

Rational Shapley Values

DAVID S. WATSON, Department of Statistical Science, University College London, UK

Explaining the predictions of opaque machine learning algorithms is an important and challenging task, especially as complex models are increasingly used to assist in high-stakes decisions such as those arising in healthcare and finance. Most popular tools for post-hoc explainable artificial intelligence (XAI) are either insensitive to context (e.g., feature attributions) or difficult to summarize (e.g., counterfactuals). In this paper, I introduce *rational Shapley values*, a novel XAI method that synthesizes and extends these seemingly incompatible approaches in a rigorous, flexible manner. I leverage tools from decision theory and causal modeling to formalize and implement a pragmatic approach that resolves a number of known challenges in XAI. By pairing the distribution of random variables with the appropriate reference class for a given explanation task, I illustrate through theory and experiments how user goals and knowledge can inform and constrain the solution set in an iterative fashion. The method compares favorably to state of the art XAI tools in a range of quantitative and qualitative comparisons.

CCS Concepts: • **Theory of computation** → **Machine learning theory**; • **Mathematics of computing** → **Probability and statistics**.

Additional Key Words and Phrases: Explainable artificial intelligence, Interpretable machine learning, Shapley values, Counterfactuals, Decision theory

ACM Reference Format:

David S. Watson. 2022. Rational Shapley Values. In *2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22)*, June 21–24, 2022, Seoul, Republic of Korea. ACM, New York, NY, USA, 18 pages. <https://doi.org/10.1145/3531146.3533170>

1 INTRODUCTION

Machine learning algorithms are increasingly used in a variety of high-stakes domains, from credit scoring to medical diagnosis. However, many of the most popular and powerful models are *opaque*, inasmuch as end users cannot understand the reasoning that goes into individual predictions. A subdiscipline of computational statistics has rapidly emerged to address this issue, variously referred to as *interpretable machine learning* or *explainable artificial intelligence* (XAI). For good introductory surveys, see, e.g., [2, 17, 33].

The most prominent methods in this burgeoning area are post-hoc, model-agnostic tools such as feature attributions and counterfactuals. Both approaches have recently come under fire for failing to properly handle dependencies between covariates. Most algorithms for computing Shapley values – a popular feature attribution technique, formally defined in Sect. 2 – implicitly treat predictors as mutually independent, assigning positive weight to potentially impossible data permutations [16, 22]. Much recent work in counterfactual explanation has focused on how to model the impact of interventions on downstream variables, a reformulation that requires an explicitly causal approach [21, 29].

Others have pushed back against these critiques. Sundararajan and Najmi [42] demonstrate that conditioning on covariates can misleadingly assign importance to irrelevant features. Janzing et al. [19] point out that standard supervised

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.

Manuscript submitted to ACM

learning algorithms do not explicitly model dependencies between features, and so intervening to set predictors to some fixed value is not just permissible but proper when computing Shapley values.

Often lost in these debates is the seemingly obvious realization that the “right” explanation depends crucially on who is asking and why. Current explainability methods tend to treat explanations as objective deliverables that can be computed from just two ingredients: a target function and an accompanying dataset of real or simulated samples. This ignores a basic fact long acknowledged in philosophy and social science – that *explanations are fundamentally pragmatic*. To succeed, they must be tailored to the inquiring agent, who requests an explanation on the basis of certain beliefs and interests.

If pragmatic approaches are few and far between in the XAI literature, perhaps this is because it is not obvious how we should formalize abstract, subjective notions like beliefs and interests. In this paper, I use tools from Bayesian decision theory to do just that. I make three main contributions: (1) I combine disparate work in feature attributions and counterfactuals, demonstrating how the two can join forces to create more useful and flexible explanations. (2) I extend the axiomatic guarantees of Shapley values, allowing user beliefs and preferences to guide the search for a relevant subspace of contrastive baselines. (3) I implement an expected utility maximization procedure that optionally incorporates causal information to compute optimal attribution sets for individual agents.

The remainder of this paper is structured as follows. In Sect. 2, I review important background material in XAI, with a focus on Shapley values and counterfactuals, as well as the essential formalisms of decision theory and structural causal models. In Sect. 3, I compare several different reference distributions, propose a new desideratum for Shapley values, and introduce an algorithm for computing them. In Sect. 4, I put this method to work on a number of benchmark datasets. Following a discussion in Sect. 5, Sect. 6 concludes.

2 BACKGROUND

In what follows, I use uppercase italics to represent variables, e.g. X ; lowercase italics to represent their values, e.g. x ; uppercase boldface to represent matrices, e.g. \mathbf{X} ; lowercase boldface to represent vectors, e.g. \mathbf{x} ; and calligraphic type to represent the support of a distribution, e.g. \mathcal{X} .

2.1 XAI Methods

Originally developed in the context of cooperative game theory [38], Shapley values have been adapted for model interpretability by numerous authors. They represent the contribution of each feature toward a particular prediction. Let $\mathbf{x}_i \in \mathcal{X} \subseteq \mathbb{R}^d$ denote an input datapoint and $f(\mathbf{x}_i) = \hat{y}_i \in \mathcal{Y} \subseteq \mathbb{R}$ the corresponding output of our target function f for each sample $i \in [n] = \{1, \dots, n\}$.¹ Shapley values express this value as a sum:

$$f(\mathbf{x}_i) = \sum_{j=0}^d \phi_j(i), \quad (1)$$

where $\phi_0(i)$ represents i 's “baseline” (i.e., the default output of f ; more on this shortly) and $\phi_j(i)$ the weight assigned to feature X_j at point \mathbf{x}_i . Feature attribution methods that satisfy Eq. 1 are said to be *efficient*. Let $v : [n] \times 2^d \rightarrow \mathbb{R}$ be a value function such that $v(i, S)$ is the payoff associated with feature subset $S \subseteq [d]$ and $v(i, \{\emptyset\}) = 0$ for all $i \in [n]$. The

¹The model f may be a classifier, although in this case we typically require class probabilities and work on the logit scale. For ease of exposition, in this section we will assume f is a regressor.

Shapley value $\phi_j(i)$ is given by a weighted average of j 's marginal contributions to all subsets that exclude it:

$$\phi_j(i) = \frac{1}{d!} \sum_{S \subseteq [d] \setminus \{j\}} |S|!(d - |S| - 1)! v(i, S \cup \{j\}) - v(i, S). \quad (2)$$

Several properties of this value are worth noting.

- *Linearity.* For all $j \in [d]$, all $a, b \in \mathbb{R}$, and value functions v_1, v_2 , we have:

$$\phi_j(av_1 + bv_2) = a\phi_j(v_1) + b\phi_j(v_2),$$

where $\phi_j(v)$ makes explicit the dependence of attributions on value functions. (Indeed, ϕ_j only depends on i through v .)

- *Sensitivity.* If $v(i, S \cup \{j\}) = v(i, S)$ for all $S \subseteq [d] \setminus \{j\}$, then $\phi_j(i) = 0$.
- *Symmetry.* If $v(i, S \cup \{j\}) = v(i, S \cup \{k\})$ for all $S \subseteq [d] \setminus \{j, k\}$, then $\phi_j(i) = \phi_k(i)$.

It can be shown that Shapley values are the unique solution to the attribution problem that satisfies efficiency, linearity, sensitivity, and symmetry [38]. Computing exact Shapley values is NP-hard, although numerous efficient approximations have been proposed [42].

