

Data Science Principles for Interpretable and Explainable AI

KRIS SANKARAN^{1,*}

¹*Department of Statistics, University of Wisconsin - Madison, Madison, WI, United States*

Abstract

Society’s capacity for algorithmic problem-solving has never been greater. Artificial Intelligence is now applied across more domains than ever, a consequence of powerful abstractions, abundant data, and accessible software. As capabilities have expanded, so have risks, with models often deployed without fully understanding their potential impacts. Interpretable and interactive machine learning aims to make complex models more transparent and controllable, enhancing user agency. This review synthesizes key principles from the growing literature in this field. We first introduce precise vocabulary for discussing interpretability, like the distinction between glass box and explainable models. We then explore connections to classical statistical and design principles, like parsimony and the gulfs of interaction. Basic explainability techniques – including learned embeddings, integrated gradients, and concept bottlenecks – are illustrated with a simple case study. We also review criteria for objectively evaluating interpretability approaches. Throughout, we underscore the importance of considering audience goals when designing interactive data-driven systems. Finally, we outline open challenges and discuss the potential role of data science in addressing them. Code to reproduce all examples can be found at <https://go.wisc.edu/3k1ewe>.

Keywords *explainability; Human Computer Interaction; interpretability; trustworthy machine learning*

1 Introduction

The success of Artificial Intelligence (AI) stems in part from its ability to abstract away details of its context. When a model is applied to predict responses from a collection of features, it matters little whether the response is the shape of a galaxy or the next word in a sentence. Such abstraction has enabled the design of versatile methods with remarkably diverse applications. Nonetheless, models are never used in a vacuum. Models affect human well-being, and it is important to systematically study the human-model interface. How can we ensure that a model behaves acceptably in unforeseen circumstances? How can we prevent harm and unfair treatment? How can we align model behavior with our expectations? The field of interpretable machine learning and explainable AI (XAI) has emerged to address these questions, allowing model-based abstractions to be used wisely within messy, real-world contexts. This literature has engaged researchers from various backgrounds, from theorists seeking to characterize the fundamental limits of explainability techniques (Williamson and Feng, 2020; Bilodeau et al., 2024) to philosophers and psychologists critiquing popular discourses surrounding them (Krishnan, 2019; Nussberger et al., 2022).

* Email: ksankaran@wisc.edu.

What can go wrong without interpretability? Consider these examples:

- Caruana et al. (2015) aimed to help a hospital triage incoming pneumonia patients, ensuring that high-risk patients are admitted while allowing lower-risk patients to receive outpatient care. Using a dataset of past pneumonia patients, their model achieved an area under the receiver operating characteristic curve (AUC) of 0.857 in predicting patient mortality based on lab and medical profiles. Surprisingly, an important feature analysis revealed that asthmatic patients had lower predicted probability of death. This arose because the original training data were observational. Historically, asthmatic patients were treated more aggressively than those without asthma, and therefore enjoyed lower mortality. Directly deploying such a model would put asthma patients at risk. Fortunately, their model was editable, so the asthma association could be removed before deployment.
- Gu et al. (2019) designed stickers that, when attached to stop signs, caused object detector models used by self-driving cars to misclassify them as speed limits. This real-world adversarial attack exploits the sensitivity of models to small (and, to humans, seemingly arbitrary) input perturbations. More interpretable models could yield systems with more predictable behavior.
- In Google’s first public demo of the generative AI chatbot Bard, it was asked, “What new discoveries from the James Webb Space Telescope (JWST) can I tell my 9 year old about?” The chatbot responded with three seemingly legitimate facts. Only after the demo was shared online did astronomers point out that the third fact, “JWST took the very first pictures of a planet outside of our own solar system,” was in fact untrue — the first exoplanet had been imaged in 2004 (Kundaliya, 2023). AI chatbots can “hallucinate” falsehoods, even when their output seems authoritative. Research on XAI could shed light on how these systems generate responses, allowing them to be treated less like oracles and more like search engines.

These examples beg the question: What makes a model interpretable? Directly answering this questions is complex, so let’s consider a simpler one: What makes a data visualization effective? This question has been studied for decades (Cleveland, 1993; Tufte, 2001; Agrawala et al., 2011; Sedlmair et al., 2012), and every data scientist has firsthand experience with it. A good visualization streamlines a taxing cognitive operation into a perceptual one. It helps when the visual encoding – the mapping from sample properties to graphical elements – uses representations that are already familiar or easily learnable. Further, the graphical elements must be legible and well-annotated. We also tend to learn more from information-dense visualizations (Tufte, 2001; Oppermann and Munzner, 2022), since they prevent oversummarization and can highlight details for follow-up study. Similarly, an interpretable model can be broken down into relevant components, each of which can be assigned meaning. Instead of information density (showing more of the data), interpretability relies on faithfulness (showing more of the model).

The parallels run deeper. As in a good visualization, the data provenance of a trustworthy model can be traced back to the original measurement mechanism — beautiful design and high-accuracy have little value otherwise. Moreover, like visualization, interpretability must be tailored to an audience and the problems they need solved (Lipton, 2018). There are different levels of data literacy, and visual representations may be familiar to some audiences but not others. Similarly, AI models are employed across a range of problem domains, necessitating validation in realistic settings (see Section 3). Finally, effective visualizations push readers beyond passive consumption — they inspire deeper exploration of complexity. Likewise, interpretability can support the transition from automated to augmented decision-making (Heer, 2019), enhancing rather than substituting human reason.

Model interpretability can be approached with the same nuance that is already routine in

data visualization. The data visualization literature has identified systematic design approaches and formal evaluation criteria, cautioning against blanket statements about entire approaches. Indeed, even information density has been thoughtfully critiqued (Borkin et al., 2013). Similarly, discussions of interpretability can go beyond dichotomies about models being either glass or black boxes. To help us navigate the gray areas of this growing field, Section 1.1 introduces relevant vocabulary. Section 2 outlines representative techniques and applies them to a simple simulation example. This section, and the code that accompanies it (<https://go.wisc.edu/3k1ewe>), can serve as a tutorial on the practical implementation of interpretability techniques. Finally, Sections 3–4 examine the conceptual and technical questions required for effective evaluation and progress as the AI landscape evolves.

1.1 Vocabulary

The terms used in interpretability research can be confusing, because though terms like “explainable” are familiar in a colloquial sense, they have a precise technical meaning in the literature. Indeed, one difficulty is that many people agree that model interpretability is important — they just have different ideas about what it looks like. Therefore, it helps to have vocabulary that distinguishes key properties while maintaining an appropriate level of abstraction.

An important distinction is between an *intrinsically interpretable model* and an *explainability technique* (Murdoch et al., 2019; Rudin, 2019). An intrinsically interpretable model is one that, by virtue of its design, is easy to accurately describe and alter. For this reason, they are often called glass boxes. A canonical example is a sparse linear model. For any new example, a prediction can be constructed in a single pass over the features with nonzero coefficients. In contrast, explainability techniques are designed to improve our mental models of arbitrary black box models. A common example is the partial dependence profile (Friedman, 2001), defined as $PDP(x_d) = \frac{1}{N} \sum_{i=1}^N \hat{f}(x_d, \mathbf{x}_{i(-d)})$, which describes the influence of the d^{th} feature when holding all others fixed at values observed in the original data $\mathbf{x}_1, \dots, \mathbf{x}_N$. Though originally designed to summarize gradient boosting fits \hat{f} , the fact that it only requires $(\mathbf{x}, \hat{f}(\mathbf{x}))$ pairs means it applies to any supervised learner. To summarize, interpretable models are like glass boxes whose inner workings are transparent, while explainability techniques are systematic ways of analyzing the outputs emerging from black boxes.

We can relate these approaches to data science and design principles; see Table 1. First, consider intrinsically interpretable models. When we call a linear model interpretable, we are invoking *parsimony* and *simulatability*. Parsimony means that predictions can be traced back to a few model components, each of which comes with a simple story attached. For example, sparse linear models have few nonzero coefficients, and the relationship between each coefficient and the output can be concisely described. A similar principle applies to generalized additive modeling, which have few allowable interactions (Caruana et al., 2015), or in latent variable models, which have few underlying factors (Sankaran and Holmes, 2023). Simulatability formalizes the idea that predictions can be manually reconstructed from model descriptions. For example, in decision trees and falling rules lists, this can be done by answering a sequence of yes-no questions. Finally, since interpretation requires human interaction, it can be affected by the gulfs of interaction (Hutchins et al., 1985). Ideally, an interpretability method should allow users to quickly query properties of the underlying model. Further, the method’s outputs should be immediately relevant to the query, requiring no further cognitive processing. To the extent that a method has reached these two ideals, it has reduced the gulfs of execution and evaluation, respectively.

Table 1: The core data science principles underlying interpretability. They can guide the development of more trustworthy and transparent AI models.

Principle	Discussion
Parsimony	To facilitate interpretation, the total number of relevant model components should be relatively small. For example, in ℓ_1 -regularized methods, the number of coefficients is kept small, and decision trees with fewer splits are more interpretable than those with many.
Simulatability	Given a sample and a model description, how easy would it be for a user to manually derive the associated prediction? Regression models with few coefficients and approaches like rule lists tend to be more simulatable than those that involve more intensive or multistep pooling of evidence.
Sensitivity	How robust are predictions to small changes in either the data or the model? Ideas from robust statistics can inform how we understand variable importance and stability in more complex AI models.
Navigating Scales	Local interpretations concern the role of individual samples in generating predictions, while global interpretations contribute to understanding a model overall. Interpretation at both scales can guide the appropriate real-world model use.
Interaction Gulfs	The “Gulf of Execution” is the time it takes for an interactive program to respond to a user’s input, while the “Gulf of Evaluation” is the time it takes for the user to understand the output and specify a new change. Both can influence the practicality of an interpretability method.

Unfortunately, even when restricting focus to glass box models, these definitions are still somewhat ambiguous, leading to issues with evaluation. For example, simulatability alone does not quantify the user effort required to form predictions; it also ignores their accuracy. Linear models become unwieldy when they have many nonzero coefficients (Poursabzi-Sangdeh et al., 2021), and deep decision trees take time to parse. Similarly, latent variables are only interpretable if they can be related to existing knowledge.

Among explainability techniques, a key question is scale, and this leads to the distinction between *global* and *local explanations* (Lundberg et al., 2020). Global explanations summarize the overall structure of a model. Partial dependence profiles are an example of a global explanation, because they describe \hat{f} across the range of possible inputs. In contrast, local explanations support reasoning about an individual sample’s prediction. For example, when a sentiment analysis model classifies a particular hotel review as negative, we might be curious about which phrases were most important in that classification. Global explanations shed light on the high-level properties of a model while local explanations dissect specific instances of a model’s prediction. Local explanations are built on the observation that perturbations, either of the model or the data, provide valuable context for understanding model-based decisions. Sensitivity, a principle with a long history in statistics (Tukey, 1959; Huber, 1964; Holmes, 2017), is a fruitful path to explainability. Nonetheless, we caution that differences between local and global can blur in practice (Dandl et al., 2023; Herbringer et al., 2022). For example, concept bottleneck models

(Section 2.3.4), globally restrict the space of possible prediction functions, forcing them to pass through an interpretable concept bottleneck. However, editing the bottleneck for specific examples allows for local interventions. As detailed in Section 2.3.4, if sensitive attributes are used as concepts, then their coefficients can be set to zero, ensuring that they don't influence the final prediction, even indirectly through related features. In light of these ambiguities, it is helpful to think of the vocabulary in this section as highlighting the main axes of variation in the space of techniques, a map to guide application and evaluation of interpretability methods.

2 Methods

How is interpretability put into practice? We next review approaches for direct interpretability and XAI. Rather than attempting to review all proposals from this rapidly growing literature, we prioritize the simple but timeless ideas that lie behind larger classes of methods. To ground the discussion, we include a simulation example where the generative mechanism is simple to communicate, but complex enough to warrant analysis with AI.

2.1 Simulation Design

The simulation is inspired by the types of longitudinal studies often encountered in microbiome data science (Jeganathan et al., 2018; Silverman et al., 2018; Kodikara et al., 2022). Consider following 500 participants for two months, gathering microbiome samples each day. At the study's conclusion, each participant has their health checked, and they are assigned to either healthy or disease groups. For example, the participants may have contracted HIV or developed Type I Diabetes, as in (Gosmann et al., 2017) and (Kostic et al., 2015), respectively. The question of interest is whether there are systematic differences in the microbiome trajectories between groups. These might have diagnostic potential and could shed light on the underlying mechanisms of disease development. For example, does the disease group have lower levels of a protective species? Do they exhibit unstable dynamics for any important microbial subcommunities? For either fundamental biology or diagnostic potential, any models applied to the data must be either interpretable or explainable – classification alone has little value unless it can be incorporated into a medically coherent narrative.

In our hypothetical scenario, we simulate $N \in \{500, 1000, 5000, 10000\}$ participants across $T = 50$ time points, with each sample containing abundances of $D = 144$ microbial species. Figure 1 illustrates trajectories for a selection of taxa and participants. The generative mechanism is as follows: we create $K = 25$ community trajectories $\mathbf{C}_k \in \mathbb{R}^{T \times D}$, for $k = 1, \dots, K$ by sampling:

$$\mathbf{C}_{k,d} = \begin{cases} \text{Unif}[0, 0.01]^T & \text{with probability } 0.7 \\ \mathbf{f}_{kd}^{\text{increase}} & \text{with probability } 0.1 \\ \mathbf{f}_{kd}^{\text{decrease}} & \text{with probability } 0.1 \\ \mathbf{f}_{kd}^{\text{bloom}} & \text{with probability } 0.1. \end{cases} \quad (1)$$

The species-wise random trajectories \mathbf{f}_{kd} are defined as follows. For random increases, we compute the cumulative sum of random nonnegative weights. Decreases, as in Fig 1b, follow the same

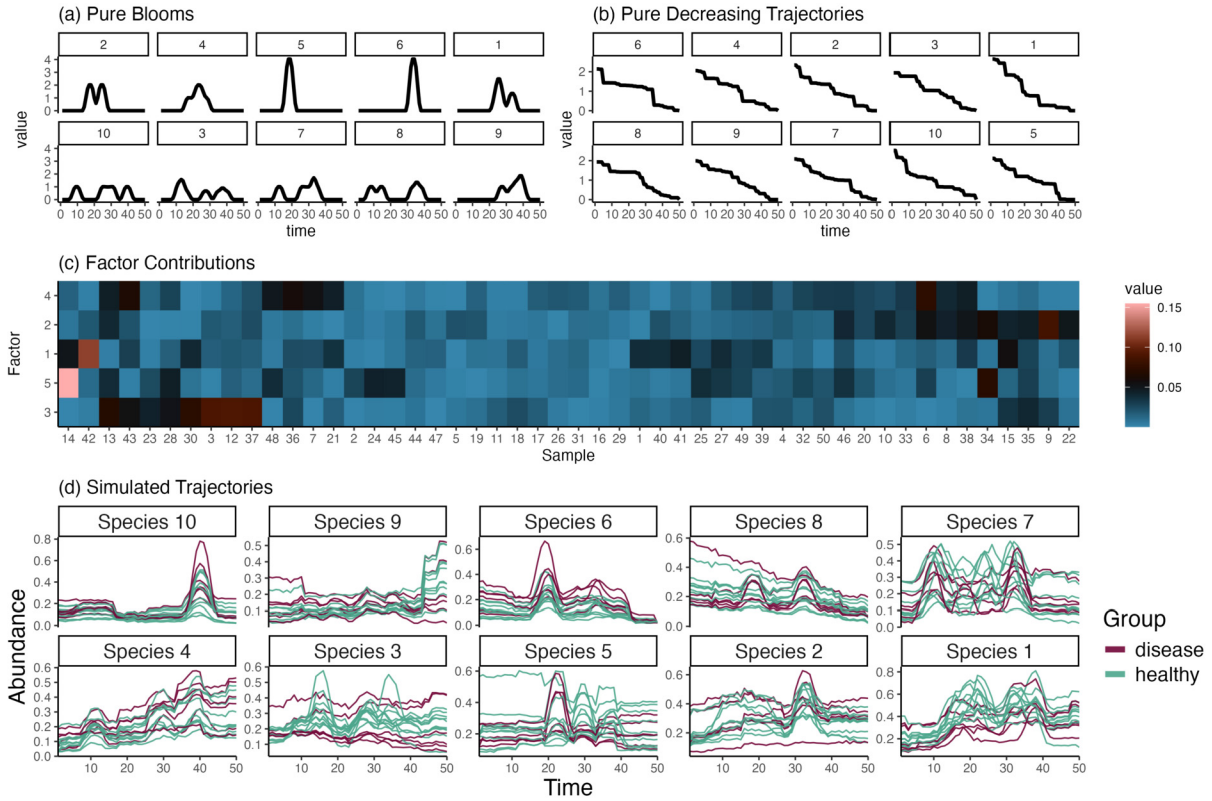


Figure 1: Example simulated trajectories. (a - b). Trajectories for two latent types. Each sample is a mixture of these types. Even for a fixed type, samples may have different variants, e.g., blooms at different times. (c) Observed trajectories are mixtures of these pure underlying trajectory types. The contributions of five latent trajectory patterns for each sample are shown in this heatmap. (d) Healthy and disease samples generated through this process for a subset of samples and taxa.

process but with a sign reversal. Specifically, the T coordinates of $\mathbf{f}_{kd}^{\text{increase}}$ are defined using,

$$\tilde{f}_{kdt}^{\text{increase}} = \sum_{t'=1}^t u_{t'} \quad (2)$$

$$(u_1, \dots, u_T) \sim \text{Dirichlet}(\lambda_u \mathbf{1}_T), \quad (3)$$

and then renormalizing $\tilde{\mathbf{f}}_{kd}$ to sum to T . We set $\lambda_u = 0.3$ to allow for occasional large jumps. Random blooms $\mathbf{f}_{kd}^{\text{bloom}}$ like those in Fig 1a are formed using,

$$\tilde{f}_{kdt}^{\text{bloom}} = \sum_{t^*} K_{r,L}(t - t^*) \quad (4)$$

$$t_1^*, \dots, t_{n_{\text{bloom}}}^* | n_{\text{bloom}} \sim \text{Unif}[L, T - L] \quad (5)$$

$$n_{\text{bloom}} \sim \text{Poi}(\lambda_{\text{bloom}}), \quad (6)$$

where K_r is a Tukey window function (Van Boxtel et al., 2021) with bandwidth $r = 0.9$ and window size $L = 9$. The trajectory is then renormalized as with $\mathbf{f}_{kd}^{\text{increase}}$. Given dictionaries \mathbf{C}_k

of community trajectories, samples $\mathbf{x}_i \in \mathbb{R}^{T \times D}$ for subject i are sampled,

$$\mathbf{x}_i = \sum_{k=1}^K \theta_{ik} \mathbf{C}_k \quad (7)$$

$$(\theta_{i1}, \dots, \theta_{iK}) \sim \text{Dir}(\lambda_\theta \mathbf{1}_K). \quad (8)$$

To generate class assignments y_i , we cluster the community weights θ_{ik} into 24 clusters. These clusters are then randomly divided into 12 healthy and 12 disease groups. Note that each cluster contains a mixture of the $K = 25$ community trajectories. This approach therefore links subjects with similar underlying community trajectories into similar health outcomes in a way that is highly nonlinear.

In practice, the steps of the generative mechanism would not be known. Visually examining example trajectories in Figure 1d does not suggest clear choices about which microbiome features are relevant for disease classification. Each taxon exhibits a complex, nonlinear trend over time, and while some participants show “blooms” at similar times, it’s unclear whether these are linked to disease. Further, the sheer number of taxa makes drawing conclusions through manual inspection impractical. Computational approaches are necessary to distinguish disease from healthy trajectories and identify taxonomic or temporal features of interest, which could then be validated in follow-up studies involving direct interventions on features in a controlled setting.

2.2 Intrinsically Interpretable Models

We can often learn salient, non-obvious characteristics of the data by applying a intrinsically interpretable model. Some have argued that with sufficient ingenuity, such models can rival the accuracy and efficiency of any black box (Rudin, 2019). Even otherwise, the results of intrinsically interpretable analyses can guide the application of more sophisticated models later on. With this in mind, we briefly review sparse logistic regression (Friedman et al., 2010) and decision trees (Loh, 2014), apply them to the simulation, and interpret the results.

2.2.1 Sparse Logistic Regression

Sparse logistic regression solves the optimization problem,

$$\hat{\boldsymbol{\beta}} = \arg \min_{\boldsymbol{\beta} \in \mathbb{R}^D} \sum_{i=1}^N \ell(y_i | \mathbf{x}_i, \boldsymbol{\beta}) + \lambda \|\boldsymbol{\beta}\|_1,$$

where $\ell(y_i | \mathbf{x}_i, \boldsymbol{\beta}) = \log \left((1 + \exp(\mathbf{x}_i^T \boldsymbol{\beta} y_i))^{-1} \exp(\mathbf{x}_i^T \boldsymbol{\beta} y_i) \right)$. This optimization identifies a sparse coefficient vector $\hat{\boldsymbol{\beta}}$ that maximizes the log-likelihood of pairs (\mathbf{x}_i, y_i) . The parameter λ controls sparsity: higher values yield simpler but less expressive models. The combination of sparsity and linearity ensures interpretability. Sparsity allows many features within \mathbf{x}_i to be ignored during prediction, while linearity ensures that when features do contribute, they enter the model transparently. Specifically, a one-unit increase in the d^{th} -coordinate x_d changes the relative risk for class $y = 1$ by $\exp(\beta_d)$.

2.2.2 Decision Trees

In contrast, decision trees do not solve a single optimization problem. Instead, they are defined through a recursive algorithm. Initially, we scan all D dimensions of \mathbf{x}_i and determine thresholds t_d such that $\mathbb{I}(x_{id} > t_d)$ maximizes classification accuracy. From these candidates, we select the optimal d^* , defining the first partition along the axis d^* . The two elements of the partition correspond to two leaves of a tree split according to x_{id^*} . Next, for a tree with L leaves, we iteratively split each element along axis d at a threshold t_d^l , for all pairs of leaves l and dimensions d , selecting the split that maximizes accuracy. This process stops once model complexity outweighs performance on a holdout set.

Interpretability here stems from parsimony and simulatability. The final tree has only as many partition elements as leaves, and with sufficient cost on model complexity, this can be easily visualized as a single tree. Further, the axis-aligned splits yield straightforward descriptions for each partition element, recoverable by tracing a path from the root to the leaf. Finally, predictions remain constant within each element, changing only across partition boundaries.

2.2.3 Advances

Sparse and tree-based models are the foundation for many proposals for interpretable machine learning (Xin et al., 2022; Zhong et al., 2023). For example, Zeng et al. (2016) introduced fast decision trees that match black box models in recidivism prediction, a domain with significant moral ramifications where interpretability is crucial. Similarly, Hazimeh et al. (2020) designed a decision tree variant that competes effectively with deep learning in various operations research tasks. By leveraging advances in combinatorial optimization, their approach allows interpretable trees to be learned from very large datasets without excessive compute demands.

Similarly, ℓ_1 -regularization is central to many modern interpretability workflows, even in problems that initially seem distant from traditional statistical learning frameworks. For example, Fridovich-Keil et al. (2022) developed an interpretable approach to the view synthesis problem, where many two-dimensional views of an object are consolidated into a coherent three-dimensional reconstruction. Instead of the typical deep encoder-decoder architecture, they used ℓ_1 -regularized regression on the appropriate Fourier representation. This approach achieved better accuracy than competing deep learning approaches with far less sensitivity to hyperparameter selection, and could be trained at a fraction of the compute cost.

2.2.4 Application to the Simulation

We next apply these intrinsically interpretable models to the simulated trajectories from Section 2.1. We explore two approaches. First, we concatenate time points for each participant into a single $T \times D$ -dimensional vector. This method can uncover temporally localized predictive signals. Second, we derive time series summaries for each taxon, reducing the number of predictors to $S \times D$ for S different summaries. In our study, we focus solely on linear trends and curvature. For each subject i and taxon d , we fit a linear model $x_{itd} \sim \mu_{id} + \beta_{id}t$, with time as a predictor. The estimated $\hat{\beta}_{id}$ summarizes the linear trend for that trajectory. Similarly, we use the sum-of-squares of the temporal second differences to summarize curvature: $\frac{1}{T-2} \sum_{t=2}^{T-1} (x_{itd+1} - x_{itd-1})^2$. Note that summarization derives potentially relevant features that are not simple linear combinations of the original inputs.

The holdout accuracy for the sparse logistic regression and decision tree models under these two representations of \mathbf{x}_i and across a range of sample sizes N is given in Table 2. Regularization

Table 2: In and out-of-sample accuracies of models applied to the simulation study. For cross-validated models, training time was averaged across folds. Training time was computed on a 2022 MacBook Pro with 16GB RAM and an M2 GPU. Models differ in their inherent interpretability, manual feature curation effort, and generalization performance. Deep learning models can overfit the training sample while still generalizing well. Intrinsically interpretable models can be substantially improved through effective featurization.

Data	Model	Samples ($\times 1000$)	In-Sample Accuracy	Out-of-Sample Accuracy	Training Time (s)
Original	Sparse Logistic Reg.	0.5	81.2%	78.0%	3.9
		1	82.1%	81.0%	15.9
		5	78.1%	77.9%	22.9
		10	77.5%	77.0%	37.83
	Decision Tree	0.5	91.6%	70.2%	95.9
		1	91.3%	76.5%	497.9
		5	91.4%	75.7%	3255.3
		10	90.3%	78.4%	7676.9
	Transformer	0.5	100%	86.4%	91.2
		1	90.7%	76.0%	246.6
		5	96.6%	86.4%	1167.6
		10	87.6%	87.7%	1964.4
	Concept Bottleneck	0.5	98.4%	84.8%	110.9
		1	99.2%	73.2%	211.6
		5	90.4%	88.2%	892.8
		10	92.5%	89.0%	1854.0
Featurized	Sparse Logistic Reg.	0.5	92.6%	87.8%	0.4
		1	88.8%	85.4%	8.1
		5	86.5%	85.3%	32.7
		10	85.9%	85.5%	119.5
	Decision Tree	0.5	74.4%	69.6%	1.8
		1	87.9%	70.6%	14.7
		5	90.2%	75.1%	185.5
		10	88.1%	77.7%	770.8

λ and tree complexity parameters were chosen with four-fold cross validation. For these data, the sparse logistic regression model generally outperforms the decision tree. Moreover, manually derived summary statistics yield better performance compared to concatenating all inputs. This suggests a general takeaway: a simple model on the appropriate representation can achieve good performance, and performance is contingent on the representation. The widely invoked “interpretability–accuracy tradeoff” overlooks this nuance. Further, Table 2 includes the training time for the models, illustrating a broad range of runtimes among interpretable models. Simple doesn’t always mean fast or scalable.

With respect to interpretability, when $N = 500$, the sparse logistic regression model on concatenated inputs with the optimal λ used 37 taxon-by-time point features, while the decision tree needed 13 splits. In contrast, with the summary statistic featurization, the optimal sparse

2.3 Deep Learning and XAI

The sparse logistic regression with derived summary statistics seems to be our best option in the case study. However, we should be concerned by the fact that we manually define our features – what other potentially predictive features might we be missing out on? Black box models are appealing because they automate the representation learning process (Bengio, 2009; Bengio et al., 2013). For example, in computer vision, deep learning models can identify higher-order, semantically meaningful features from raw pixel inputs (Zeiler and Fergus, 2013; Simonyan et al., 2014; Yosinski et al., 2015; Achtabat et al., 2023; Bykov et al., 2023). They can learn features that activate on images of eyes simply by being given training data of faces and an appropriate loss function (Ross et al., 2017). Applying a deep learning model to our microbiome data could reveal predictive temporal and taxonomic features (e.g., the presence of particular subcommunities) that we had not previously considered. The challenge, then, would be to explain these models in a way that advances scientific knowledge and ensures accurate medical diagnoses. To demonstrate the process, let’s explore using a transformer encoder model in our simulation.

2.3.1 Transformers

We first review the mechanics of transformer architectures in deep learning. Beyond our case study, transformers are particularly relevant to XAI because they have become a universal building block of modern AI. What used to be multiple specialized architectures – like CNNs for vision, LSTMs for text, and autoregressive networks for audio – have now become consolidated into different forms of transformers. At a high-level, the transformer is an architecture for deriving features from sequential data. It works by recombining representations of individual sequence elements through a series of “attention” operations, which prioritize segments of the sequence relevant to the current element’s updated representation. It takes low-level representations of individual sequence elements, like one-hot encodings of individual words, and then modifies them to reflect larger sequential context.

Specifically, consider a sequence of tokens $\mathbf{x}_1, \dots, \mathbf{x}_T$. These tokens can be embedded into a sequence $\mathbf{z}_1, \dots, \mathbf{z}_T$ of high-dimensional (e.g., 512) real-valued vectors. In a word-based language model, \mathbf{x}_t could be one-hot encodings of the words from the language’s vocabulary (Pennington et al., 2014), while in a vision problem, \mathbf{x}_t might represent neighboring patches from an image (Kolesnikov et al., 2021). The \mathbf{z}_t vectors are meant to capture the high-level features of the original t^{th} token (word or patch) within its larger context (sentence or image). Therefore, they are much more valuable for downstream analysis than word identity or pixel values considered in isolation.

Transformer models use a mechanism called self-attention to *transform* the original \mathbf{x}_t into the \mathbf{z}_t (Vaswani et al., 2017; Lee, 2021). If we stack the \mathbf{x}_t into $\mathbf{X} \in \mathbb{R}^{T \times D}$, then self-attention has the form,

$$\frac{1}{\sqrt{D_A}} (\mathbf{X} \mathbf{W}_q \mathbf{W}_k^T \mathbf{X}^T) \mathbf{X} \mathbf{W}_v,$$

for trainable parameters $\mathbf{W}_q \in \mathbb{R}^{D \times D_A}$, $\mathbf{W}_k \in \mathbb{R}^{D \times D_A}$, $\mathbf{W}_v \in \mathbb{R}^{D \times D_v}$. The term in parenthesis has dimension $T \times T$ and summarizes the extent to which coordinates t' in the output “attend” to coordinates t in the input. The matrix \mathbf{W}_v allows attention to return a linearly transformed version of \mathbf{X} , which can often be helpful for increasing/decreasing output dimensions. Typically, several of these transformations are applied in parallel, concatenated into an object called “multihead

attention,” and passed through a softmax nonlinearity. The output of this complete process is the matrix of embeddings with rows $\mathbf{z}_1, \dots, \mathbf{z}_T$.

The embeddings \mathbf{z}_t are difficult to interpret because their derivation is neither parsimonious nor simulatable. Empirically, the representations have been shown to encapsulate high-level features associated with each token’s context. However, the number of updates, the fact that the update can “attend” to all other tokens simultaneously, and the fact that the attention mechanism is nonlinear all contribute to the difficulty in interpreting these embeddings as well as any decision rules derived from them. Nonetheless, several approaches have emerged for enhancing the interpretability of these types of deep, embedding-based models (Erhan et al., 2010; Wongsuphasawat et al., 2018). We review three approaches below: global explainability, integrated gradients, and concept bottleneck models.

2.3.2 Embedding Visualization

Embedding visualizations treat the learned representations \mathbf{z}_t as data and apply Exploratory Data Analysis techniques to better understand their properties. Methods like principal components analysis, canonical correlation analysis, or UMAP are often applied to the collection of \mathbf{z}_i (Nguyen and Holmes, 2019; Raghu et al., 2021, 2019). Since the embeddings can be studied simultaneously across the full dataset, this is an example of a global explainability technique.

For example, Coenen et al. (2019) used embeddings to analyze the BERT language model (Devlin et al., 2019). In one experiment, they examined sentences containing words that can have multiple meanings depending on context, like, “fair.” By computing embeddings for instances of “fair” within these sentences, it was observed that although the word remained the same, its embeddings distinctly clustered in a two-dimensional principal components projection. This clustering reflected the context dependence of the word’s meaning, with clusters representing the legal (“fair use”), mathematical (“fair coin”), and civic (“world’s fair”) uses. This qualitative study can also be accompanied by quantitative analysis. For example, Coenen et al. (2019) and Hewitt and Manning (2019) used linear probes to quantify the extent to which grammatically structure (represented by parse tree representations) could be reconstructed from BERT embeddings. Figure 3 in Section 2.2.4 gives embedding visualization examples in the context of our simulation study. In particular, Figure 3d parallels the approach of Coenen et al. (2019), focusing on a species instead of a word.

2.3.3 Integrated Gradients

When seeking local explanations, embedding visualization lacks sufficient specificity. We often have questions in mind that require more than model-based summaries – for example, what was it about a hotel review that led to its classification as negative sentiment, or how could a loan applicant modify their application so that it can be approved? Addressing these questions requires thinking carefully about the model’s sensitivity to perturbation. Many methods formalize this intuition – influence functions (Koh and Liang, 2017), Shapely values (Lundberg and Lee, 2017; Fumagalli et al., 2023), LIME (Ribeiro et al., 2016), GradCAM (Selvaraju et al., 2017), and Integrated Gradients (IG) (Sundararajan et al., 2017) for example. Here, we dive into the mechanics of IG.

The integrated gradient of a sample \mathbf{x}_i with respect to class y defined by:

$$\text{IG}(\mathbf{x}_i) = (\mathbf{x}_i - \mathbf{x}_0) \times \int_{\alpha=0}^1 \frac{\partial f_y(\mathbf{x}_0 + \alpha \times (\mathbf{x}_i - \mathbf{x}_0))}{\partial \mathbf{x}_i} d\alpha, \quad (9)$$

where \mathbf{x}_0 is a reference sample (e.g., an all-black “empty” image) and f_y is the model’s predicted probability for the input being assigned to class y . The gradient $\frac{\partial f_y(\mathbf{u})}{\partial \mathbf{u}}$ represents how infinitesimal changes in each coordinate of \mathbf{u} affects the probability of assigning the input to class y . At first, this appears to be sufficient for local explanation – large gradient coordinates correspond to words or pixels with the largest influence on the class assignment. However, this often fails in practice due to saturation. Specifically, deep learning classifiers often use a multiclass logistic function in the final layer, and for large input magnitudes, this function’s gradients approach zero, rendering explanations based on $\frac{\partial f_y(\mathbf{u})}{\partial \mathbf{u}}$ unsatisfactory.

The solution proposed by IG is to consider versions of the input $\mathbf{x}_0 + \alpha \times (\mathbf{x}_i - \mathbf{x}_0)$ that have been shrunk down to \mathbf{x}_0 , hence avoiding vanishing gradients. Since it might not be clear how much to shrink at first, IG considers scaling along an entire sequence of $\alpha \in [0, 1]$. Integrating gradients over this range gives a more faithful description of the i th coordinate’s contribution to classification as y . For example, in Sturmfels et al. (2020), gradients for the object of interest (a bird) are much larger for small α values, ensuring that the final IG for the bird class places appropriate weight on that class.

Integrated gradients and its local explanation counterparts are widely used in XAI, but they are also among the most controversial (Lundstrom et al., 2022). Both theoretical and empirical results suggest that explanations do not necessarily respect expectations. For example, Adebayo et al. (2018) observed that saliency maps, including those from IG, derived from randomly initialized neural network classifiers look closely resemble those from properly trained ones. A theoretical study by Bilodeau et al. (2024) found that locally similar functions could nonetheless be assigned different explanations at their query points, indicating that global properties can influence local explanations. These alarming findings have renewed interest in the interpretability community to develop local explainability techniques with formal guarantees.

2.3.4 Concept Bottleneck Models

The explainability techniques we have discussed so far depend only on the model and either the original training data, for embedding analysis, or the specific example to be explained, for integrated gradients. This is both a strength and a limitation. On the one hand, it ensures generality, the same explanation method can be effective across various applications. On the other, it limits the incorporation of domain-specific language into the explanation. Concept Bottleneck Models (CBMs) (Koh et al., 2020) offer an abstraction for infusing domain-specific language into models. CBM “concepts” bridge model and human representations of relevant data features. The overarching goal of this class of models is to create representations that can be manipulated by both the model (as high-dimensional, numerical vectors) and experts (as community-derived concepts).

CBMs leverage expert knowledge to form annotations useful for both model training and explanation. During training, each sample has the form $(\mathbf{x}_i, y_i, \mathbf{c}_i)$. \mathbf{x}_i and y_i are the usual supervised model input and output for sample i . $\mathbf{c}_i \in \{0, 1\}^K$ represents a dense concept annotation. For example, \mathbf{x}_i could be the image of a bird, y_i its species label, and \mathbf{c}_i a summary of interpretable features like wing color or beak length. The CBM learns a model $\mathbf{x}_i \rightarrow \mathbf{c}_i \rightarrow y_i$, compressing low-level measurements in \mathbf{x}_i into concept annotations \mathbf{c}_i before making a prediction y_i . Koh et al. (2020) propose several strategies for implementing this bottleneck, including joint training of functions f and g to minimize the weighted sum of concept L_1 and label L_2

losses:

$$\hat{f}, \hat{g} := \arg \min_{f, g} \sum_i L_1(\mathbf{c}_i, f(\mathbf{x}_i)) + \lambda L_2(y_i, g(\mathbf{c}_i)), \quad (10)$$

for some $\lambda > 0$.

On a test sample \mathbf{x}^* , predictions are made using $\hat{\mathbf{c}} = f(\mathbf{x}^*)$ and $\hat{y} = g(\hat{\mathbf{c}}) = g(f(\mathbf{x}^*))$. Note that ground truth concept labels \mathbf{c}_i are not needed at this stage. Typically, a simple model g , like multiclass logistic regression, is used. CBMs automate human-interpretable feature extraction, bridging manual feature engineering with black box representation learning. While manually designing a “wing color” feature extractor would be tedious, it is helpful to know that the model f learns a proxy for such a feature, and its output can be used like any human engineered feature. Moreover, this approach streamlines counterfactual reasoning. For example, we can ask what the model would have predicted had the wing been yellow instead of red, simply by intervening on the coordinate of \mathbf{c}_i associated with wing color. This is important in fairness or recourse analysis (Mehrabi et al., 2021; Karimi et al., 2022). For example, if we know that a feature should not be used for prediction (e.g., the race of a loan applicant), then we can zero out the coefficient from that concept in the model g .

Listing 1: An example implementation of a Concept Bottleneck Model. The variable \mathbf{c} represents the concept labels, which are transformed into class predictions using a shallow multilayer perceptron, `self.mlp`. See the full implementation at the repository <https://go.wisc.edu/3k1ewe>.

```
class ConceptBottleneck(nn.Module):
    """
    Learn Concepts and Classes for Sequences
    """
    def __init__(self, n_embd=144, n_positions=50, n_layer=6, n_concept=25, n_class=2):
        super(ConceptBottleneck, self).__init__()
        self.n_concept = n_concept
        self.n_class = n_class
        config = GPT2Config(n_embd=n_embd, n_positions=n_positions, n_layer=n_layer)
        self.backbone = GPT2Model(config)
        self.concept = nn.Linear(n_embd * n_positions, self.n_concept)
        self.mlp = nn.Sequential(
            nn.Linear(self.n_concept, self.n_concept),
            nn.ReLU(),
            nn.Linear(self.n_concept, self.n_concept),
            nn.ReLU(),
            nn.Linear(self.n_concept, self.n_concept),
            nn.ReLU(),
            nn.Linear(self.n_concept, self.n_class - 1)
        )

    def forward(self, x):
        z = self.backbone(inputs_embeds=x)
        c = self.concept(z.last_hidden_state.view(x.shape[0], -1))
        return c, self.mlp(c)
```

2.3.5 Application to the Simulation

With these XAI techniques, we can apply black box models to our microbiome case study and still draw scientific insight from them. The simulated dataset is well-suited to a transformer

encoder architecture, where each sample corresponds to a word and the full trajectory to a sentence. The final healthy vs. disease classification is analogous to text sentiment classification. Guided by these parallels, for each N we trained a GPT-2 architecture with a random sample of $0.75N$ simulated participants; the remaining $0.25N$ are reserved for validation. We used default hyperparameters from a public implementation with the two modifications. First, we changed the input dimension to 144 to one-hot encode the species. Second, we reduced the number of layers to 6 from the default of 12, since the sample sizes are relatively small. After training this model for 70 epochs, the model achieves a holdout accuracy ranging from 76.0% ($N = 1000$) to 87.7% ($N = 10000$). This performance is noteworthy considering that in several cases it exceeds the best models with handcrafted features (cf. Table 2), which were designed with the true generative mechanism in mind. Indeed, considering the close alignment between handcrafted features and the generative mechanism, we suspect that the Bayes error for this problem is between 10 and 15%. The ability to match or exceed a handcrafted model underscores the appeal of modern deep learning. It is not necessarily about achieving better performance than interpretable counterparts on all problems; rather, it comes from the potential to bypasses feature engineering – a traditionally labor-intensive and time-consuming step – while achieving comparable performance. Moreover, this simulation highlights the attractive scaling of deep learning models in both out-of-sample accuracy and running time, since performance continues to improve with larger sample sizes. This makes them valuable in applications where frequent errors cannot be tolerated, even if they could be explained. This is often the case in technological systems that have to operate well even with limited human oversight, like flight control (Chai et al., 2021) or online recommender systems (Resnick and Varian, 1997). Indeed, in these settings, reliability and robustness may be more important than interpretability alone.

The remaining challenge lies in understanding how the models achieved their performances. We consider each of the techniques discussed above, focusing on the models trained with $N = 500$ samples. Fig 3 offers two global views of the data. Fig 3a gives the sparse PCA projections derived from the scaled N subjects \times (D taxa \times T timepoints) matrix of community profiles over time, while Fig 3b shows the analog for the N subjects \times L embeddings transformer representations. The embeddings in (b) more clearly distinguish healthy and diseased groups, capturing more of the total variance within the top two dimensions. This is expected, since the final classification uses a linear hyperplane, encouraging the transformer to learn a linearly separable representation. To better understand the features encoded by the transformer, we focus on the subset of embeddings associated with species 21. Fig 3c shows sparse PCA applied to these embeddings. We have highlighted samples that lie near the linear interpolation between samples 110 and 378 in the original embedding space; their original trajectories are shown in panel (d). Along this interpolation, the peaks and valleys become less pronounced, suggesting that the model has learned to differentiate curves with different trajectory properties. This is a subtle feature, and one that we know by design is related to the response. In contrast, some visually prominent features unrelated to disease, like the maximum value of the series, are appropriately disregarded.

We next study local explanations using IG. Based on the input structure, IG returns a value for each time point-by-taxon combination. Example values are given in Fig 4. In each panel, we can identify time points that informed the final classification and the direction in which perturbations change the target class label. Blue squares denote that an increased abundance at that time point increases the probability of correct classification; red circles suggests the opposite. Larger symbols imply larger changes in probability. For example, for subject 6, the first ten time points and the valley surrounding time 20 were most important. Increases in

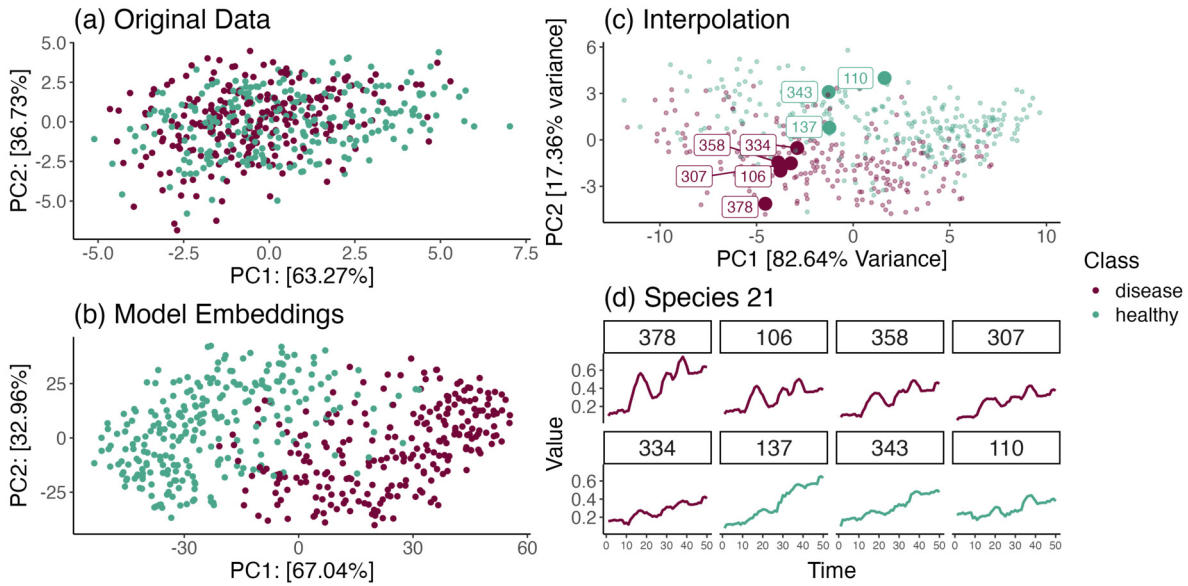


Figure 3: Global explanation of the simulation’s transformer model via embeddings. (a) Healthy and disease classes generally overlap in the original data. Each point gives the projection of a subject onto the top two sparse PCA directions. (b) Transformer representations more clearly separate the two classes. (c) A sequence of points along the linear interpolation between samples 110 and 378 in embedding space restricted to features from species 21. (d) Trajectories associated with each labeled point in (c).

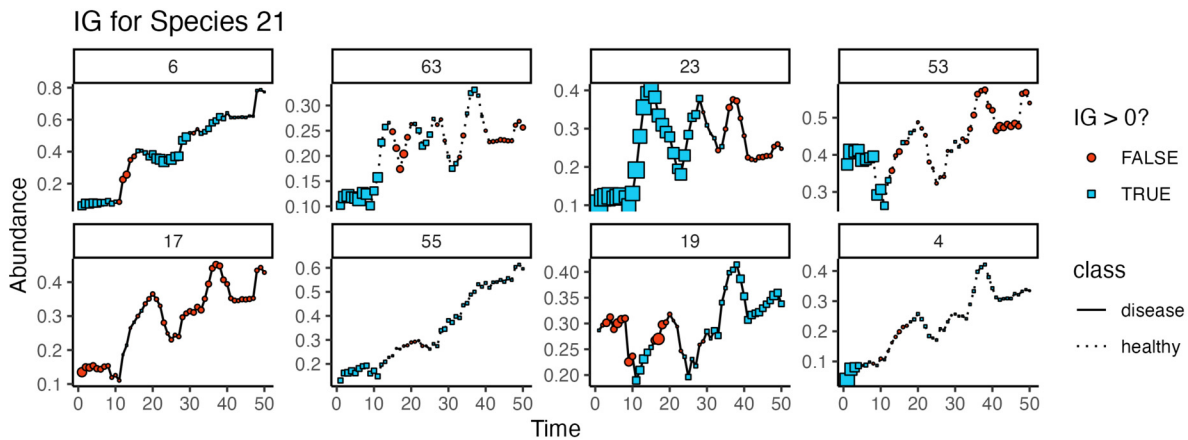


Figure 4: Integrated gradients for the transformer model overlaid on a subset of samples from the simulation study. Each curve represents the trajectory for species 21, with symbols indicating the size and sign of the integrated gradient at that sample and time point. The subset of relevant time points can vary from sample to sample.

either of these intervals increase the probability of correctness. In contrast, for subject 63, an increase in the valley surrounding time 20 decreases this probability. Since these subjects are from opposite classes, deeper valleys around time 20 seem to be related to health. However, beware that these conclusions are essentially local, and this pattern need not hold across all

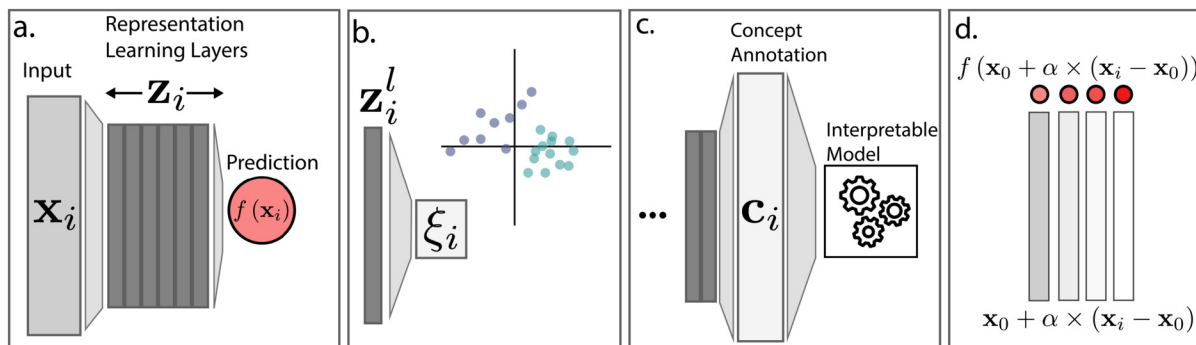


Figure 5: Summary of XAI techniques. (a) Deep learning models transform input \mathbf{x}_i into higher-level semantic representations. The final layer is used for prediction. (b) Global embeddings represent a layer of interest with lower-dimensional vectors \mathbf{z}_i , which can be interpreted alongside sample-level metadata, like point color in this example. (c) CBMs only use the black box to predict dense, concept-level annotations, not the final response. An interpretable model translates predicted concepts into the prediction (d) In IG, a sample is subject to a sequence of perturbations, and the impact on the final prediction is recorded to identify important coordinates.

samples for IG to declare it important. In general, IG effectively localizes the responsibility for a classified instance down to a subset of time points, aiding error analysis and uncovering subsets of features that could flip the class of a prediction.

Next, we take a closer look at the use of CBMs in our case study. Recall that concepts are a form of dense annotation for each sample. We extend our original generative mechanism to provide sample-level concepts. For each θ_{ik} , we define a binary vector by thresholding: $\mathbf{c}_i = \mathbf{1}\{\theta_{ik} > t\}$. Since each θ_{ik} corresponds to a latent trajectory, this acts as a coarse label for sample i , indicating which of the k underlying trajectories it is most influenced by. Next, we adapt our transformer to predict both the binary pattern of \mathbf{c}_i and the final disease status, adding a concept-level loss that computes a logistic loss coordinatewise (see Listing 1). The results are displayed in Table 2. Notably, using this concept definition, the model performs better than the original, unconstrained transformer, suggesting closer alignment to the original generative mechanism. In general, better CBMs require denser concept annotation. Note, however, that recent work has proposed variations that discover concepts de novo or identify cheap concept proxies following model training (Yuksekgonul et al., 2023; Ghorbani et al., 2019).

3 Evaluation

Developing effective benchmarks for interpretability methods is an active area of research, critical for objectively comparing proposals. However, designing effective criteria faces several challenges,

- **Operationalization:** Interpretability encompasses several technical definitions, like parsimony and simulatability, making it difficult to capture precisely with a single metric (Hedström et al., 2023b,a).
- **Output Types:** Different classes of methods return different types of outputs, like attribution maps and concepts, which are hard to compare directly.
- **Context Dependence:** Method performance must be gauged relative to the audience and tasks

they were designed for. Nonetheless, a good benchmark should generalize across tasks.

- **Ambiguous Ground Truth:** Unlike supervised learning, establishing ground truth interpretability output (e.g., the “correct” feature attribution map) is often impossible (Arras et al., 2022; Guidotti, 2021).
- **Human-Computer Interaction:** The success of an interpretability method depends on how it interfaces with users (Baniecki et al., 2021; Baniecki and Biecek, 2019). Simplified experiments may not reflect real-world complexity, while real-world observations may be confounded by factors that don’t generalize across applications.

These challenges mirror those encountered in data visualization, which has developed various evaluation protocols to guide systematic progress. Unlike traditional machine learning, where isolated computational benchmarks have led to significant progress, we expect evaluation in interpretability will depend on triangulating multiple sources of experimental evidence, like in data visualization.

Interpretability evaluation methods fall into two general categories: tailored benchmark datasets and user studies with controlled tasks. Within each class, there are variations that account for the challenges above in different ways. For dataset benchmarks, one approach is ablation (Hooker et al., 2019; Wang et al., 2022). Specifically, for methods that output feature importance attributions, this involves retraining models with versions of the data where purported important regions are masked out. Deterioration in model performance suggests that the attributions accurately reflect regions that the original model relied upon. Note that retraining is necessary. The ablated images are not drawn from the same distribution as the training data, so deterioration in performance of the original model could be due to either removal of important features (the criteria of interest) or domain shift (a confounding factor). Another data-driven benchmarking approach is to design synthetic datasets. These datasets are generated using known mechanisms, allowing for later evaluation of explanations (Liu et al., 2021; Zhou et al., 2022). For example, the funnybirds benchmark creates imaginary cartoon birds where “species” classes are derived from attributes, like wing, beak, and tail type, that are controlled during data generation (Hesse et al., 2023). Given a specific instance, an ideal explanation should focus on the parts of the bird that distinguish it from related classes. To make these benchmarks more realistic, variation in viewpoint and illumination can be introduced.

In user studies, one approach is to have participants make pairwise preference judgments between explanations (Jeyakumar et al., 2020). However, subjective preference may not align with objective task performance. To this end, more specific experimental tasks can be designed, evaluating user decisions with and without support from interpretability outputs (Ma et al., 2024; Colin et al., 2022). For example, to assess simulatability, participants could be asked to guess a model’s output on an example before and after viewing an explanation. Alternatively, they could be asked to guess how predictions might change under perturbations to the input – effective explanations should help in making accurate guesses. Finally, integrating interpretability methods into workflow-specific interfaces makes it possible to study their performance within tailored tasks. For example, Wu et al. (2021) designed an interface to aid auditors in generating counterfactual examples that flip model classifications. This is an instance of a counterfactual explanation (Mothilal et al., 2020; Guidotti, 2022) that helps to characterize sensitivity to out-of-distribution shifts. Influential examples that participants discovered could be incorporated into training to improve generalization. Different choices in the frontend interface and backend interpretability engine led to differential success in discovering surprising counterfactuals (those with small perturbations in input that led to large change in class predictions). In this way, interface-oriented user studies can more precisely gauge “human-AI collaboration.” Indeed,

these experiments have uncovered counterintuitive phenomena, like instances where task performance worsens with AI explanations — human “collaborators” can become less critical when shown an explanation (Bansal et al., 2021). Finally, a key challenge in user studies is selecting an appropriate participant pool with representative tasks (Bućinca et al., 2020). For example, an interpretability method in healthcare may be more effective for patients than for doctors. Narrow tasks can more directly guide deployment while broader ones are easier to run at scale.

There is often overlap between benchmark datasets and user study evaluation. For example, Casper et al. (2023) created a dataset to understand the effectiveness of explanations in aiding model debugging. The dataset introduced distractor patches to a small fraction of images, leading to failures on the test set. For example, a small strawberry was pasted into 1 in 3000 images in both training and test sets, and in each “poisoned” training (but not test) image, the class was switched to “goose,” leading to failures at test time. Participants were presented with misclassified test images, along with various explanations. Their accuracy and speed in identifying the poisoning mechanism served to quantify the utility of these explanations for debugging.

Evidently, while there is no one way to evaluate interpretability methods, there is a natural hierarchy of tests to which proposals can be subjected. General-purpose proposals should first pass ablation and synthetic data benchmarks. Success in these criteria justifies investing resources in user studies with more realistic audience and task designs (Rong et al., 2024). Future interpretability methods may be evaluated similarly to new therapies in medicine: *in silico* associations (ablation and synthetic data studies) can guide *in vitro* experiments (user studies with generic audiences) which, if successful, lead to follow-up clinical trials (user studies with representative audiences).

4 Discussion

This review has surveyed the significant advances in interpretability over recent decades. In closing, we highlight key ongoing shifts in the AI landscape and their implications for interpretability. One important trend is the rise of multimodal models, which are trained on mixed data modalities (Ngiam et al., 2011; Chen et al., 2021). This has facilitated novel applications, like automatic open vocabulary segmentation (Zou et al., 2023), where, rather than annotating each image pixel with one of a preset collection of class labels, the system can generate segmentation masks for arbitrary text strings (e.g., “The pedestrian that is crossing the street.”). While single modality methods remain applicable to individual model components (e.g., analyzing the word embeddings in a vision language model), there is a pressing need for interpretability techniques that offer an integrated view of how these models merge information across sources. This could help developers train models that effectively balance evidence from complementary views, enhancing overall system robustness. Statistical methods for data integration (Ding et al., 2022) could offer a principled approach for analyzing models with these more complex, interacting components.

Another significant trend is the emergence of foundation models (Bommasani et al., 2021). These models are trained on corpora spanning petabytes of data, demonstrate a capacity for in-context learning. They can immediately solve tasks that previously required additional labeled data and fine-tuning. In this sense, foundation models are generalists, allowing users across diverse domains to benefit from a single, albeit intensive, modeling investment. Therefore, understanding how to interpret these models and their role in downstream specialized tasks is an important area for study. It is no longer wise to explain AI systems as if they were trained in

isolation; they must be considered part of a broader AI ecosystem. Further, Xie et al. (2022) and Teh (2019) have pointed out connections between in-context learning and Bayesian posterior updating. This suggests that statistical thinking may resolve attribution questions when data are drawn from heterogeneous sources, which is important for resolving data provenance and ownership.

Third, as new legal and social norms around model-based decision-making are established, we will likely encounter scenarios where even comprehensive explanation are unsatisfactory, and where high-performing glass boxes become critical (Nannini et al., 2023). In certain applications, reliability and interpretability will outweigh performance alone. For example, in health-care and legal contexts, simple decision-making criteria are ideal, and in science, generalizable insight holds greater value than out-of-sample performance alone. Though recent progress in deep learning has sparked the community’s imagination about potentially more human-like or human-surpassing systems, we should perhaps consider futures where systems only marginally outperform current ones but are significantly safer and transparent.

Finally, we ask: how can we make interpretability a more interactive process? Interactivity is critical for ensuring that our systems augment intelligence, empowering users to become creative and critical problem solvers (Baniecki et al., 2023; Sokol and Flach, 2020). Neglecting this problem may lead to an overreliance on oracles, even among experts. For example, doctors should feel comfortable rejecting AI recommendations, even if there is a chance their decision turns out to be incorrect. A prerequisite for such interactivity is shared language (Heer, 2019; Crabbe et al., 2021; Kim, 2022) – it is difficult to interact without communication. CBMs are already a step in this direction, where concepts serve as a representation linking both human and machine reasoning. The goal of a shared language isn’t as fantastical as it might seem; programming, after all, is an exercise in communicating with machines, with the programming language as the shared representation.

Looking ahead, we imagine a dynamic where interpretability is central to the interactive design and development of AI systems. Returning to our visualization analogy, it is already common for national newspapers to feature sophisticated interactive visualizations, and college freshman across various disciplines are taught data literacy and visual exploration. In the future, interpretability can help AI models be communicated and developed in a similarly democratic way, making it easier for data to shape human judgment and creativity.

Supplementary Material

Code to reproduce our simulation experiment can be found at <https://go.wisc.edu/3k1ewe>. The repository’s README file describes how to generate the dataset, fit each model, and create the visualizations. The data, compiled notebooks, intermediate outputs, and log files can be accessed at <https://go.wisc.edu/v623lq>.

Acknowledgment

The author thanks Claire Donnat, John Cherian, Pratheepa Jeganathan, Laura Symul, Susan Holmes, and the organizers and participants of the workshop “Statistical Aspects of Trustworthy Machine Learning” for insightful comments on early versions of this work. He also thanks two reviewers and the editor for suggestions that improved the manuscript.

Funding

This work was supported in part by grant number R01GM152744 from the National Institute of General Medical Sciences.

References

- Achtibat R, Dreyer M, Eisenbraun I, Bosse S, Wiegand T, Samek W, et al. (2023). From attribution maps to human-understandable explanations through concept relevance propagation. *Nature Machine Intelligence*, 5(9): 1006–1019. <https://doi.org/10.1038/s42256-023-00711-8>
- Adebayo J, Gilmer J, Muelly M, Goodfellow I, Hardt M, Kim B (2018). Sanity checks for saliency maps. In: *Proceedings of the 32nd International Conference on Neural Information Processing Systems, NIPS'18*, (S Bengio H Wallach, H Larochelle, K Grauman, N Cesa-Bianchi, eds.), 9525–9536. Curran Associates Inc., Red, Hook, NY, USA.
- Agrawala M, Li W, Berthouzoz F (2011). Design principles for visual communication. *Communications of the ACM*, 54(4): 60–69. <https://doi.org/10.1145/1924421.1924439>
- Arras L, Osman A, Samek W (2022). CLEVR-XAI: A benchmark dataset for the ground truth evaluation of neural network explanations. *Information Fusion*, 81: 14–40. <https://doi.org/10.1016/j.inffus.2021.11.008>
- Baniecki H, Biecek P (2019). modelstudio: Interactive studio with explanations for ml predictive models. *The Journal of Open Source Software*, 4(43): 1798. <https://doi.org/10.21105/joss.01798>
- Baniecki H, Kretowicz W, Piątysek P, Wiśniewski J, Biecek P (2021). dalex: Responsible machine learning with interactive explainability and fairness in python. *Journal of Machine Learning Research*, 22(214): 1–7.
- Baniecki H, Parzych D, Biecek P (2023). The grammar of interactive explanatory model analysis. *Data Mining and Knowledge Discovery*, 38: 1–37.
- Bansal G, Wu T, Zhou J, Fok R, Nushi B, Kamar E, et al. (2021). Does the whole exceed its parts? The effect of AI explanations on complementary team performance. In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems, CHI '21*. (Y Kitamura, A Quigley, K Isbister, T Igarashi, P Bjorn, S Drucker, eds.), ACM.
- Bengio Y (2009). *Learning Deep Architectures for AI*. NOW.
- Bengio Y, Courville A, Vincent P (2013). Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8): 1798–1828. <https://doi.org/10.1109/TPAMI.2013.50>
- Bilodeau B, Jaques N, Koh PW, Kim B (2024). Impossibility theorems for feature attribution. *Proceedings of the National Academy of Sciences*, 121(2). <https://doi.org/10.1073/pnas.2304406120>
- Bommasani R, Hudson DA, Adeli E, Altman R, Arora S, von Arx S, et al. (2021). On the opportunities and risks of foundation models.
- Borkin MA, Vo AA, Bylinskii Z, Isola P, Sunkavalli S, Oliva A, et al. (2013). What makes a visualization memorable? *IEEE Transactions on Visualization and Computer Graphics*, 19(12): 2306–2315. <https://doi.org/10.1109/TVCG.2013.234>
- Buçinca Z, Lin P, Gajos KZ, Glassman EL (2020). Proxy tasks and subjective measures can be misleading in evaluating explainable AI systems. In: *Proceedings of the 25th International Conference on Intelligent User Interfaces, IUI '20*, (F Paternò, N Oliver, C Conati, L Spano,

- N Tintarev, eds.). ACM.
- Bykov K, Kopf L, Nakajima S, Kloft M, Höhne MMC (2023). Labeling neural representations with inverse recognition.
- Caruana R, Lou Y, Gehrke J, Koch P, Sturm M, Elhadad N (2015). Intelligible models for healthcare: Predicting pneumonia risk and hospital 30-day readmission. In: *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '15*, (L Cao, C Zhang, T Joachims, G Webb, D D Margineantu, G Williams, eds), 1721–1730. Association for Computing Machinery, New York, NY, USA.
- Casper S, Bu T, Li Y, Li J, Zhang K, Hariharan K, et al. (2023). Red teaming deep neural networks with feature synthesis tools. In: *Advances in Neural Information Processing Systems* (A Oh, T Naumann, A Globerson, K Saenko, M Hardt, S Levine, eds.), volume 36, 80470–80516. Curran Associates, Inc.
- Chai R, Tsourdos A, Savvaris A, Chai S, Xia Y, Philip Chen C (2021). Review of advanced guidance and control algorithms for space/aerospace vehicles. *Progress in Aerospace Sciences*, 122: 100696. <https://doi.org/10.1016/j.paerosci.2021.100696>
- Chen S, Guhur PL, Schmid C, Laptev I (2021). History aware multimodal transformer for vision-and-language navigation. In: *Advances in Neural Information Processing Systems* (M Ranzato, A Beygelzimer, Y Dauphin, PS Liang, JW Vaughan, eds.), volume 34, 5834–5847. Curran Associates, Inc.
- Cleveland WS (1993). *Visualizing Data*. Hobart Press.
- Coenen A, Reif E, Yuan A, Kim B, Pearce A, Viégas F, et al. (2019). *Visualizing and Measuring the Geometry of BERT*. Curran Associates Inc., Red, Hook, NY, USA.
- Colin J, T Cadene R FEL, Serre T (2022). What I cannot predict, I do not understand: A human-centered evaluation framework for explainability methods. In: *Advances in Neural Information Processing Systems* (S Koyejo, S Mohamed, A Agarwal, D Belgrave, K Cho, A Oh, eds.), volume 35, 2832–2845. Curran Associates, Inc.
- Crabbe J, Qian Z, Imrie F, van der Schaar M (2021). Explaining latent representations with a corpus of examples. In: *Advances in Neural Information Processing Systems* (M Ranzato, A Beygelzimer, Y Dauphin, PS Liang, JW Vaughan, eds.), volume 34, 12154–12166. Curran Associates, Inc.
- Dandl S, Casalicchio G, Bischl B, Bothmann L (2023). Interpretable regional descriptors: Hyperbox-based local explanations. In: *Machine Learning and Knowledge Discovery in Databases: Research Track, Lecture Notes in Computer Science*, (D Koutra, C Plant, M Gomez Rodriguez, E Baralis, F Bonchi, eds.), 479–495. Springer Nature, Switzerland, Cham.
- Devlin J, Chang MW, Lee K, Toutanova K (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In: *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)* (J Burstein, C Doran, T Solorio, eds.), 4171–4186. Association for Computational Linguistics, Minneapolis, Minnesota.
- Ding DY, Li S, Narasimhan B, Tibshirani R (2022). Cooperative learning for multiview analysis. *Proceedings of the National Academy of Sciences*, 119(38).
- Erhan D, Bengio Y, Courville A, Manzagol PA, Vincent P, Bengio S (2010). Why does unsupervised pre-training help deep learning? *Journal of Machine Learning Research*, 11(19): 625–660.
- Fridovich-Keil S, Yu A, Tancik M, Chen Q, Recht B, Kanazawa A (2022). Plenoxels: Radiance fields without neural networks. In: *2022 IEEE/CVF Conference on Computer Vision and*

- Pattern Recognition (CVPR)*, (R Chellappa, J Matas, L Quan, M Shah, eds). IEEE.
- Friedman J, Hastie T, Tibshirani R (2010). Regularization paths for generalized linear models via coordinate descent. *Journal of Statistical Software*, 33(1): 1–22. <https://doi.org/10.18637/jss.v033.i01>
- Friedman JH (2001). Greedy function approximation: A gradient boosting machine. *The Annals of Statistics*, 29(5). 1189–1232. <https://doi.org/10.1214/aos/1013203451>
- Fumagalli F, Muschalik M, Kolpaczki P, Hüllermeier E, Hammer B (2023). SHAP-IQ: Unified approximation of any-order shapley interactions.
- Ghorbani A, Wexler J, Zou J, Kim B (2019). *Towards Automatic Concept-Based Explanations*. Curran Associates Inc., Red, Hook, NY, USA.
- Gosmann C, Anahtar MN, Handley SA, Farcasanu M, Abu-Ali G, Bowman BA, et al. (2017). Lactobacillus-deficient cervicovaginal bacterial communities are associated with increased HIV acquisition in young South African women. *Immunity*, 46(1): 29–37. <https://doi.org/10.1016/j.immuni.2016.12.013>
- Gu T, Liu K, Dolan-Gavitt B, Garg S (2019). Badnets: Evaluating backdooring attacks on deep neural networks. *IEEE Access*, 7: 47230–47244. <https://doi.org/10.1109/ACCESS.2019.2909068>
- Guidotti R (2021). Evaluating local explanation methods on ground truth. *Artificial Intelligence*, 291(103428): 103428. <https://doi.org/10.1016/j.artint.2020.103428>
- Guidotti R (2022). Counterfactual explanations and how to find them: literature review and benchmarking. *Data Mining and Knowledge Discovery*. 28: 2770–2770.
- Hazimeh H, Ponomareva N, Mol P, Tan Z, Mazumder R (2020). The tree ensemble layer: Differentiability meets conditional computation. In: *Proceedings of the 37th International Conference on Machine Learning, ICML '20*, (H Daumé, A Singh, eds.). [JMLR.org](https://doi.org/10.1145/3458456).
- Hedström A, Bommer P, Samek W W, Lapuschkin S, Höhne MMC (2023a). The meta-evaluation problem in explainable AI: Identifying reliable estimators with MetaQuantus.
- Hedström A, Weber L, Krakowczyk D, Bareeva D, Motzkus F, Samek W, et al. (2023b). Quantus: An explainable AI toolkit for responsible evaluation of neural network explanations and beyond. *Journal of Machine Learning Research*, 24(34): 1–11.
- Heer J (2019). Agency plus automation: Designing artificial intelligence into interactive systems. *Proceedings of the National Academy of Sciences*, 116(6): 1844–1850. <https://doi.org/10.1073/pnas.1807184115>
- Herbinger J, Bischl B, Casalicchio G (2022). REPID: Regional effect plots with implicit interaction detection. 10209–10233.
- Hesse R, Schaub-Meyer S, Roth S (2023). FunnyBirds: A synthetic vision dataset for a part-based analysis of explainable AI methods. In: *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, (J Kosecka, J Ponce, C Schmid, A Zisserman, eds.). IEEE.
- Hewitt J, Manning CD (2019). A structural probe for finding syntax in word representations. In: *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)* (J Burstein, C Doran, T Solorio, eds.), 4129–4138. Association for Computational Linguistics, Minneapolis, Minnesota.
- Holmes S (2017). Statistical proof? The problem of irreproducibility. *Bulletin of the American Mathematical Society*, 55(1): 31–55. <https://doi.org/10.1090/bull/1597>
- Hooker S, Erhan D, Kindermans PJ, Kim B (2019). A benchmark for interpretability methods in deep neural networks. In: *Advances in Neural Information Processing Systems* (H Wal-

- lach, H Larochelle, A Beygelzimer, F Alché-Buc, E Fox, R Garnett, eds.), volume 32. Curran Associates, Inc.
- Huber PJ (1964). Robust estimation of a location parameter. *The Annals of Mathematical Statistics*, 35(1): 73–101. <https://doi.org/10.1214/aoms/1177703732>
- Hutchins EL, Hollan JD, Norman DA (1985). Direct manipulation interfaces. *Human-Computer Interaction*, 1(4): 311–338. https://doi.org/10.1207/s15327051hci0104_2
- Jeganathan P, Callahan BJ, Proctor DM, Relman DA, Holmes SP (2018). The block bootstrap method for longitudinal microbiome data.
- Jeyakumar JV, Noor J, Cheng YH, Garcia L, Srivastava M (2020). How can I explain this to you? An empirical study of deep neural network explanation methods. In: *Advances in Neural Information Processing Systems*, volume 33, (H Larochelle, M Ranzato, R Hadsell, MF Balcan, H Lin, eds.), 4211–4222. Curran Associates, