ for } \boldsymbol{\alpha} \in \mathbb{R}^R \right\} \quad (21)$$

s.t. $\mathcal{H}_k^D \subset \mathcal{H}_k$ as was done in (Fisher et al., 2019). Since these spaces are still considerably expressive, it is common to apply regularization when learning models from them. From the Rashomon perspective, this implies studying the Rashomon Set

$$\mathcal{R}(\mathcal{H}_k^D, \epsilon) := \left\{ h_{\boldsymbol{\alpha}} \in \mathcal{H}_k^D : \widehat{\mathcal{L}}_D(h_{\boldsymbol{\alpha}}) + \lambda \|h_{\boldsymbol{\alpha}}\|^2 \leq \epsilon \right\}, \quad (22)$$

where $\lambda > 0$ is a regularization hyper-parameter that is fine-tuned by cross-validation and $\|h_{\boldsymbol{\alpha}}\|^2 = \langle h_{\boldsymbol{\alpha}}, h_{\boldsymbol{\alpha}} \rangle_{\mathcal{H}_k} = \sum_{i,j=1}^R \alpha_i \alpha_j k(\mathbf{r}^{(i)}, \mathbf{r}^{(j)})$ is the functional norm induced by the scalar product on \mathcal{H}_k . We let $\mathbf{K} \in \mathbb{R}^{R \times R}$ be the symmetric positive definite matrix of kernel evaluations on the dictionary $\mathbf{K}[i, j] = k(\mathbf{r}^{(i)}, \mathbf{r}^{(j)})$. The regularized least-square solution is

$$\boldsymbol{\alpha}_D = (\mathbf{K} + \lambda R \mathbf{I})^{-1} \mathbf{y}. \quad (23)$$

Given this notation, we can present the Rashomon Set of Kernel Ridge Regression.

Definition 8 (Rashomon Set for Kernel Ridge Regression) *Let \mathcal{H}_k^D be the space induced by the kernel k and dictionary D , ℓ be the squared loss, $\lambda > 0$ be a regularization hyper-parameter, and $\boldsymbol{\alpha}_D$ be the solution of the regularized least-square. If one uses the performance threshold $\epsilon \geq \widehat{\mathcal{L}}_D(h_{\boldsymbol{\alpha}_D}) + \lambda \|h_{\boldsymbol{\alpha}_D}\|^2$, then the Rashomon set $\mathcal{R}(\mathcal{H}_k^D, \epsilon)$ contains all models $h_{\boldsymbol{\alpha}}$ s.t.*

$$(\boldsymbol{\alpha} - \boldsymbol{\alpha}_D)^T (\mathbf{K}/R + \lambda \mathbf{I}) \mathbf{K} (\boldsymbol{\alpha} - \boldsymbol{\alpha}_D) \leq \epsilon - \widehat{\mathcal{L}}_D(h_{\boldsymbol{\alpha}_D}) - \lambda \|h_{\boldsymbol{\alpha}_D}\|^2. \quad (24)$$

We see that the Rashomon Set is isomorphic to an ellipsoid in \mathbb{R}^R .

The proof is mutatis mutandis like the proof for Ridge Regression in Semenova et al. (2019) but with Kernel Ridge instead.

6.2 Asserting Model Consensus

Unlike the previous section, the model h_{α} is no longer additive, and hence there is no universal way to assign a score ϕ_i to each input feature when explaining a model decision. Hence we must rely on either SHAP or Integrated Gradient, which are two principled approaches for computing said scores. Because the exponential burden of Shapley values has not yet been solved for kernel methods, SHAP was not used and we instead employed the Integrated Gradient with a single baseline input \mathbf{z} . Henceforth, assuming the kernel is continuous and has continuous partial derivatives ($k \in \mathbb{C}^1(\mathcal{X} \times \mathcal{X})$), we compute the IG as follows.


$$\begin{aligned} \phi_i(h_{\alpha}, \mathbf{x}, \mathbf{z}) &= (x_i - z_i) \int_0^1 \frac{\partial h_{\alpha}}{\partial x_i} \Big|_{t\mathbf{x}+(1-t)\mathbf{z}} dt \\ &= \sum_{j=1}^R \alpha_j \underbrace{\left[(x_i - z_i) \int_0^1 \frac{\partial k(\cdot, \mathbf{r}^{(j)})}{\partial x_i} \Big|_{t\mathbf{x}+(1-t)\mathbf{z}} dt \right]}_{\phi_{ij}} = \sum_{j=1}^R \alpha_j \phi_{ij}, \end{aligned} \quad (25)$$

which is a linear function of the coefficients α . Consequently, asserting a consensus on IG feature attributions will again amount to optimizing a linear function over an ellipsoid so we can leverage results from the previous section. The only additional step required for Kernel Ridge is to pre-compute the path integrals ϕ_{ij} with common quadrature methods.

6.3 Criminal Recidivism Prediction

COMPAS is a proprietary model currently employed in the United States to predict the risk of recidivism from individuals that were recently arrested. These risks are encoded as scores going from 1 (low-risk) to 10 (high-risk). The use of this automated tool in the justice system is driven by the promise of providing objective information to judges based on empirical data, thus circumventing human biases. Still, the strong reliance of models on historical data means they can reproduce/perpetuate past injustices. To test such claims, ProPublica has collected several thousands of COMPAS scores from 2013-2014 in the Florida Broward County (Larson et al., 2016). In the resulting article, several pairs of Caucasian and African-American defendants are presented along with their COMPAS scores, the former often being lower than the latter despite the Caucasian defendant having a longer criminal history. These examples of pairs along with the subsequent analysis from the article seem to imply that the proprietary model depends on race. However, the methodology of ProPublica has been heavily criticized alongside the claim that COMPAS depends explicitly on race (Rudin et al., 2018). Hence, there may exist alternative explanations besides race for the discrepancy between scores, so it is pertinent to study the feature attributions of the whole Rashomon Set of reasonable models for predicting COMPAS scores.

To analyze the dependencies of risk scores on the various features, we repeated the experiments of Fisher et al. (2019) where a Kernel Ridge Regression model was fitted directly on the 1-10 scores from the ProPublica dataset. The same features were employed while adding two additional ones related to juvenile misdemeanors and felonies. We utilized Polynomial Kernel $k(\mathbf{x}, \mathbf{x}') = (\gamma \langle \mathbf{x}, \mathbf{x}' \rangle + 1)^p$ with degree $p = 3$ and the Gaussian Kernel $k(\mathbf{x}, \mathbf{x}') = \exp(-\gamma \|\mathbf{x} - \mathbf{x}'\|^2)$. The kernel scale hyper-parameter γ and the regularization



(a) Hyper-parameter tuning for γ and λ and Gaussian Kernels. Top horizontal line shows the error of the predictor returning the mean, while the bottom line shows the error of a Random Forest with default hyper-parameters. (b) Convergence of the Gap Error relative to the number of steps in the quadrature. We observe a second-order convergence. Hence, augmenting the number of steps by a factor 10 reduced the error by a factor 100.

Figure 4

Input	Name	Score	Race	Age	Priors	Charge
\mathbf{x}	Robert Cannon	6	African-American	22	0	Misdemeanor
\mathbf{z}	James Rivelli	3	Caucasian	54	3	Felony

Table 2: Comparison of the COMPAS scores of two individuals.

factor λ were fine-tuned with 5-fold cross-validation, see the results for Gaussian Kernels in Figure 4 (a). Similar results were obtained with Polynomial Kernels. The resulting test set RMSE was 2.11 for Gaussian Kernels and 2.12 for Polynomial Kernels. We note that the performances are worst than Fisher et al. (2019) because, unlike them, we predict the recidivism risk scores and not the risk scores for **violent** recidivism, which seem easier to predict. The reason we did not predict the same COMPAS scores was that we wanted to study the scores which were actually discussed in the ProPublica article.