There is some ambiguity as to how one should calculate payoffs on a proper subset of features, since f requires d -dimensional input. Let $R = [d] \setminus S$, such that we can rewrite any \mathbf{x}_i as a pair of subvectors $(\mathbf{x}_i^S, \mathbf{x}_i^R)$. Then the payoff for feature subset S takes the form of a conditional expectation, with \mathbf{x}_i^S held fixed while \mathbf{x}^R varies. Following Merrick and Taly [30], I consider a general form for the value function, parametrized by a distribution \mathcal{D}_R :

$$v_{\mathcal{D}_R}(i, S) = \mathbb{E}_{\mathbf{x}^R \sim \mathcal{D}_R} [f(\mathbf{x}_i^S, \mathbf{x}^R)]. \quad (3)$$

Popular options for \mathcal{D}_R include the marginal $P(\mathbf{X}^R)$, as in Lundberg and Lee [28]; the conditional $P(\mathbf{X}^R | \mathbf{x}_i^S)$, as in Aas et al. [1]; and the interventional $P(\mathbf{X}^R | do(\mathbf{x}_i^S))$, as in Heskes et al. [18]. I revisit the distinction between these distributions in Sect. 3.1.

Counterfactual explanations, unlike Shapley values, do not produce feature weights. Instead, they identify nearby datapoints with alternative outcomes – the eponymous *counterfactuals* – intended to explain predictions by highlighting some (preferably small) set of input perturbations. The optimization problem can be expressed as

$$\mathbf{x}^* = \operatorname{argmin}_{\tilde{\mathbf{x}} \in \operatorname{CF}(\mathbf{x}_i)} d(\mathbf{x}_i, \tilde{\mathbf{x}}), \quad (4)$$

where $\operatorname{CF}(\mathbf{x}_i)$ denotes a counterfactual space, where $f(\mathbf{x}_i) \neq f(\tilde{\mathbf{x}})$ (for classification) or $f(\mathbf{x}_i) = f(\tilde{\mathbf{x}}) + \delta$ (for regression), where δ is a hyperparameter and d is some distance measure (unspecified for now). The method was originally introduced by Wachter et al. [45], who use generative adversarial networks (GANs) to solve Eq. 4. Others have proposed a variety of alternatives designed to ensure that counterfactuals are coherent and/or actionable [20, 35, 37, 43].

2.2 Bayesian Decision Theory

Why do we use XAI tools in the first place? Watson and Floridi [47] list three reasons: (1) to audit for potential bias; (2) to validate performance, guarding against unexpected errors; and (3) to discover underlying mechanisms of the data generating process. Another use case, with elements of (1) and (3) yet distinctly analyzed in the algorithmic recourse setting, is (4) to recommend actions so as to alter predicted outcomes. In all four cases, the aim is ultimately to make some sort of *decision* – be it about whether to sue a firm, deploy an algorithm, or perform an experiment.

These motivations may overlap at the edges but they represent distinct tasks based on different assumptions and requiring their own explanatory methodologies. This heterogeneity is largely ignored by current feature attribution and counterfactual approaches, which implicitly assume a sort of *explanatory objectivism*. According to this view, the quality of an explanation is independent of its context.

Decades of research in philosophy and social science has made such objectivism untenable [15, 26, 31]. If the view has any adherents today, perhaps it is because the alternative is often misconstrued as *relativism*, an anything-goes slippery slope that inevitably leads to what Feyerabend [13] (approvingly) refers to as “epistemological anarchism”. This, however, is a false dichotomy. I reject the polar extremes of objectivism and relativism in favor of *pragmatism*, which holds neither that there exists some single ideal explanation nor that all explanations are equally valid, but rather that various explanations may be more or less appropriate depending on who requests them and why.

To operationalize this insight, I rely on the formal apparatus of decision theory. Let A and H denote finite sets of actions and outcomes, respectively. An agent’s preferences over action-outcome pairs induce a partial ordering that we represent with a utility function $u : A \times H \rightarrow \mathbb{R}$. Let $p(\cdot)$ denote a credence function over outcomes such that $p(H) = 1$ and $p(h|E)$ denotes the agent’s subjective degree of belief in some $h \in H$ given a body of evidence $E = \{e_1, \dots, e_n\}$. The expected utility of each action $a \in A$ is given by a weighted average over hypotheses:

$$\mathbb{E}_{h \sim p} [u(a, h)|E] = \sum_j p(h_j|E) u(a, h_j). \quad (5)$$

If u satisfies the utility axioms (i.e., completeness, transitivity, continuity, and independence) and p satisfies the probability axioms (i.e., non-negativity, unit measure, and σ -additivity), then we say that the agent is *rational*, and it can be shown that she will tend to maximize expected utility [44]. That is, she will always choose some optimal action:

$$a^* = \operatorname{argmax}_{a \in A} \mathbb{E}_{h \sim p} [u(a, h)|E] \quad (6)$$

from a set of alternatives.

This expected utility framework is very general, and applies to explanation instances of all four types cited above. For instance, in an auditing example where a data subject believes her loan application was denied due to some sensitive attribute(s), we might have $A = \{\text{sue}, \neg\text{sue}\}$ and $H = \{\text{biased}, \neg\text{biased}\}$, with evidence provided by some XAI tool and fairness criteria [4]. In algorithmic recourse, we typically consider more diverse action sets. For example, a credit scoring case may include $A = \{\text{education}, \text{income}, \text{credit}\}$, where increasing any subset of those variables incurs some cost but may improve outcomes $H = \{\text{approve}, \neg\text{approve}\}$.

A key consideration, often ignored by authors in XAI, is the difference between model- and system-level explanations [46]. At the model level, for instance, it is typically appropriate to treat all features as mutually independent and observe outputs on real and/or synthetic inputs that vary in some systematic fashion. At the system level, by contrast, we cannot afford to ignore causal relationships between features, as an intervention on one variable may have downstream effects on others. In these cases, nature constrains the input space in accordance with certain structural dependencies.

This distinction is formalized using tools from the causal modeling literature [34]. A structural causal model (SCM) is a tuple $\mathcal{M} = \langle U, V, \mathcal{F} \rangle$, where U is a set of exogenous variables, i.e. background conditions; V is a set of endogenous variables, i.e. observed states; and \mathcal{F} is a set of structural equations, one for each $V_j \in V$. These functions map causes to direct effects in a deterministic fashion, $f_j : pa(V_j) \rightarrow V_j$, where $pa(\cdot)$ denotes the “parents” of some endogenous variable(s). The union of V_j ’s direct and indirect causes (parents, grandparents, etc.) is called its “ancestors”, denoted

$an(V_j)$. Pearl’s *do*-calculus [2000, Ch. 3.4] provides a provably complete set of rules for identifying causal effects from observational data where possible [39].

A deterministic SCM can be expanded to accommodate stochastic relationships by placing a joint probability distribution on background conditions, $p(U)$. By equating this with the credence function introduced above, we may encode an agent’s uncertainty with respect to a causal system, which has implications for expected utility if $H \in \mathcal{V}$. This decision theoretic approach to causality is explicitly endorsed by Dawid [8, 9, 10], who argues that treatment policies should be designed not by considering differences between potential outcomes, but instead by minimizing expected loss as computed over a set of observed data. A similar premise underlies recent work in causal reinforcement learning [24, 25], where the goal is to find optimal interventions by maximizing expected rewards over trials.

3 RATIONAL SHAPLEY VALUES

My basic strategy is to allow user inputs to inform the generation of counterfactuals and estimation of Shapley values, thereby resulting in custom feature attributions that are more useful than either method on its own. Intuitively, this gives agents the tools to pose not just generic (i.e., objectivist) questions of the form “Why did model f predict outcome \hat{y}_i for input \mathbf{x}_i ?”, but more targeted (i.e., pragmatic) questions, such as: (1) Why did f predict \hat{y}_i as opposed to \tilde{y} for \mathbf{x}_i ? (2) Why did f predict \hat{y}_i as opposed to \tilde{y} for \mathbf{x}_i , given \mathbf{x}_i^S ? (3) Why did f predict \hat{y}_i as opposed to \tilde{y} for \mathbf{x}_i , given \mathbf{x}_i^S and certain constraints on \mathbf{X}^R ? Unlike counterfactuals, which may confuse users by generating a large number of synthetic datapoints with no clear takeaway message, the proposed method provides unambiguous weights for all variables of interest. Unlike feature attributions, which ignore pragmatic information by construction, the present approach is *rational* in the decision theoretic sense. I will make these claims more precise and apparent below. First, however, I introduce the notion of a *relevant subspace* to formalize the hierarchy of specification implied by the questions above.

Definition 1 (Relevant subspace). *Let $\mathcal{Y}_0, \mathcal{Y}_1$ be a partition of \mathcal{Y} into baseline and contrastive outcomes, respectively, with $\hat{y}_i \in \mathcal{Y}_0$. Let $S \subseteq [d]$ denote a (potentially empty) conditioning set with complement $R = [d] \setminus S$. Let \mathcal{D}_S be the distribution of possible values for the conditioning set. Let \mathcal{D}_R be a distribution encoding probable variation in \mathbf{X}^R as a function of \mathbf{x}^S . Define $\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$. We say that counterfactual sample $\tilde{\mathbf{z}} = (\tilde{\mathbf{x}}^S, \tilde{\mathbf{x}}^R, \tilde{y} = f(\tilde{\mathbf{x}}^S, \tilde{\mathbf{x}}^R))$ is relevant for the inquiring agent if and only if: (i) $\tilde{y} \in \mathcal{Y}_1$; (ii) $\tilde{\mathbf{x}}^S \sim \mathcal{D}_S$; and (iii) $\tilde{\mathbf{x}}^R \sim \mathcal{D}_R$. Any space $\tilde{\mathcal{Z}} \subseteq \mathcal{Z}$ that meets these criteria is a relevant subspace.*

The issue with popular Shapley value approximators (e.g., SHAP [28]) is that baselines are fixed by the data centroid. Attributions are only designed to account for the discrepancy between $f(\mathbf{x}_i)$ and the empirical mean $\phi_0(i) = n^{-1} \sum_{i=1}^n f(\mathbf{x}_i)$. Moreover, it is common to use the marginal $P(\mathbf{X}^R)$ as the reference distribution \mathcal{D}_R , thereby breaking associations between the features in S and R , resulting in a value function that implicitly compares each x_{ij} to the mean of the corresponding variable X_j . Yet these values might not fall within the relevant subspace for a given agent, who may, for instance, want to know why she received an average credit score instead of a better-than-average one. Or perhaps she wants to compare her predictions only to those of others in a similar age range and income bracket, effectively setting attributions for those features to zero.

Counterfactuals naturally handle such cases, identifying real or synthetic datapoints that differ from the explanans in specific ways. However, there are several issues with the counterfactual approach. First, the cost function in Eq. 4 is almost always equated with some distance metric in practice, despite much discussion of more generic alternatives [20, 43, 48]. This is an unsatisfying notion of costs, which can and should reflect the preferences of inquiring agents much more flexibly. For instance, relevant action sets may range over interventions exogenous to the feature space, where the notion of distance is often inapplicable. Second, there is the question of how many counterfactuals to provide.

Eq. 4 suggests that just a single sample would suffice; however, such a procedure overestimates the certainty with which such a datapoint is generated. Can we really be sure that counterfactual x^* is to be preferred to counterfactual x' just because the former is marginally closer in feature space to the input x_i ? Even if we are willing to accept a purely distance-based notion of cost, it is not at all obvious that very small distances should be taken seriously when there are so many potential sources of noise – in the measurements themselves, in the finite training data, and in whatever approximations were necessary to efficiently generate candidate counterfactuals. That is why some authors prefer to sample a large, diverse coalition of counterfactuals [32], so that users can survey the various paths through which a prediction may change. Yet this approach raises new problems of intelligibility. How should an agent make sense of a relevant subspace with dozens or hundreds of samples within a tolerable distance of x_i ? How should she prioritize among different counterfactuals or integrate this information into a high-level summary?

Feature attribution methods are well suited to this task. Shapley values provide an efficient and intuitive account of the information encoded in a large set of counterfactuals, regardless of sample size. I therefore propose a hybrid method in which data are sampled from a relevant subspace, and Shapley values computed only with respect to such samples. The crucial observation is that nothing in the definition of Shapley values (Eqs. 1, 2, 3) precludes us from shifting the reference distribution from an entire observational dataset to a particular region with certain desirable properties – i.e., a relevant subspace.

3.1 Selecting the Reference Distribution

I take it more or less for granted that agents requesting explanations for algorithmic predictions have some contrastive outcome \mathcal{Y}_1 in mind. They are also likely to have some intuition about the conditioning set S , which will presumably include immutable features and/or any variables that the agent does not want to enter into the explanation if it can be avoided. The precise distribution for these features may not be available *a priori*, but conditioning on the observed x_i^S is probably a good start when the data allow for it. The trickiest component to evaluate then is the reference distribution \mathcal{D}_R , which has implications both for how counterfactuals are computed (since this helps define the relevant subspace) and for how Shapley values are estimated (since this helps define the value function).

I consider three possibilities: (1) marginal: $P(X^R)$; (2) conditional: $P(X^R|x_i^S)$; and (3) interventional: $P(X^R|do(x_i^S))$. The first point to observe is that the three distributions are not equivalent. (1) comes apart from the other two whenever statistical dependencies are present, i.e. $X^R \not\perp X^S$. Options (2) and (3) come apart whenever causal dependencies are present, i.e. if $an(X^R) \cap X^S \neq \{\emptyset\}$. I want to be clear that the differences between these distributions are in no sense normative. None is generally “better” than the others, or more useful, or more broadly applicable. Rather, each is the right answer to a different question. What follows is a brief discussion of the advantages and disadvantages associated with each value of \mathcal{D}_R , as well some heuristic guidelines on when to use which.

Marginal. The marginal distribution is preferable when the goal is a model-level explanation. This may be the case in certain instances of auditing or validation, where the analyst seeks merely to discover what the model has learned without any further restrictions. In this case, we may not care too much about impossible data perturbations (e.g., teenage grandmothers) since we simply want to recreate a decision boundary in model space. The extent to which this model space matches reality is another question altogether, presumably one that data scientists should take seriously, both before and throughout deployment. Once the model is in the wild, however, it may be of independent interest to users and regulators how it performs on a wide range of hypothetical cases, including those off the true data manifold.

Conditional. The conditional distribution is better suited to explanations at the system level, where impossible data permutations could lead to issues. This may be the case, for instance, in certain instances of validation, where we do

not wish to punish a model for failing to extrapolate to data points far from its training set. Such methods are especially valuable when joint distributions may be estimated with reasonable accuracy and causal relationships are unknown or potentially inapplicable (e.g., in image classification tasks). Estimating joint densities is a hard problem in general, especially when data are high-dimensional and/or of mixed variable types.

Interventional. The interventional distribution is optimal when seeking explanations at the system level for causal data generating processes. With access to the underlying SCM, there can be no more accurate estimator than that defined by (3). These methods are required when seeking to use XAI for discovery and/or planning, as both sorts of actions invariably rely on real-world mechanisms that cannot be approximated by either a purely marginal approach (which ignores all dependencies) or a conditional one (which fails to distinguish between correlation and causation). Of course, complete causal information is almost never available, which means that in practice analysts must make strong assumptions and/or make do with workable approximations.

I employ marginal, conditional, and interventional value functions in all experiments below (see Sect. 4). To compute interventional Shapley values, I rely on partial orderings, as these are often the most readily available to analysts at test time. A partial ordering of features implies a chain graph in which links are composed of all variables in a single causal group. For instance, users may not know the complete graph structure of the set {age, sex, income, savings, credit} but they may be more confident about the partial ordering {age, sex} \rightarrow {income, savings} \rightarrow {credit}. Lauritzen and Richardson [23] demonstrate how to factorize the probability distribution for a DAG of chain components. Heskes et al. [18] combine this procedure with a few steps of *do*-calculus to compute interventional Shapley values from observational data.

3.2 A New Axiom

Say two agents are identical along all recorded variables for some credit scoring task, and both receive lower scores than they had hoped. Counterfactuals identify two minimal perturbations sufficient to improve predictions: (1) increase education by one unit; or (2) increase savings by one unit. Let us stipulate that these actions are equivalent in terms of time and money. However, they differ in another important respect – the first agent wants to go back to school and the second wants to save more each month. In this case, *preferences alone* determine which of two explanations is optimal.

We may devise a similar example to show the impact of differing credences. Say, for instance, that two labs are planning expensive gene knockout experiments to test the predictions of a gene regulatory network inference algorithm. Their utilities may be identical – the goal for both groups is to maximize true discoveries while minimizing false positives – but they are working with different sets of evidence after preliminary, as yet unpublished results indicated to the first lab that one set of pathways is a probable dead end. In this case, *beliefs alone* determine which of two explanations is optimal.

What these examples demonstrate is that utilities and credences matter when attempting to explain model predictions. With default software like SHAP, feature attributions for these credit loan applicants or research labs would not generally differ – same inputs, same outputs. However, by explicitly incorporating pragmatic information, we can secure custom explanations at no cost to the axiomatic guarantees of Shapley values. In fact, we can extend the current desiderata.

Observe that for fixed value function v , model f , and input \mathbf{x}_i , feature attributions $\phi = \{\phi_j\}_{j=1}^d$ vary only as a function of contrastive baseline $\tilde{\mathbf{z}}$. I make this dependence explicit moving forward, writing $\phi(\tilde{\mathbf{z}})$. Define a reward function for the inquiring agent as the expected utility of the utility-maximizing action a^* , conditional on a given set of

feature attributions:

$$r(\tilde{z}) = \mathbb{E}_{h \sim p} [u(a^*, h) | \phi(\tilde{z})]. \quad (7)$$

The axiomatic constraint that characterizes pragmatic feature attributions can now be stated as follows. For the proof of Thm. 3.1, see Appx. A.

Definition 2 (Rationality). *Let $r(\cdot)$ be the reward function for a rational agent, i.e. one who behaves in accordance with the utility and probability axioms. Let $\phi(\tilde{z})$ be a feature attribution vector computed with respect to subspace \tilde{Z} . Then $\phi(\tilde{z})$ is rational if and only if, for any alternative subspace \mathcal{Z}' , $\mathbb{E}_{\tilde{z} \sim \tilde{Z}} [r(\tilde{z})] \geq \mathbb{E}_{\tilde{z}' \sim \mathcal{Z}'} [r(\tilde{z}')]$. In other words, feature attributions are rational to the extent that they tend to maximize expected rewards for the inquiring agent.*

THEOREM 3.1. *When the relevant subspace is nonempty, rational Shapley values are the unique additive feature attribution method that satisfies efficiency, linearity, sensitivity, symmetry, and rationality.*

3.3 Rational Shapley Value Algorithm

Algorithm 1: Rational Shapley Values

Input: Input datapoint x_i , set of candidate subspaces $\{\langle \mathcal{D}_S, \mathcal{D}_R, \mathcal{Y}_1 \rangle_k\}_{k=1}^m$, utility function u , credence function p

Output: $\phi(\tilde{z}^*)$

Compute classical Shapley values $\phi(z)$ using mean $\phi_0 = n^{-1} \sum_{i=1}^n f(x_i)$ and reference \mathcal{D}_R

Initialize $r(z) \leftarrow \sum_j p(h_j | \phi(z)) u(a^*, h_j)$

for $k = \{1, \dots, m\}$ **do**

 Draw data $\tilde{z}_k \sim \langle \mathcal{D}_S, \mathcal{D}_R, \mathcal{Y}_1 \rangle_k$

if $\tilde{z}_k = \{\emptyset\}$ **then**

 | Fail

else

 | Compute $\phi(\tilde{z}_k)$

 | Record $r(\tilde{z}_k) \leftarrow \sum_j p(h_j | \phi(\tilde{z}_k)) u(a^*, h_j)$

$\tilde{z}^* \leftarrow \operatorname{argmax}_{\tilde{z}} r(\tilde{z})$

I reframe the objective function of Eq. 4 with a slight twist. Our goal is not to minimize the cost or distance between two individual datapoints, but rather to compute a subspace that maximizes rewards. That is:

$$\mathcal{Z}^* = \operatorname{argmax}_{\tilde{Z} \subseteq \mathcal{Z}} \mathbb{E}_{\tilde{z} \sim \tilde{Z}} [r(\tilde{z})]. \quad (8)$$

This differs from (and improves upon) Eq. 4 in several respects. First, it results not in a single point but in a region of space, which is stabler and more informative. Second, it explicitly incorporates preferences and credences via the functions that define $r(\cdot)$. Finally, since rewards are conditioned upon a feature attribution vector, this target combines elements of both counterfactuals and feature attributions, which is preferable to either alternative alone for all the reasons argued above.

There are many ways one could go about computing this subspace in practice, and no single method is optimal in general. See Karimi et al. [20] for a recent review. In the following experiments, I follow Wexler et al. [49] in relying on observational samples. These subspaces have the advantage of being fast to find and guaranteed to lie on the true data manifold (up to potential sampling or measurement errors). They avoid the assumptions and approximations of more convoluted optimization techniques. The main disadvantage of this method is that it may provide low coverage for

regions of the feature space where data are undersampled. However, this poses issues for other approaches as well, as no method can confidently draw realistic points from low-density regions of the feature space.

The basic strategy is schematized in pseudocode (see Alg. 1). The user postulates a set of candidate subspaces, computing a sequence of rational Shapley vectors and comparing relative rewards. Evaluating $r(\tilde{z})$ with precision requires explicit utilities and (conditional) credences, which may not be generally available. This can be done informally by comparing feature attribution vectors on an ordinal basis, such that rewards are ranked rather than directly quantified. This procedure is equivalent to the cardinal alternative under reasonable assumptions about agentic rationality, as noted above. The algorithm is deliberately silent on how to select candidate subspaces, as this is irreducibly context-dependent. Credence updates may be calculated in closed form under some parametric assumptions, or else via Monte Carlo. The examples below illustrate how the method works in practice.

4 EXPERIMENTS

In this section, I describe results from a number of experiments on benchmark datasets. I compare performance against baselines using marginal, conditional, and interventional value functions, respectively labelled MSV, CSV, and ISV. Code for reproducing all results and figures can be found at https://github.com/dswatson/rational_shapley.

4.1 Auditing: COMPAS

The COMPAS algorithm is a statistical model used to assign risk scores to defendants awaiting trial. High risk scores are associated with elevated probability of recidivism, which judges in nine US states use to help decide whether to let defendants out on bail while awaiting trial. A 2016 report by ProPublica alleged that the COMPAS algorithm was racially biased against African Americans [3]; subsequent analysis by independent researchers has not always corroborated those findings [14, 36].

Though the creators of COMPAS have published technical reports defending their model, they have made neither their training data nor code publicly available, so attempts to recreate the algorithm and audit for racial bias are typically based on data gathered by ProPublica on some 12,000 individuals arrested in Broward County, Florida between 2013 and 2014. Using a random forest [5], I regress violent risk scores on the following features: age, number of priors, and whether the present charge is a felony (all deemed “admissible”); as well as race and sex (deemed “inadmissible”). For the purposes of this experiment, I take the trained model as the true target.

MSV is computed via Monte Carlo [41], with 2000 simulations per attribution. CSV is computed using the TreeSHAP method [27], since random forests are composed of individual regression trees. For ISV, I presume a simple partial ordering in which demographic variables (age, sex, and race) are root nodes, upon which all other predictors depend. The mean predicted response for all defendants in the dataset is $\phi_0 = 3.69$.

I focus specifically on high-risk defendants, namely the 264 subjects with risk scores at or above the 95th percentile according to the COMPAS approximator. Some 94% of these defendants are African American, compared to just 60% in the general dataset, which immediately raises questions about racial bias. MSV, CSV, and ISV all look fairly similar in this case (see Fig. 1). Age and priors generally receive the highest feature attributions in this group, consistent with the findings of other researchers. Race also receives nonzero attribution throughout. Notably, race receives higher attribution under the ISV reference than MSV, since MSV can only detect direct effects, and race is presumed to have both direct and indirect effects given the partial ordering. In this example, I presume that auditors are working at the model level – i.e., evaluating the patterns learned by the algorithm, rightly or wrongly – in which case the marginal distribution is best suited to their purpose.

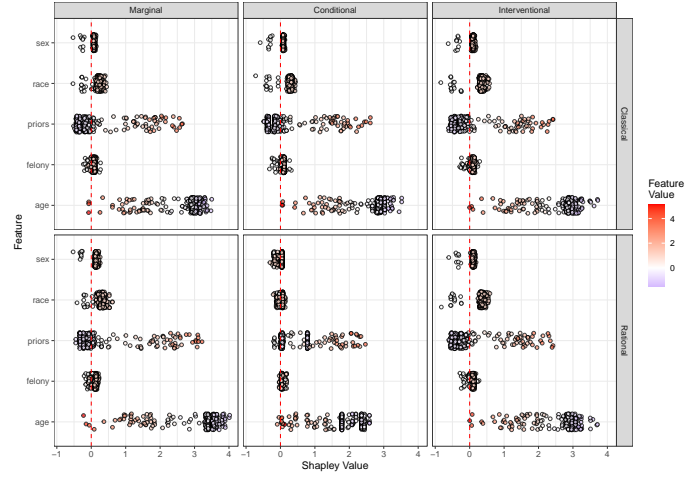


Fig. 1. Classical and rational Shapley values for the Broward County COMPAS dataset, computed with marginal, conditional, and interventional reference distributions. Continuous feature values are z-scored for visualization.

To compute rational Shapley values, I consider a relevant subspace of defendants in the same age range (at most 33 years old) with at least one prior, yet who fell in the bottom half of defendants by predicted risk score. All other features are allowed to vary unconstrained. 356 subjects in the data meet these criteria, and therefore constitute the reference sample for this experiment. This is too large a group of counterfactuals to “explain” anything on their own, even if they were ranked by some measure (e.g., distance from the target input). What the task requires is a summary of differences across features, a job for which Shapley values are uniquely well-suited. The racial breakdown of defendants in the relevant subspace is markedly different from our high-risk group – just 48% are African American – further reinforcing concerns about potential bias in the COMPAS data. The resulting Shapley values look fairly similar to their classical counterparts, although priors and age appear to dominate the explanations under conditional and interventional reference distributions. Using the marginal reference, however – which, as noted above, is most relevant for auditing scenarios – we find greater evidence of racial bias using rational Shapley values. The racial attribution gap between Black and white defendants is significantly greater in the relevant subspace than it is across the complete dataset ($t = 3.874, p < 0.001$).

Consider the case of a young Black defendant deciding whether to file suit against the makers of COMPAS. Despite having just a single prior offense, this 21-year-old was placed in the highest risk group according to the algorithm. A utility matrix for this individual is given in Table 1. Assume that the defendant in question assigns a uniform prior over H to begin with. We say that h_1 is corroborated to the extent that race receives greater positive attribution for African Americans than it does for defendants of other races. Then although both classical and rational Shapley values recommend the same action (a_1), the latter does so with higher expected reward, since $p(h_1|\phi(\tilde{z})) > p(h_1|\phi(z))$. Rational Shapley values are therefore preferable in this case, as theory suggests.

4.2 Discovery: Medical Diagnosis

As an example of XAI for discovery, I consider the diabetes dataset originally described by [12]. The data consist of 442 patients and 10 predictors – age, sex, body mass index (BMI), blood pressure (MAP), and six blood serum measurements,

Table 1. Utility matrix for a 21-year-old African American defendant with a single prior, predicted to be high risk and deciding whether to sue the makers of COMPAS.

	h_1 : biased	h_2 : \neg biased
a_1 : sue	5	-1
a_2 : \neg sue	0	0




Fig. 2. Classical and rational Shapley values for the diabetes dataset, computed with marginal, conditional, and interventional reference distributions. Continuous feature values are z-scored for visualization.

including three cholesterol-related variables (LDL, HDL, and TC). The goal is to predict disease progression after one year, for which I use an elastic net regression [50]. Note that exact Shapley values can be analytically calculated for linear models under the assumption of feature independence. However, exploratory data analysis reveals strong correlations between some features in the data (especially blood serum measurements), and causal dependencies between features such as BMI and MAP are well established. In such a case, the analytic formulae for local explanation in linear models are inapplicable. For ISV, I assume the partial ordering $\{\text{age, sex, bmi}\} \rightarrow \{\text{map, ldl, hdl, tc}\} \rightarrow \{\text{tch, ltg, glt}\}$. This treats age, sex, and BMI as root nodes; MAP and cholesterol as causally intermediate; and remaining blood serum measurements are downstream. Such a DAG admittedly oversimplifies several complex biochemical processes but is broadly consistent with known structural relationships.

I focus on patients with especially poor prognoses, as these are typically the subjects of greatest concern to clinicians. Specifically, I examine the top decile of patients by disease activity ($\hat{y} \geq 265$). Classical Shapley values for these patients are visualized under all three value functions in Fig. 2. CSV and ISV are fairly similar here, although both are quite different from the marginal alternative. For instance, TC receives relatively large attributions under the marginal reference distribution, but nearly none under conditional or interventional distributions, owing to its strong collinearity with other cholesterol measures.

For rational Shapley analysis, I consider a subspace consisting of older men (i.e., those in the top quartile by age), who are generally at higher risk than other groups in this dataset. Since age and sex are (more or less) unactionable, a rational agent will likely want prognostic explanations that place little or no weight on such features. We find in

Table 2. Complete feature vector for both Bert and Ernie in the diabetes dataset.

sex	age	bmi	map	tc	ldl	hdl	tch	ltg	glu
1	-0.064	0.096	0.105	-0.003	-0.005	-0.007	-0.003	0.023	0.073

Table 3. Utility matrices for Bert (left) and Ernie (right). Bert would like to reduce his risk of diabetes without lowering his carb intake; Ernie wants the same outcome without lowering his fat intake.

Bert	$h_1 : \hat{y} < 265$	$h_2 : \hat{y} \geq 265$	Ernie	$h_1 : \hat{y} < 265$	$h_2 : \hat{y} \geq 265$
a_1 : low-fat	5	-1	a_1 : low-fat	1	-6
a_2 : low-carb	1	-6	a_2 : low-carb	5	-1

Table 4. Rational Shapley vectors for Bert (top) and Ernie (bottom), computed from the same input vector (see Table 2) but with respect to different relevant subspaces.

sex	age	bmi	map	tc	ldl	hdl	tch	ltg	glu
-3.789	-4.867	94.922	29.241	1.763	0.849	4.053	6.566	13.019	4.324
-1.291	0.114	42.604	66.942	2.187	1.129	-0.167	-3.532	11.136	9.903

Fig. 2 that BMI and MAP receive large rational attributions on average, reflecting their greater variance in the relevant subspace. Conditional and interventional reference distributions are especially valuable here, as the analytical goal almost certainly requires a system-level approach and causal dependencies cannot therefore be ignored. Alternative structural assumptions may lead to somewhat different attributions, but focusing on a relevant subspace has effectively driven the Shapley values for unactionable variables like sex and age toward zero. The result is that explanations are provided primarily in terms of lifestyle variables like BMI and MAP, which agents can work to minimize in order to improve their disease prognosis.

Consider the case of two high-risk patients – call them Bert and Ernie – who share identical feature vectors (see Table 2; note that the negative value of age is a quirk of the scaling procedure.) Despite their striking similarities, however, Bert and Ernie differ in their culinary tastes. Bert was raised on carbohydrate-rich foods, and cannot imagine living without bread, pasta, and potatoes; Ernie, by contrast, is a voracious carnivore who eats bacon for breakfast, cold cuts for lunch, and steak for dinner. Their respective utility matrices are given by Table 3, where we assume for the sake of this example that only two dietary interventions are under consideration: a low-carb diet and a low-fat diet. Evidence suggests that the latter is slightly more effective for weight loss, while the former has the added benefit of reducing blood pressure [6]. Classical Shapley values, which place the greatest weight on BMI, would tend to favour a_1 regardless of agentive preferences. However, rational Shapley values can distinguish between optimal explanations based entirely on utilities, since changes to either BMI or MAP are sufficient to bring about the desired outcome, albeit at differing costs to Bert and Ernie. Thus we find the following two feature attribution vectors (see Table 4), computed using the conditional reference distribution on inputs that differ only with respect to utilities u .

Observe that, among the many differences between these two vectors, they flip the relative importance of BMI and MAP, suggesting two alternative paths toward reducing disease risk. A counterfactual approach with a well-designed cost function could in principle also identify different explanations for these two agents, but it could not additionally provide complete feature attributions summarising the relative impact of other predictors. Such a synthesis is unique to the rational Shapley approach.




Fig. 3. Classical and rational Shapley values for the German credit dataset, computed with marginal, conditional, and interventional reference distributions. Continuous feature values are z-scored for visualization.

4.3 Recourse: Credit Scoring

As a final example, I examine the recourse setting, where the goal is to advise an agent what interventions are necessary and/or sufficient to secure a desired prediction. As Karimi et al. [21] have shown, this task requires causal information. More specifically, they prove that noncausal recommendations, such as those provided by Ustun et al. [43], guarantee recourse if and only if the treatment variables have no descendants in the underlying causal graph. For this experiment, I use the German credit dataset from the UCI Machine Learning Repository [11]. To simplify the presentation, I restrict focus to just seven of the most informative features: age A , gender G , marital status M , job J , savings S , loan amount L , and duration D . The outcome Y is binary (loan approved/denied), with a base rate of 70%. I train a support vector machine (SVM) [7] to predict Y using a radial basis kernel.

I assume a simple partial ordering in which demographic variables are causally antecedent to both financial predictors and loan application details: $\{A, G, M\} \rightarrow \{L, D, J, S\}$. The UCI website notes that false positives are more costly for banks than false negatives, and therefore recommends penalizing the former at five times the rate of the latter. Post-processing the SVM, an optimal decision threshold is obtained at 0.74, and I therefore label all and only those points with predicted probabilities at or above this value as 1. The resulting model is about 69% accurate on the true positives and 73% accurate on the true negatives. I focus on the 189 borderline applicants with predicted success probabilities on the interval $[0.7, 0.74)$. These are subjects who could plausibly benefit from even a slight improvement in their model predictions. Classical Shapley values are unable to help such applicants, since the mean response in the full dataset is $\phi_0 = 0.7$, i.e. the lower bound of the interval. (Irrational) feature attributions for these subjects are therefore guaranteed to sum to some value on $[0, 0.04)$, explaining why they did slightly better than average instead of why they were ultimately unsuccessful.

The most important component of the relevant subspace is the contrastive outcome \mathcal{Y}_1 , as the borderline applicants appear fairly representative along all predictors of interest. I therefore stratify by predicted response, zooming in on the top quartile of applicants ($\hat{y} \geq 0.76$). Bivariate analyses suggest that such applicants are for the most part financially

Table 5. Complete feature vector for Ruth in the German credit dataset.

gender	marital	age	amount	duration	savings	job
0	0	26	3181	26	290	2

Table 6. Utility matrix for Ruth, whose demographic and financial information place her just below the loan approval threshold for a hypothetical bank.

	$h_1 : \hat{y} \geq 0.74$	$h_2 : \hat{y} < 0.74$
a_1 : Increase S	1	-2
a_2 : Decrease L	2	-1
a_3 : Decrease D	2	-1

Table 7. Rational Shapley values for Ruth in the German credit dataset.

gender	marital	age	amount	duration	savings	job
0	0	-0.168	0.061	-0.115	0.013	0.128

secure and request small loans of relatively brief duration. The classical and rational Shapley vectors look fairly similar in this experiment, except for the crucial difference in offset. (See Fig. 3; note that the response variable is transformed to the logit scale for easier visualization.)

Consider the case of a loan applicant, call her Ruth, whose predicted success probability is 0.73 – just under the decision boundary. Her input feature vector can be found in Table 5, where values are reported on the original (pre-transformation) scale. She is looking to buy a new home and believes her high-skilled job and decent savings (both well above their respective median values in the dataset) should make her a strong applicant. She is willing to make some changes if necessary – specifically, to increase her savings, decrease the size of her requested loan, or pay it back sooner – but in any case she would like to act quickly, since she fears the house she wants to buy will not be on the market for long. Ruth’s utility matrix is given in Table 6. We find here that she prefers changing the terms of her application to increasing her savings in general, as this is likely the faster route to bank approval. She is indifferent between changing the loan amount or duration. As noted above, computing Shapley values from the entire dataset is pointless in this case, as Ruth’s predicted outcome is already above the base rate. Counterfactuals could potentially aid her search for an explanation, but only by either providing an artificially narrow set of pathways to approval or overwhelming her with an unnecessarily large reference class of successful applicants.

Rational Shapley values, computed from a relevant subspace of successful loan applicants who share Ruth’s demographic characteristics, provide a summary of feature attributions within this reference class (see Table 7; note that calculations are performed in logit space.) Interestingly, we find that her requested loan amount, though above average for the relevant subspace, did not hurt Ruth’s application. On the contrary, it appears to boost her chances somewhat, a result we might not expect from surveying the counterfactuals directly. The loan duration, however – about double the average for the relevant subspace – brought her success probability down considerably. This is a highly actionable piece of information, inasmuch as it guides Ruth toward algorithmic recourse that can push her application over the decision boundary. Of course, if the change is too onerous for Ruth, then she is free to resample from another subspace of successful applicants with longer loan duration on average. The procedure may continue like this indefinitely, with agents testing out new hypotheses in an iterative, exploratory fashion, updating their rewards accordingly.

5 DISCUSSION

The skeptical reader may plausibly object that the examples above are fairly neat and straightforward, with their small utility matrices of well-defined user preferences. Reality, of course, is far messier. Complete sets of actions and outcomes may not be known upfront, let alone utility and credence functions defined thereon. Potential interventions may be far more numerous than these experiments permit, and associated outcomes entirely uncertain. Can rational Shapley values scale to larger, more complicated instances of algorithmic explanation?

In a word, yes. These cases are merely illustrative, following on the back of theoretical results establishing the viability of a pragmatic synthesis between feature attributions and counterfactuals that works with various different reference distributions. The limiting factors in these experiments were complications around efficient Monte Carlo sampling, conditional probability estimation, and reliable partial orderings. The first can be entirely resolved with greater computational power. The second requires some care, but may in principle be addressed with flexible methods such as variational autoencoders. The third relies on background knowledge, which may be substantial in some areas, and should in any case grow over time. Improvements in any or all of these areas can help extend the rational Shapley method to larger, more complex cases.

Another potential challenge to the proposed method is that it is vulnerable to confirmation bias. If users can scan the data in search of attribution vectors that make their preferred outcomes more likely, what's to stop them from finding such vectors even when do they do not exist? There are three answers to this charge. First, it is entirely possible that the relevant subspace be empty – that is, no observed or synthetic datapoint meets the threefold criteria specified by $\langle \mathcal{D}_S, \mathcal{D}_R, \mathcal{Y}_1 \rangle$. Such a failure is highly informative, as it demarcates a realm of (im)possibility for the inquiring agent. Thus it is simply false to allege that users will always find what they are looking for. Second, in cases where the relevant subspace is nonempty but sparsely populated, we should expect estimates to be unstable. Any good inference procedure should take such uncertainty into account, making it difficult for an opportunistic user's desired outcome to pass severe tests on the basis of just a small handful of outlying points. I have not said much about testing here, but it is simple in principle to extend the experiments above with frequentist or Bayesian methods tailored to the particular task at hand [40]. Third, there is no guarantee that agents will generally concur on proper values of \mathcal{D}_S , \mathcal{D}_R , and \mathcal{Y}_1 . This makes it all the more important to be explicit about which values of each input go into any given explanation. Such transparency will make it harder for adversarial agents to game the system and easier for those acting in good faith to come to consensus on particular cases.

6 CONCLUSION

The two most popular XAI tools available for local, post-hoc explanations – feature attributions and counterfactuals – each have certain advantages. However, they tend to provide different answers in particular cases that can confuse and/or mislead end users. In this paper, I have shown how the two can be synthesized into a single method that preserves the best elements of each. Rational Shapley values preserve and extend the axiomatic guarantees of their classical forebears, giving users the flexibility to compute additive feature attributions in a fast, flexible manner. Marginal, conditional, and interventional value functions were all evaluated, and the applicability of each was compared across different use cases. By formalizing the task in an expected utility framework, I was able to demonstrate how and why individual agents may rationally seek different explanations for the same model predictions. The resulting explanations are concise, intuitive, and thoroughly pragmatic.

ACKNOWLEDGMENTS

This research was supported by ONR grant N62909-19-1-2096.

REFERENCES

- [1] Kjersti Aas, Martin Jullum, and Anders Løland. 2021. Explaining individual predictions when features are dependent: More accurate approximations to Shapley values. *Artif. Intell.* 298 (2021), 103502.
- [2] Amina Adadi and Mohammed Berrada. 2018. Peeking inside the black-box: A survey on explainable artificial intelligence (XAI). *IEEE Access* 6 (2018), 52138–52160.
- [3] J Angwin, J Larson, S Mattu, and L Kirchner. 2016. *Machine bias*. Technical Report. ProPublica. <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>
- [4] Solon Barocas, Moritz Hardt, and Arvind Narayanan. 2019. *Fairness and Machine Learning*. fairmlbook.org. <http://www.fairmlbook.org>.
- [5] Leo Breiman. 2001. Random Forests. *Mach. Learn.* 45, 1 (2001), 1–33.
- [6] Grant D Brinkworth, Manny Noakes, Jonathan D Buckley, Jennifer B Keogh, and Peter M Clifton. 2009. Long-term effects of a very-low-carbohydrate weight loss diet compared with an isocaloric low-fat diet after 12 mo. *Am. J. Clin. Nutr.* 90, 1 (2009), 23–32.
- [7] Corinna Cortes and Vladimir Vapnik. 1995. Support-vector networks. *Mach. Learn.* 20, 3 (1995), 273–297.
- [8] Philip Dawid. 2012. The Decision-Theoretic Approach to Causal Inference. In *Causality: Statistical Perspectives and Applications*, Carlo Berzuini, Philip Dawid, and Luisa Bernardinelli (Eds.). Wiley, London, Chapter 4, 25–42.
- [9] Philip Dawid. 2015. Statistical Causality from a Decision-Theoretic Perspective. *Annu. Rev. Stat. Appl.* 2, 1 (2015), 273–303.
- [10] Philip Dawid. 2020. Decision-theoretic foundations for statistical causality. (2020). *arXiv preprint, 2004.12493*.
- [11] Dheera Dua and Casey Graff. 2017. *UCI Machine Learning Repository*. University of California, School of Information and Computer Science, Irvine, CA.
- [12] Bradley Efron, Trevor Hastie, Iain Johnstone, and Robert Tibshirani. 2004. Least Angle Regression. *Ann. Stat.* 32, 2 (2004), 407–499.
- [13] Paul Feyerabend. 1975. *Against Method*. New Left Books, London.
- [14] Aaron Fisher, Cynthia Rudin, and Francesca Dominici. 2019. All Models are Wrong, but Many are Useful: Learning a Variable’s Importance by Studying an Entire Class of Prediction Models Simultaneously. *J. Mach. Learn. Res.* 20, 177 (2019), 1–81.
- [15] Luciano Floridi. 2019. *The Logic of Information*. Oxford University Press, Oxford.
- [16] Christopher Frye, Ilya Feige, and Colin Rowat. 2020. Asymmetric Shapley values: Incorporating causal knowledge into model-agnostic explainability. In *Advances in Neural Information Processing Systems*.
- [17] Riccardo Guidotti, Anna Monreale, Salvatore Ruggieri, Franco Turini, Fosca Giannotti, and Dino Pedreschi. 2018. A Survey of Methods for Explaining Black Box Models. *Comput. Surveys* 51, 5 (2018), 1–42.
- [18] Tom Heskes, Evi Sijben, Ioan Gabriel Bucur, and Tom Claassen. 2020. Causal Shapley Values: Exploiting Causal Knowledge to Explain Individual Predictions of Complex Models. In *Advances in Neural Information Processing Systems*.
- [19] Dominik Janzing, Lenon Minorics, and Patrick Bloebaum. 2020. Feature relevance quantification in explainable AI: A causal problem. In *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics*, Vol. 108. 2907–2916.
- [20] Amir-Hossein Karimi, Gilles Barthe, Bernhard Schölkopf, and Isabel Valera. 2020. A survey of algorithmic recourse: Definitions, formulations, solutions, and prospects. (2020). *arXiv preprint, 2010.04050*.
- [21] Amir-Hossein Karimi, Bernhard Schölkopf, and Isabel Valera. 2021. Algorithmic Recourse: From Counterfactual Explanations to Interventions. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. 353–362.
- [22] Indra Kumar, Suresh Venkatasubramanian, Carlos Scheidegger, and Sorelle Friedler. 2020. Problems with Shapley-value-based explanations as feature importance measures. In *Proceedings of the 37th International Conference on Machine Learning (ICML ’20)*. 1–10.
- [23] Steffen L Lauritzen and Thomas S Richardson. 2002. Chain graph models and their causal interpretations. *J. Royal Stat. Soc. Ser. B Methodol.* 64, 3 (2002), 321–348.
- [24] Sanghack Lee and Elias Bareinboim. 2018. Structural Causal Bandits: Where to Intervene?. In *Advances in Neural Information Processing Systems*, Vol. 31. 2568–2578.
- [25] Sanghack Lee and Elias Bareinboim. 2019. Structural Causal Bandits with Non-Manipulable Variables. *Proceedings of the AAAI Conference on Artificial Intelligence* 33, 1 (2019), 4164–4172.
- [26] Catherine Legg and Christopher Hookway. 2020. Pragmatism. In *The Stanford Encyclopedia of Philosophy* (fall 2020 ed.), Edward N Zalta (Ed.). Metaphysics Research Lab, Stanford University.
- [27] Scott M Lundberg, Gabriel Erion, Hugh Chen, Alex DeGrave, Jordan M Prutkin, Bala Nair, Ronit Katz, Jonathan Himmelfarb, Nisha Bansal, and Su-In Lee. 2020. From local explanations to global understanding with explainable AI for trees. *Nat. Mach. Intell.* 2, 1 (2020), 56–67.
- [28] Scott M Lundberg and Su-In Lee. 2017. A Unified Approach to Interpreting Model Predictions. In *Advances in Neural Information Processing Systems*. 4765–4774.
- [29] Divyat Mahajan, Chenhao Tan, and Amit Sharma. 2019. Preserving Causal Constraints in Counterfactual Explanations for Machine Learning Classifiers. In *CausalML: Machine Learning and Causal Inference for Improved Decision Making Workshop, NeurIPS 2019*.

- [30] Luke Merrick and Ankur Taly. 2020. The Explanation Game: Explaining Machine Learning Models Using Shapley Values. In *Machine Learning and Knowledge Extraction - 4th International Cross-Domain Conference (CD-MAKE)*. Springer, 17–38.
- [31] Tim Miller. 2019. Explanation in artificial intelligence: Insights from the social sciences. *Artif. Intell.* 267 (2019), 1–38.
- [32] Ramaravind K. Mothilal, Amit Sharma, and Chenhao Tan. 2020. Explaining Machine Learning Classifiers through Diverse Counterfactual Explanations. In *Proceedings of the International Conference on Fairness, Accountability, and Transparency in Machine Learning* (Barcelona, Spain). 607–617.
- [33] W James Murdoch, Chandan Singh, Karl Kumbier, Reza Abbasi-Asl, and Bin Yu. 2019. Definitions, methods, and applications in interpretable machine learning. *Proceedings of the National Academy of Sciences* 116, 44 (2019), 22071 – 22080.
- [34] Judea Pearl. 2000. *Causality: Models, reasoning, and inference*. Cambridge University Press, New York.
- [35] Rafael Poyiadzi, Kacper Sokol, Raul Santos-Rodriguez, Tijl De Bie, and Peter Flach. 2020. FACE: Feasible and actionable counterfactual explanations. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*. 344–350.
- [36] Cynthia Rudin, Caroline Wang, and Beau Coker. 2020. The Age of Secrecy and Unfairness in Recidivism Prediction. *Harvard Data Sci. Rev.* 2, 1 (2020).
- [37] Chris Russell. 2019. Efficient Search for Diverse Coherent Explanations. In *Proceedings of the International Conference on Fairness, Accountability, and Transparency*. 20–28.
- [38] Lloyd Shapley. 1953. A Value for n -Person Games. In *Contributions to the Theory of Games*. Princeton University Press, Princeton, Chapter 17, 307–317.
- [39] Ilya Shpitser and Judea Pearl. 2008. Complete Identification Methods for the Causal Hierarchy. *J. Mach. Learn. Res.* 9 (2008), 1941–1979.
- [40] Dylan Slack, Anna Hilgard, Sameer Singh, and Himabindu Lakkaraju. 2021. Reliable Post hoc Explanations: Modeling Uncertainty in Explainability. In *Advances in Neural Information Processing Systems*, Vol. 34. 9391–9404.
- [41] Erik Štrumbelj and Igor Kononenko. 2014. Explaining prediction models and individual predictions with feature contributions. *Knowledge and Information Systems* 41, 3 (2014), 647–665.
- [42] Mukund Sundararajan and Amir Najmi. 2019. The many Shapley values for model explanation. In *Proceedings of the ACM Conference*.
- [43] Berk Ustun, Alexander Spangher, and Yang Liu. 2019. Actionable Recourse in Linear Classification. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*. 10–19.
- [44] John von Neumann and Oskar Morgenstern. 1944. *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, NJ.
- [45] Sandra Wachter, Brent Mittelstadt, and Chris Russell. 2018. Counterfactual explanations without opening the black box: Automated decisions and the GDPR. *Harvard J. Law Technol.* 31, 2 (2018), 841–887.
- [46] David S Watson. 2022. Conceptual challenges for interpretable machine learning. *Synthese* 200, 2 (2022), 65.
- [47] David S Watson and Luciano Floridi. 2021. The explanation game: a formal framework for interpretable machine learning. *Synthese* 198, 10 (2021), 9211–9242.
- [48] David S Watson, Limor Gultchin, Ankur Taly, and Luciano Floridi. 2022. Local Explanations via Necessity and Sufficiency: Unifying Theory and Practice. *Minds Mach.* 32, 1 (2022), 185–218.
- [49] J Wexler, M Pushkarna, T Bolukbasi, M Wattenberg, F Viégas, and J Wilson. 2020. The What-If Tool: Interactive Probing of Machine Learning Models. *IEEE Transactions on Visualization and Computer Graphics* 26, 1 (2020), 56–65.
- [50] Hui Zou and Trevor Hastie. 2005. Regularization and Variable Selection via the Elastic Net. *J. Royal Stat. Soc. Ser. B Methodol.* 67, 2 (2005), 301–320.

A APPENDIX

The following is a proof of Thm. 3.1, which states that rational Shapley values are the unique additive feature attribution method that satisfies efficiency, linearity, sensitivity, symmetry, and rationality.

PROOF. Since uniqueness is already well established for the classical axioms, all that remains is to show that rational Shapley values (a) satisfy the additional rationality axiom and (b) do not violate any of the classical axioms.

Take (a) first. Assume there exists some \mathcal{Z}' that generates greater expected rewards for an agent than the relevant subspace $\tilde{\mathcal{Z}}$. Then conditioning on the corresponding Shapley values $\phi(z')$ must lead to an action of greater expected utility than the action recommended by $\phi(\tilde{z})$. This means that the values for some $\phi_j(\tilde{z})$ are either too large or too small. But this can only arise from a misspecification of (i) \mathcal{Y}_1 , (ii) \mathcal{D}_S , or (iii) \mathcal{D}_R , which together permit either too much variation in X_j (resulting in inflated values of $|\phi_j(\tilde{z})|$), or too little (resulting in deflated values of $|\phi_j(\tilde{z})|$). By Def. 1, any $\tilde{\mathcal{Z}}$ that fails to satisfy one of these criteria is not a relevant subspace. Thus we have a contradiction.

One can verify (b) by confirming that rational Shapley values do not deviate from the classical formulae (see Eqs. 1, 2, 3). The only difference between the current proposal and more familiar alternatives is their respective methods

of specifying baseline expectations ϕ_0 and reference distributions \mathcal{D}_R . Shifting these two parameters changes the interpretation of resulting Shapley vectors, but does not alter their essential character. Indeed, much recent work on Shapley values has been devoted to exploring alternative reference distributions; for more on the plurality of Shapley values, see [30, 42]. Any procedure that inputs valid values for ϕ_0 and \mathcal{D}_R and conforms to the characteristic equations enjoys the same properties. \square