After fitting the models, we identified a pair of Caucasian/African-American individuals who were highlighted in the ProPublica piece and applied our explainability framework on them. More specifically, we compared Robert Cannon and James Rivelli, see Table 2. James Rivelli is a 54-year-old Caucasian man who was arrested for shoplifting. Despite having a criminal record with three priors, he was assigned a low COMPAS score. In contrast, Robert Cannon, a 22-year-old African-American charged with petit theft, was assigned a high risk of recidivism. Letting Robert be the input of \mathbf{x} and James be the input \mathbf{z} , we observe that the differences in scores are also present for the Kernel Ridge models: $h(\mathbf{x}) = 4.9$ and $h(\mathbf{z}) = 2.5$ for Gaussian Kernels, and $h(\mathbf{x}) = 4.9$ and $h(\mathbf{z}) = 2.5$ for Polynomial Kernels. Therefore, we have a prediction gap $G(h, \mathbf{x}) = h(\mathbf{x}) - h(\mathbf{z})$ that is positive.

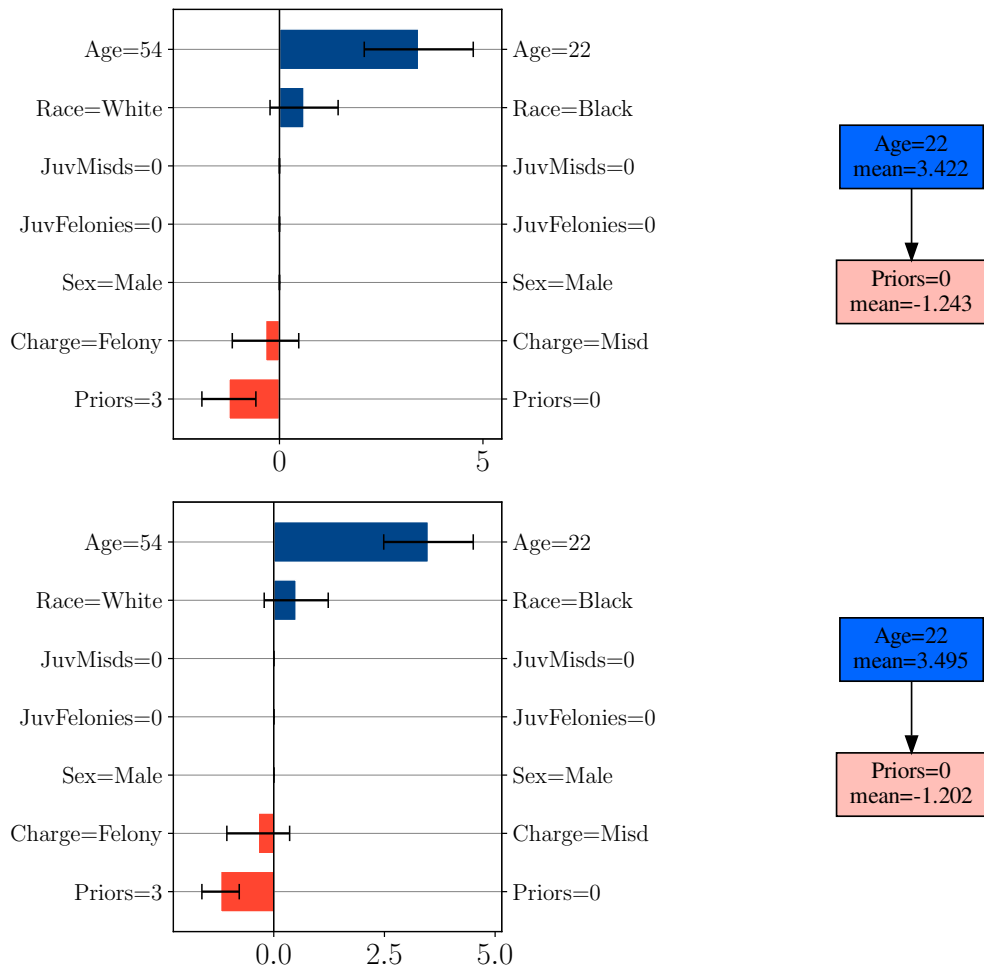


Figure 5: Feature attributions comparing Robert Cannon to James Rivelli. (Top) Gaussian Kernels. (Bottom) Polynomial Kernels. The features on the left of the bar charts represent James while the values on the right represent Robert.

Given the historical racism in the United States, it is very tempting to look at these two individuals and say that Robert Cannon is predicted to have a higher risk “because of his race”. Still, there may exist a diversity of alternative explanations for this discrepancy, which we can study by exploring the Rashomon Set of our Kernel Ridge models. To obtain feature attributions, the Integrated Gradient was employed using Robert as the input of interest \mathbf{x} and James as the reference input \mathbf{z} . Since computing the IG feature attributions requires estimating the integrals of Equation 25 with quadratures, we ended up with estimates $\hat{\phi}(h, \mathbf{x})$ of the real attributions $\phi(h, \mathbf{x})$. We characterized the estimation error of this discretization by reporting the Gap Error

$$\left| \sum_{i=1}^d \hat{\phi}_i(h, \mathbf{x}) - G(h, \mathbf{x}) \right|, \quad (26)$$

and used it as a proxy of how well $\widehat{\phi}(h, \mathbf{x})$ approximates $\phi(h, \mathbf{x})$. By simplicity, the Trapezoid quadrature was implemented, see Figure 4 (b) for the convergence of the Gap Error as the number of steps in the quadrature increases. We note that, as expected, the quadrature converges to the second order. For the remainder of the analysis, we have employed quadratures with 1000 steps.

Seeing that the regularized loss $\widehat{\mathcal{L}}_D(h_{\alpha}) + \lambda \|h_{\alpha}\|^2$ does not have interpretable units, we decided to set the error tolerance ϵ to a multiple of the regularized least-squared solution. More specifically, we set $\epsilon = 1.02 \times [\widehat{\mathcal{L}}_D(h_{\alpha_D}) + \lambda \|h_{\alpha_D}\|^2] \approx 4.27$ (an increase of 2% of the minimum loss value). Unlike the previous section, we did not conduct a full line search over the possible values of ϵ . Rather, by setting it to a small value, we only wish to prove the existence of competing models that disagree on their explanation for the discrepancy between James and Robert.

Figure 5 presents the feature attributions across the Rashomon Sets $\mathcal{R}(\mathcal{H}_k^D, 4.27)$ of Gaussian and Polynomial Kernels. Since the results are consistent across the two types of Kernels, we will only discuss Gaussian Kernels. Inspecting the top bar plot, we see that, according to the regularized least-square solution α_D plotted as the blue/red bars, the features **Age=22** and **Race=Black** have positive attributions while the features **Charge=Misdemeanor** and **Prior=0** have negative attributions. This suggests that one of the possible explanations for the high risk of Robert is racial discrimination toward African-Americans. However, when we additionally consider the opinion of models with slightly worst performance on the training data, some of our previous statements on feature attribution cease to hold. Importantly, there exist competing models that yield a negative attribution to the feature **Race=Black**, and therefore there are reasonable explanations for the disparity between Robert and James, that do not rely on Robert being African-American. Even when considering the whole Rashomon Set $\mathcal{R}(\mathcal{H}_k^D, 4.27)$, there remain statements on which models reach consensus. Notably, the attribution of the feature **Age=22** remains positive and has maximum importance. The only way for age to be a negligible factor in explaining the difference in scores between Robert and James would be to consider models with worst performances than what we considered.

These observations are concordant with previous work of Rudin et al. (2018) which hypothesizes that COMPAS depends strongly on age and (at most) weakly on race. Nonetheless, our analysis must not be taken as absolute facts about the proprietary model COMPAS. This is because we do not have access to the model and we are surrogating it with Kernel Ridge models fitted on 7 features. The original COMPAS model, on the contrary, takes 137 different factors into consideration to produce a score (Rudin et al., 2018). Our analysis is more of a proof of concept that our explainability framework can make sense of the feature attributions of competing models and that it can highlight the diversity of explanations for the discrepancies between two individuals.

7. Application to Random Forests

7.1 Rashomon Set

A Random Forest (RF) is an ensemble of independently trained decision trees whose predictions are averaged to yield the final predictions Breiman (2001a). To increase diversity, each tree is trained on a different bootstrap sample of the original dataset and each inner split is done on a random subset of features. We let s represent the random seed encoding all pseudo-random processes in the training of a single tree t_s . If \mathcal{S} is a distribution over all possible seeds on a computer, the theoretical definition of a RF is

$$h(\mathbf{x}) = \mathbb{E}_{s \sim \mathcal{S}}[t_s(\mathbf{x})]. \quad (27)$$

For simplicity, we will assume that the set of possible seeds is finite and of size M . Then, a reasonable choice of distribution over seeds is the uniform over M seeds *i.e.* $\mathcal{S} = U(\{1, 2, \dots, M\})$. The number of seeds M will be fixed to 1000 for reasons we will see later. In practice, the expectation $\mathbb{E}_{s \sim \mathcal{S}}$ has to be approximated using Monte-Carlo sampling. Given $m < M$, we sub-sample m seeds uniformly at random $S \sim \mathcal{S}^m$, and return the sample average as our estimate of the RF

$$h_S(\mathbf{x}) = \frac{1}{m} \sum_{s \in S} t_s(\mathbf{x}). \quad (28)$$

By the weak law of large numbers, the estimated RF should converge to the true RF (cf. Equation 27) as m increases. Since sampling m seeds out of M with/without replacement assigns a non-zero probability to any subset of m seeds, we conceptualize the space of all possible RFs as the collection of all subsets of trees.

Definition 9 Given a large set $\mathcal{T} = \{t_s\}_{s=1}^M$ of M trees trained with M different seeds, the set of all possible RFs of m trees is

$$\mathcal{H}_m := \left\{ \frac{1}{m} \sum_{t \in T} t : T \subseteq \mathcal{T} \text{ and } |T| = m \right\}, \quad (29)$$

i.e. all averages of subsets of m trees from \mathcal{T} . Moreover we define $\mathcal{H}_m := \cup_{k=m}^M \mathcal{H}_k$ as all RFs with least m trees.

Figure 6 illustrates an example of space \mathcal{H}_m which highlights their combinatoric nature. Note the monotonic relation $m < m' \Rightarrow \mathcal{H}_m \supset \mathcal{H}_{m'}$: meaning that decreasing m increases the number of models considered. We can interpret \mathcal{H}_1 as the set of all possible Random Forests that can ever appear in practice and so we aim at characterizing its Rashomon Set $\mathcal{R}(\mathcal{H}_1, \epsilon)$. Such a Rashomon Set cannot be explicitly computed because of the exponential size of \mathcal{H}_1 : ($|\mathcal{H}_1| = 2^M - 1$). Still, we will see that studying the space \mathcal{H}_m for a carefully chosen m can help us characterize a large subset of the Rashomon Set.

A desirable property of the spaces \mathcal{H}_m is that optimizing a linear functional over them is tractable, as highlighted by the following proposition.

Proposition 10 Let $\mathcal{T} := \{t_s\}_{s=1}^M$ be a set of M trees, \mathcal{H}_m be the set of all subsets of at least m trees from \mathcal{T} , and $\phi : \mathcal{H}_m \rightarrow \mathbb{R}$ be a linear functional, then $\min_{h \in \mathcal{H}_m} \phi(h)$ amounts to averaging the m smallest values of $\phi(t_s)$ for $s = 1, 2, \dots, M$.




Figure 6: Example of the space \mathcal{H}_2 representing all the possible groupings of 2 decision trees out of $M = 4$.

The proof of this proposition is presented in **Appendix A.3**. Now, by letting the linear functionals ϕ be the SHAP feature attributions, we can use TreeSHAP (Lundberg et al., 2020) to compute $\phi(t_s)$ for all seeds $s = 1, 2, \dots, M$ given that M is not too large.

At this point, we assume that the desired tolerance on error ϵ has been fixed and so we wish to identify a value $m(\epsilon)$ that guarantees that $\mathcal{H}_{m(\epsilon)} \subseteq \mathcal{R}(\mathcal{H}_1, \epsilon)$. This value $m(\epsilon)$ should be as small as possible so that the space $\mathcal{H}_{m(\epsilon)}$ is as large as possible. With this goal in mind, we let ℓ be the 0-1 loss or the squared loss and define

$$\epsilon^+(m) = \frac{1}{N} \sum_{i=1}^N \max_{h \in \mathcal{H}_m} \ell(h(\mathbf{x}^{(i)}), y^{(i)}), \quad (30)$$

which can be computed efficiently at any m by leveraging **Proposition 10** to identify the two most extreme predictions at any given $\mathbf{x}^{(i)}$. Note that $\max_{h \in \mathcal{H}_m} \widehat{\mathcal{L}}_S(h) \leq \epsilon^+(m) \forall m$ so that $\epsilon^+(m)$ is a pessimistic measure of the performance of all RFs containing at least m trees. Given an absolute tolerance ϵ on the empirical loss, we search for the smallest number of trees m we can keep while ensuring that $\epsilon^+(m) \leq \epsilon$

$$\begin{aligned} m(\epsilon) &:= \min_{m=1,2,\dots,M} m \\ \text{s.t. } &\epsilon^+(m) \leq \epsilon. \end{aligned} \quad (31)$$

The intuition behind the computation of $m(\epsilon)$ is presented in Figure 7. Since setting $m = m(\epsilon)$ guarantees that $\max_{h \in \mathcal{H}_m} \widehat{\mathcal{L}}_S(h) \leq \epsilon^+(m) \leq \epsilon$, we have $\mathcal{H}_{m(\epsilon)} \subseteq \mathcal{R}(\mathcal{H}_1, \epsilon)$. Hence, we are going to employ $\mathcal{H}_{m(\epsilon)}$ as a conservative estimate of the Rashomon Set over which we can efficiently optimize linear functionals such as model predictions or the SHAP attributions.

We end this subsection by discussing a toy example of how to compute $\epsilon^+(m)$. We designed a regression task where the input follows a standard Gaussian and the output y is a quadratic function of x plus some noise of amplitude 0.9. A total of $M = 1000$ different seed values were used to independently generate 1000 decision trees. Figure 8 (a) shows the


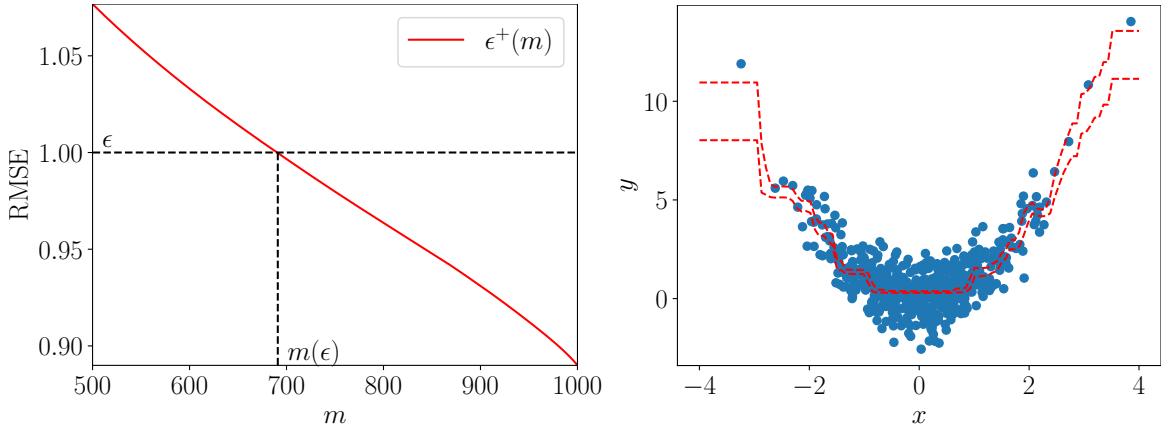


Figure 7: Intuition of how we choose m based on the error tolerance ϵ .



(a) Performance bound $\epsilon^+(m)$. Given a RMSE tolerance of $\epsilon = 1$, the smallest m we can safely consider is $m(\epsilon) = 691$. (b) Toy regression data. The min-max predictions over the hypothesis space \mathcal{H}_{691} are shown as red lines.

Figure 8

upper bound $\epsilon^+(m)$ of any RF containing at least m trees. Given a threshold on the RMSE of $\epsilon = 1$, the smallest m we can safely consider is $m(\epsilon) = 691$. Hence, we suggest to employ the set \mathcal{H}_{691} as a subset of $\mathcal{R}(1, S, \mathcal{H}_1)$. Figure 8 (b) presents the minimum and maximum predictions $\min_{h \in \mathcal{H}_{691}} h(x)$ and $\max_{h \in \mathcal{H}_{691}} h(x)$ at various values of x . We see that the min-max prediction intervals are wider in low data density regions near the boundaries. This means that there are more disagreements between individual trees on these points. Such an observation makes sense because each tree is fitted on a bootstrap sample of the dataset and therefore some trees have never seen the boundary points.

7.2 Asserting Model Consensus

Given an error tolerance ϵ , we set m to $m(\epsilon)$ and assert the consensus on \mathcal{H}_m : via optimization problems (cf. **Definition 5**) that we solve efficiently with **Proposition 10**. For example, to compute $\min_{h \in \mathcal{H}_m} \phi_i(h, \mathbf{x})$, we calculate the vector of feature attributions of all trees $[\phi_i(t_1, \mathbf{x}), \phi_i(t_2, \mathbf{x}), \dots, \phi_i(t_M, \mathbf{x})]^T$ with TreeSHAP, then we sort it and average its m smallest values. The overall complexity of this procedure w.r.t M is $\mathcal{O}(M \log M)$, which is why we must keep M at a large but reasonable value.


7.3 Income Prediction

The Census Income dataset (aka Adults) also available on the UCI repository² contains the census data of 48,842 individuals collected in 1994. It consists of a binary classification task with the goal of predicting whether or not an individual makes more ($y = 1$) or less ($y = 0$) than 50k USD per year based on 14 attributes. Out of all these features, we removed `fnlwgt` because we do not fully understand what it represents and `native-country` because it is a categorical feature with very high cardinality. We were finally left with five numerical features and seven categorical ones that were one-hot-encoded. After encoding, we were left with a data matrix of 41 columns. The data was split into train and test sets with ratios 0.8 and 0.2 respectively. The training set was used to obtain the set \mathcal{T} of M iid trees while the test set was used to compute the upper bound $\epsilon^+(m)$ on error (cf. Equation 30). Given an error tolerance ϵ , the upper bound $\epsilon^+(m)$ was used to set $m(\epsilon)$ the minimum number of trees in our RFs (cf. Equation 31). Since Adult is a classification task, the 0-1 loss was employed.

For the model, we utilized Scikit-Learn’s `RandomForestClassifiers` whose hyperparameters were tuned with a 100 steps random search and 5-fold cross-validation. Then, we trained $M = 1000$ trees in order to generate a set \mathcal{T} . The training was actually repeated 5 times so that we ended up with 5 distinct sets of 1000 trees \mathcal{T}_i with $i = 1, 2, \dots, 5$. We do not expect practitioners to fit several sets \mathcal{T}_i when applying our methodology. This was done to verify our assumption that \mathcal{T} is representative of all trees trained with bootstrapped data and random splits.

The model outputs $h(\mathbf{x}) \in [0, 1]$ must be interpreted as estimates of the conditional probabilities of y given \mathbf{x} and not as hard 0/1 predictions. Therefore, feature attributions should sum up to a difference in conditional probabilities. Since we are employing tree-based models, we decided to compute feature attributions with TreeSHAP. In fact, seeing that categorical features were one-hot-encoded, which is not supported in the TreeSHAP implementation of the SHAP library, we used the Partition-TreeSHAP algorithm described in (Laberge and Pequignot, 2022). The feature attribution requires a background distribution \mathcal{B} to serve as a reference and we decided to use the empirical distribution of the whole training set. Still, given the considerable size of the Adult dataset, we had to sub-sample B instances from the training set and use them to estimate Shapley values. So, we ended up explaining the models with estimates $\hat{\phi}$ rather than ground-truths ϕ . A proxy of the error made by sub-sampling is the Gap Error presented in Equation 26. Figure 9 (a) presents distributions of the Gap Errors when explaining 2000 test instances using various background sizes. We note that across five reruns, the errors start to stagnate at about

2. <https://archive.ics.uci.edu/ml/datasets/adult>



(a) Gap Errors for increasing sizes of the background dataset. Each box represents one of the five reruns.

(b) Upper bound ϵ^+ as a function of the utility of the explainability system. Each curve is a different rerun.

Figure 9

50 background instances. As extra security, we decided to run TreeSHAP with $B = 500$ background instances as our best estimates. This led to a one-hour runtime for explaining $M = 1000$ decision trees on 2000 test instances. Figure 9 (b) presents the resulting trade-offs between utility and test error. We observe that the five curves are very similar which suggests that fitting $M = 1000$ trees can be representative of all trees possibly generated for RFs. The red star indicates the utility 0.54 and error 0.17 that were selected for the sequel. We accentuate that an upper bound of 0.17 on the test miss-classification rate is an increase of approximately 0.03 relative to the best test error, which we view as a reasonable tolerance. The associated $m(\epsilon)$ was 798 meaning that the following results consider the consensus on all Random Forests with at least 798 trees. We now discuss three instances that were explained with our framework.

The first instance is an individual who makes more than 50k per year and whose predictions ranged from 0.76 to 0.86 across \mathcal{H}_{798} . Figure 10 (Top) illustrates this person’s feature attribution and the resulting partial order that encodes the statements on which there is a consensus among all models. We observe that the feature `educational-num=very large` has maximum positive importance for understanding why this individual has higher-than-average predictions. This means that all RFs of at least 798 trees put this feature in the top 1. The second most important feature is `marital-status=Married`. These two rankings coincide with the ranking we would get by averaging all 1000 trees as indicated by the bar chart. Still, the fact that these rankings remain intact even when considering the Rashomon Set rather than a single model gives us more confidence in their robustness. The three next features `age, occupation, hours-per-week` all appear at the same vertical position and are all mutually incomparable. This is a type of information that cannot be conveyed when simply averaging all models or averaging the ranks of their feature importance, as the work of (Shaikhina et al., 2021; Schulz et al., 2021) would recommend.

The second instance is a person who makes more than 50k and whose predictions range from 0.30 to 0.51. It is fair to say that this person is located near the decision boundary of

the task. Figure 10 (Middle) shows how our framework would explain the model predictions. We focus on the two features `capital-gain=0` and `workclass=Self-Employed` which both have a negative attribution according to the average model. Looking at the error bars on the bar chart, we observe that the model uncertainty is higher for `workclass` than with `capital-gain`. This means that there is more agreement among RFs that `capital-gain=0` reduced the model output. For `workclass=Self-Employed`, the model uncertainty is so high that the min-max interval crosses the origin, which implies the existence of RFs with satisfactory performance that yield a positive attribution to this feature. Our framework identified this ambiguity and hence removed the feature `workclass` from the partial order despite it having a negative attribution according to the average model.

The final instance we explain is a female individual who makes more than 50k but whose predictions ranged from 0.038 to 0.196. A hypothesis for why this person would have such low outputs is gender bias, which we can investigate in Figure 10 (Bottom). Looking at the bar chart, we note that the min-max interval of the feature `gender=female` is highly concentrated toward the origin. This means that no RF considers gender important for this individual. This is a stronger result than stating that the average model does not rely on gender for this prediction. In fact, a company being sued by this individual could argue in court that the probability any of their models rely on gender is null. Now, looking at the partial order, we can get some insights into the model predictions. We see there are many features with maximal importance and different signs for their attributions. Hence, there are many strong opposing forces pushing the decision up and down. For instance, being married drives the prediction up while not having a long education and hours per week drives the prediction down.

We end this section by resuming this experiment and our conclusions. We have fitted 1000 decision trees on Adults and have studied the Rashomon Set of all RFs containing at least 798 of the 1000 trees. By only presenting end-users with statements on which a diversity of good models agree, we have derived more nuanced insights from feature attributions. A by-product of this methodology is that it is often the same features that come up in the partial orders: `educational-num`, `marital-status`, `age`, `occupation`, `hours-per-week`, `relationship`, and `capital-gain`. In a way, since these features have the most robust feature attributions, our framework naturally focuses on them. For instance, the feature `education` is less robust to model perturbations because `education=HS-grad` leads to ambiguity regarding the sign of the feature attribution, see the bar charts of Figure 10 (Middle and Bottom). Although the partial orders often contain the features, the sign of their attribution and their relative order relations vary between instances which means that local feature attributions are not redundant and can convey information relevant to a specific individual.

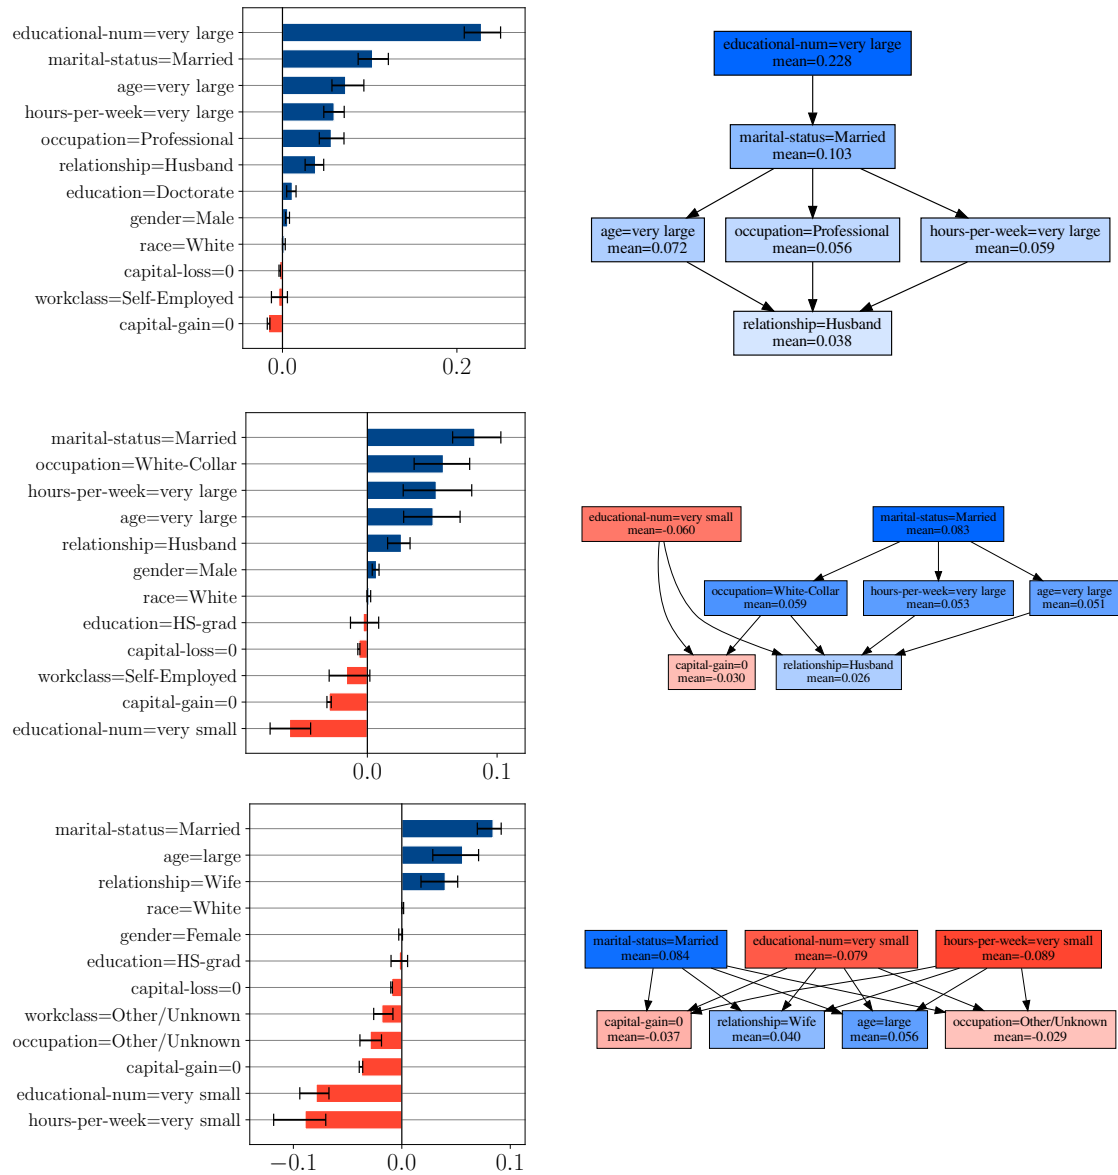


Figure 10: Feature attributions on three individuals. (Top) Person with a high prediction. (Middle) Individual near the decision boundary. (Bottom) Person with a low prediction.

8. Discussion

As suggested by our experiments, model under-specification has an important impact on feature attributions on real data, and taking into account this uncertainty seems necessary to derive reliable insight from machine learning models. Our conservative approach only retains the information on features attributions on which all models agree and still succeeds in finding partial order in this chaos. This in itself is an important observation because one could have expected the partial orders to be trivial and contain no interesting structure (no arrows).

The principal limitation of our approach is that we are currently restricted to Linear/Additive Regression, Kernel Ridge Regression, and Random Forests. It is therefore primordial to extend our work to other models, especially to more Classifiers. We envision using techniques from previous work to sub-sample models from the Rashomon Sets of Logistic Regression and Decision Trees (Dong and Rudin, 2019; Semenova et al., 2019). Once such ensembles of good models are computed, applying our framework would be as simple as sorting all models w.r.t their empirical error and asserting the consensus at any level of tolerance ϵ .

Still, there may also exist hypothesis spaces whose Rashomon Set is too large to be realistically estimated, for instance, Neural Networks. Moreover, the cost of training/explaining multiple models may be too high for practitioners to see any benefit. A potential solution to derive careful conclusions from these large models would be to employ only a few models, but train them in a way that ensures they are as diverse as possible. This application of our framework is left as future work as it involves unique and novel challenges regarding the training of Neural Networks.

The main characteristic of our approach is that we require a perfect consensus among all good models. However, when employing our methodology with a finite ensemble of models, one may wonder why not also consider statements on which a majority of models agree (or at least 90% of the models agree). As a more extreme example, a practitioner may have 1000 models and 999 of these models state something while a single one states the opposite. Our approach would abstain from making any statement in that case, which may seem unnecessarily strict. The technical justification for requiring a perfect consensus is to guarantee the transitivity of the order relations. This property is crucial for the interpretability of the feature orderings. Although there exist post-processing techniques to ensure the transitivity of order relations derived from the consensus from at least $\tau\%$ of the models (Cheng et al., 2010), it becomes hard to interpret why certain arrows are present/absent from the Hasse Diagram. Our diagrams, on the other hand, remain simple to interpret: an arrow means that all models agree and the absence of an arrow means that two models disagree. The philosophical justification for perfect consensus is that, given that the error threshold ϵ was fixed at a value that represents satisfactory performances, the single model that disagrees with the rest is still a good model, and its mere existence is enough to put into question the claim supported by the other 999 models. If this single model had worst performance than the 999 others, slightly reducing the error tolerance would remove this model from the Rashomon Set and we would reach a consensus.

Speaking of tuning the error tolerance ϵ , similar to prior work, we apply a line search over a realistic range of values and visually inspect the effect of under-specification on conclusions

drawn from models. Nonetheless, it is not clear what is a good value for ϵ , although we argue this is a limitation shared by most studies on the Rashomon Set (D’Amour et al., 2020; Semenova et al., 2019; Coker et al., 2021). We think the most promising result tackling this limitation is Proposition 7 from Fisher et al. (2019). There, the error threshold ϵ was linked to the probability that the min-max interval of a statistic over the Rashomon Set captures the statistic for the unknown best-in-class model h^* . By extending Proposition 7, one could choose ϵ large enough so that the probability of making statements about h^* using our framework becomes close to 1. This promising research direction is left as future work.

9. Conclusion

In this work, we propose a new approach to explanations in the context of model uncertainty. Rather than considering the mean attributions or the mean rank, we identify properties and relations of feature attributions that are consistent across a set of models with good performance. These logical statements about feature attribution naturally lead to a partial order of feature importance, which we show can provide more nuanced explanations than the more common total orders based on mean attributions. As such, we believe that our work opens a new perspective on post-hoc explanations in the context of model uncertainty.

In future work, we intend to study more Classifiers (Logistic Regression, Decision Trees, Neural Networks) and other local/global post-hoc explanations (LIME, Permutation Importance, SAGE). Moreover, we shall apply our methodology to more practical settings, especially those where there are clear *actionable* features on which a human subject is able to act upon. We hope that in these scenarios, the nuance introduced by partial orders will prove most beneficial.

Acknowledgements

The authors wish to thank the DEEL project CRDPJ 537462-18 funded by the National Science and Engineering Research Council of Canada (NSERC) and the Consortium for Research and Innovation in Aerospace in Québec (CRIAQ), together with its industrial partners Thales Canada inc, Bell Textron Canada Limited, CAE inc and Bombardier inc.

References

- Alejandro Barredo Arrieta, Natalia Díaz-Rodríguez, Javier Del Ser, Adrien Bennetot, Siham Tabik, Alberto Barbado, Salvador García, Sergio Gil-López, Daniel Molina, Richard Benjamins, et al. Explainable artificial intelligence (xai): Concepts, taxonomies, opportunities and challenges toward responsible ai. *Information Fusion*, 58:82–115, 2020.
- Leo Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001a.
- Leo Breiman. Statistical modeling: The two cultures (with comments and a rejoinder by the author). *Statistical science*, 16(3):199–231, 2001b.
- Weiwei Cheng, Michaël Rademaker, Bernard De Baets, and Eyke Hüllermeier. Predicting partial orders: ranking with abstention. In *Joint European conference on machine learning and knowledge discovery in databases*, pages 215–230. Springer, 2010.
- Beau Coker, Cynthia Rudin, and Gary King. A theory of statistical inference for ensuring the robustness of scientific results. *Management Science*, 67(10):6174–6197, 2021.
- Ian Covert, Scott M Lundberg, and Su-In Lee. Understanding global feature contributions with additive importance measures. *Advances in Neural Information Processing Systems*, 33:17212–17223, 2020.
- Alexander D’Amour, Katherine Heller, Dan Moldovan, Ben Adlam, Babak Alipanahi, Alex Beutel, Christina Chen, Jonathan Deaton, Jacob Eisenstein, Matthew D Hoffman, et al. Underspecification presents challenges for credibility in modern machine learning. *arXiv preprint arXiv:2011.03395*, 2020.
- Jiayun Dong and Cynthia Rudin. Variable importance clouds: A way to explore variable importance for the set of good models. *arXiv preprint arXiv:1901.03209*, 2019.
- Gabriel Erion, Joseph D Janizek, Pascal Sturmfels, Scott M Lundberg, and Su-In Lee. Improving performance of deep learning models with axiomatic attribution priors and expected gradients. *Nature Machine Intelligence*, pages 1–12, 2021.
- Geanderson Esteves, Eduardo Figueiredo, Adriano Veloso, Markos Viggiano, and Nivio Ziviani. Understanding machine learning software defect predictions. *Automated Software Engineering*, 27(3):369–392, 2020.
- Thomas Fel, David Vigouroux, Rémi Cadène, and Thomas Serre. How good is your explanation? algorithmic stability measures to assess the quality of explanations for deep neural networks. 2021.
- Aaron Fisher, Cynthia Rudin, and Francesca Dominici. All models are wrong, but many are useful: Learning a variable’s importance by studying an entire class of prediction models simultaneously. *Journal of Machine Learning Research*, 20(177):1–81, 2019.
- Trevor Hastie, Robert Tibshirani, Jerome H Friedman, and Jerome H Friedman. *The elements of statistical learning: data mining, inference, and prediction*, volume 2. Springer, 2009.

- Naira Kaieski, Cristiano Andre da Costa, Rodrigo da Rosa Righi, Priscila Schmidt Lora, and Bjoern Eskofier. Application of artificial intelligence methods in vital signs analysis of hospitalized patients: A systematic literature review. *Applied Soft Computing*, page 106612, 2020.
- Gabriel Laberge and Yann Pequignot. Understanding interventional treeshap: How and why it works. *arXiv preprint arXiv:2209.15123*, 2022.
- Jeff Larson, Surya Mattu, Lauren Kirchner, and Julia Angwin. How we analyzed the compas recidivism algorithm. *ProPublica*, May 2016. URL <https://www.propublica.org/article/how-we-analyzed-the-compas-recidivism-algorithm>.
- Scott M Lundberg and Su-In Lee. A unified approach to interpreting model predictions. In *Proceedings of the 31st international conference on neural information processing systems*, pages 4768–4777, 2017.
- Scott M Lundberg, Gabriel Erion, Hugh Chen, Alex DeGrave, Jordan M Prutkin, Bala Nair, Ronit Katz, Jonathan Himmelfarb, Nisha Bansal, and Su-In Lee. From local explanations to global understanding with explainable ai for trees. *Nature machine intelligence*, 2(1): 56–67, 2020.
- Charles Marx, Flavio Calmon, and Berk Ustun. Predictive multiplicity in classification. In *International Conference on Machine Learning*, pages 6765–6774. PMLR, 2020.
- Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. *Foundations of machine learning*. MIT press, 2018.
- Harsha Nori, Samuel Jenkins, Paul Koch, and Rich Caruana. Interpretml: A unified framework for machine learning interpretability. *arXiv preprint arXiv:1909.09223*, 2019.
- 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.
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. ” why should i trust you?” explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 1135–1144, 2016.
- Cynthia Rudin, Caroline Wang, and Beau Coker. The age of secrecy and unfairness in recidivism prediction. *arXiv preprint arXiv:1811.00731*, 2018.
- Azar Salih, Subhi T Zeebaree, Sadeeq Ameen, Ahmed Alkhyyat, and Hnan M Shukur. A survey on the role of artificial intelligence, machine learning and deep learning for cybersecurity attack detection. In *2021 7th International Engineering Conference “Research & Innovation amid Global Pandemic”(IEC)*, pages 61–66. IEEE, 2021.
- Jonas Schulz, Rafael Poyiadzi, and Raul Santos-Rodriguez. Uncertainty quantification of surrogate explanations: an ordinal consensus approach. *arXiv preprint arXiv:2111.09121*, 2021.

- Lesia Semenova, Cynthia Rudin, and Ronald Parr. A study in rashomon curves and volumes: A new perspective on generalization and model simplicity in machine learning. *arXiv preprint arXiv:1908.01755*, 2019.
- Torgyn Shaikhina, Umang Bhatt, Roxanne Zhang, Konstantinos Georgatzis, Alice Xiang, and Adrian Weller. Effects of uncertainty on the quality of feature importance explanations. *AAAI Workshop on Explainable Agency in Artificial Intelligence*, 2021.
- Lloyd S Shapley. A value for n-person games. *Contributions to the Theory of Games*, pages 307–317, 1953.
- Dylan Slack, Anna Hilgard, Sameer Singh, and Himabindu Lakkaraju. Reliable post hoc explanations: Modeling uncertainty in explainability. *Advances in Neural Information Processing Systems*, 34, 2021.
- Mukund Sundararajan, Ankur Taly, and Qiqi Yan. Axiomatic attribution for deep networks. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 3319–3328. PMLR, 2017.
- Giorgio Visani, Enrico Bagli, Federico Chesani, Alessandro Poluzzi, and Davide Capuzzo. Statistical stability indices for lime: obtaining reliable explanations for machine learning models. *Journal of the Operational Research Society*, pages 1–11, 2020.
- Zhengze Zhou, Giles Hooker, and Fei Wang. S-lime: Stabilized-lime for model explanation. *arXiv preprint arXiv:2106.07875*, 2021.

A. Proofs

A.1 Relation to Prior Work

Proposition 11 (Proposition 6) *Let $\phi(\cdot, \mathbf{x})$ be a linear feature attribution functional, and $E = \{h_k\}_{k=1}^M$ be an ensemble of M models from \mathcal{H} trained with the same stochastic learning algorithm $h_k \sim \mathcal{A}(S)$. Said feature attribution and ensemble will be employed in the methods of (Shaikhina et al., 2021; Schulz et al., 2021). Moreover let $\epsilon \geq \max\{\widehat{\mathcal{L}}_S(h_k)\}_{k=1}^M$ be an error tolerance, and let $\preceq_{\epsilon, \mathbf{x}}$ be the consensus order relation on $SA(\epsilon, \mathbf{x})$ (cf. Equation 11). If the relation $i \preceq_{\epsilon, \mathbf{x}} j$ holds, we have that i is less important than j in the two total orders of prior work (Shaikhina et al., 2021; Schulz et al., 2021).*

Proof We first note that, since $i, j \in SA(\epsilon, \mathbf{x})$, there is a consensus across the Rashomon Set that these features attributions have sign s_i and s_j respectively. As a reminder, this simplifies the expression of the feature importance : $\forall h \in \mathcal{R}(\mathcal{H}, \epsilon) \quad |\phi_i(h, \mathbf{x})| = s_i \phi_i(h, \mathbf{x})$. Additionally, our assumption that $\epsilon \geq \max\{\widehat{\mathcal{L}}_S(h_k)\}_{k=1}^M$, guarantees that $E \subseteq \mathcal{R}(\mathcal{H}, \epsilon)$. We now prove that the order relation $i \preceq_{\epsilon, \mathbf{x}} j$ is present in the two rankings from the literature.

Shaikhina et al. (2021) order features using the mean rank $\frac{1}{M} \sum_{k=1}^M \mathbf{r}[|\phi(h_k, \mathbf{x})|]$, where $\mathbf{r} : \mathbb{R}_+^d \rightarrow [d]$ is the rank function. By the definition, for any model h , we have $|\phi_i(h, \mathbf{x})| \leq |\phi_j(h, \mathbf{x})| \iff r_i[|\phi(h, \mathbf{x})|] \leq r_j[|\phi(h, \mathbf{x})|]$. Therefore,

$$\begin{aligned} i \preceq_{\epsilon, \mathbf{x}} j &\Rightarrow \forall h \in \mathcal{R}(\mathcal{H}, \epsilon) \quad |\phi_i(h, \mathbf{x})| \leq |\phi_j(h, \mathbf{x})| \\ &\Rightarrow \forall h \in E \quad |\phi_i(h, \mathbf{x})| \leq |\phi_j(h, \mathbf{x})| \\ &\Rightarrow \forall h \in E \quad r_i[|\phi(h, \mathbf{x})|] \leq r_j[|\phi(h, \mathbf{x})|] \\ &\Rightarrow \frac{1}{M} \sum_{k=1}^M r_i[|\phi(h_k, \mathbf{x})|] \leq \frac{1}{M} \sum_{k=1}^M r_j[|\phi(h_k, \mathbf{x})|], \end{aligned}$$

which means that the order relation is also supported by the mean ranks.

Schulz et al. (2021) compute the average model $h_E = \frac{1}{M} \sum_{k=1}^M h_k$ and rank features according to their importance for this model $|\phi(h_E, \mathbf{x})|$. For any $i, j \in SA(\epsilon, \mathbf{x})$, we deduce

$$\begin{aligned} i \preceq_{\epsilon, \mathbf{x}} j &\Rightarrow \forall h \in \mathcal{R}(\mathcal{H}, \epsilon) \quad |\phi_i(h, \mathbf{x})| \leq |\phi_j(h, \mathbf{x})| \\ &\Rightarrow \forall h \in \mathcal{R}(\mathcal{H}, \epsilon) \quad s_i \phi_i(h, \mathbf{x}) \leq s_j \phi_j(h, \mathbf{x}) \\ &\Rightarrow \forall h \in E \quad s_i \phi_i(h, \mathbf{x}) \leq s_j \phi_j(h, \mathbf{x}) \\ &\Rightarrow \frac{1}{M} \sum_{k=1}^M s_i \phi_i(h_k, \mathbf{x}) \leq \frac{1}{M} \sum_{k=1}^M s_j \phi_j(h_k, \mathbf{x}) \\ &\Rightarrow s_i \phi_i(h_E, \mathbf{x}) \leq s_j \phi_j(h_E, \mathbf{x}) \quad (\text{By Linearity of } \phi) \\ &\Rightarrow |\phi_i(h_E, \mathbf{x})| \leq |\phi_j(h_E, \mathbf{x})|, \quad (\text{By Linearity of } \phi, s_i = \text{sign}[\phi_i(h_E, \mathbf{x})]) \end{aligned}$$

thus proving that the order relation is also present when explaining the average model. \blacksquare




Figure 11: Mapping an ellipsoid to the unit sphere.

A.2 Optimization over a Ellipsoid

We study the optimization of a linear function over an ellipsoid

$$\begin{aligned} \max_{\boldsymbol{\omega}} \quad & \mathbf{a}^T \boldsymbol{\omega} \\ \text{with} \quad & (\boldsymbol{\omega} - \boldsymbol{\omega}_S)^T \mathbf{A} (\boldsymbol{\omega} - \boldsymbol{\omega}_S) \leq \epsilon - \widehat{\mathcal{L}}_S(\boldsymbol{\omega}_S), \end{aligned} \quad (32)$$

which is necessary to compute the feature attribution consensus on the Rashomon Set of Linear/Additive Regression and Kernel Ridge Regression. Solving this problem can be done efficiently with a Cholesky decomposition of $\mathbf{A} = \mathbf{A}^{\frac{1}{2}} \mathbf{A}^{\frac{1}{2}T}$, which we know exists since \mathbf{A} is positive definite. Letting $\mathbf{A}^{-\frac{1}{2}} := (\mathbf{A}^{\frac{1}{2}})^{-1}$, we also get $\mathbf{A}^{-1} = \mathbf{A}^{-\frac{1}{2}T} \mathbf{A}^{-\frac{1}{2}}$. Now, it is always possible to map an ellipsoid back to the unit sphere by defining a new variable

$$\mathbf{z} := (\epsilon - \widehat{\mathcal{L}}_S(\boldsymbol{\omega}_S))^{-\frac{1}{2}} \mathbf{A}^{\frac{1}{2}T} (\boldsymbol{\omega} - \boldsymbol{\omega}_S), \quad (33)$$

see Figure 11. Applying the inverse change of variable to $\boldsymbol{\omega}$ in Equation 32, we get

$$\begin{aligned} \mathbf{a}^T \boldsymbol{\omega} &= \mathbf{a}^T ((\epsilon - \widehat{\mathcal{L}}_S(\boldsymbol{\omega}_S))^{\frac{1}{2}} \mathbf{A}^{-\frac{1}{2}T} \mathbf{z} + \boldsymbol{\omega}_S) \\ &= \sqrt{\epsilon - \widehat{\mathcal{L}}_S(\boldsymbol{\omega}_S)} \underbrace{\mathbf{a}^T \mathbf{A}^{-\frac{1}{2}T}}_{\mathbf{a}'^T} \mathbf{z} + \mathbf{a}^T \boldsymbol{\omega}_S, \end{aligned} \quad (34)$$

leading to the optimization problem

$$\begin{aligned} \max_{\mathbf{z}} \quad & \sqrt{\epsilon - \widehat{\mathcal{L}}_S(\boldsymbol{\omega}_S)} \mathbf{a}'^T \mathbf{z} + \mathbf{a}^T \boldsymbol{\omega}_S \\ \text{with} \quad & \mathbf{z}^T \mathbf{z} \leq 1, \end{aligned} \quad (35)$$

Importantly, the optimization problems of Equations 32 and 35 both reach the same optimal values. Since $\mathbf{a}'^T \mathbf{z}$ is a scalar product with a unit vector, it reaches its maximum objective value $\|\mathbf{a}'\|$ when the vector \mathbf{z} points in the exact direction of \mathbf{a}' : when $\mathbf{z}^* = \mathbf{a}' / \|\mathbf{a}'\|$. The maximum and minimum values of the objective are therefore $\pm \sqrt{\epsilon - \widehat{\mathcal{L}}_S(\boldsymbol{\omega}_S)} \|\mathbf{a}'\| + \mathbf{a}^T \boldsymbol{\omega}_S$.

A.3 Random Forests

Proposition 12 (Proposition 10) *Let $\mathcal{T} := \{t_s\}_{s=1}^M$ be a set of M trees, \mathcal{H}_m : be the set of all subsets of at least m trees from \mathcal{T} , and $\phi : \mathcal{H}_m \rightarrow \mathbb{R}$ be a linear functional, then $\min_{h \in \mathcal{H}_m} \phi(h)$ amounts to averaging the m smallest values of $\phi(t_s)$ for $s = 1, 2, \dots, M$.*

Proof We can compute the linear functional on every tree $\{\phi(t_s)\}_{s=1}^M$ and store the index of the m smallest ones in a set C_m s.t. $|C_m| = m$ and

$$s \in C_m \text{ and } s' \notin C_m \Rightarrow \phi(t_s) \leq \phi(t_{s'}). \quad (36)$$

Now, to prove to proposition, we must show that $\phi(\frac{1}{m} \sum_{s \in C_m} t_s) \leq \phi(h) \forall h \in \mathcal{H}_m$. Since $\min_{h \in \mathcal{H}_m} \phi(h) = \min_{k=m, \dots, M} \min_{h \in \mathcal{H}_k} \phi(h)$, the proof can be done two parts: first for a fixed k we prove that $\phi(\frac{1}{k} \sum_{s \in C_k} t_s) \leq \phi(h) \forall h \in \mathcal{H}_k$ and secondly prove that $\operatorname{argmin}_{k=m, \dots, M} \phi(\frac{1}{k} \sum_{s \in C_k} t_s) = m$.

Part 1 By linearity $\phi(\frac{1}{k} \sum_{s \in C_k} t_r) = \frac{1}{k} \sum_{s \in C_k} \phi(t_r)$. Also, remember that any model $h \in \mathcal{H}_k$ is associated to a subset C'_k of k seeds *i.e.* $h = \frac{1}{k} \sum_{s \in C'_k} t_r$. Importantly, since C_k and C'_k have the same size, the two sets $C_k \setminus C'_k$ and $C'_k \setminus C_k$ have a one-to-one correspondence. We get

$$\begin{aligned} \frac{1}{k} \sum_{s \in C_k} \phi(t_s) &= \frac{1}{k} \left(\sum_{s \in C_k \cap C'_k} \phi(t_s) + \sum_{s \in C_k \setminus C'_k} \phi(t_s) \right) \\ &\leq \frac{1}{k} \left(\sum_{s \in C_k \cap C'_k} \phi(t_s) + \sum_{s' \in C'_k \setminus C_k} \phi(t_{s'}) \right) \quad (\text{cf Equation 36}) \\ &= \frac{1}{k} \sum_{s \in C'_k} \phi(t_s) = \phi \left(\frac{1}{k} \sum_{s \in C'_k} t_s \right) = \phi(h). \end{aligned}$$

Part 2 We now prove that $\operatorname{argmin}_{k=m, \dots, M} \phi(\frac{1}{k} \sum_{s \in C_k} t_s) = m$. The key insight is that given $m' > m$, the set C_m contains the m smallest elements of $C_{m'}$. We get

$$\begin{aligned} \frac{1}{m'} \sum_{s \in C_{m'}} \phi(t_s) &= \frac{1}{m'} \left(\sum_{s \in C_m} \phi(t_s) + \sum_{s' \in C_{m'} \setminus C_m} \phi(t_{s'}) \right) \\ &\geq \frac{1}{m'} \left(\sum_{s \in C_m} \phi(t_s) + \sum_{s' \in C_{m'} \setminus C_m} \left[\frac{1}{m} \sum_{s \in C_m} \phi(t_s) \right] \right) \\ &= \frac{1}{m'} \left(\sum_{s \in C_m} \phi(t_s) + \frac{m' - m}{m} \sum_{s \in C_m} \phi(t_s) \right) \\ &= \frac{1}{m'} \frac{m'}{m} \sum_{s \in C_m} \phi(t_s) = \frac{1}{m} \sum_{s \in C_m} \phi(t_s), \end{aligned}$$

which ends the proof. ■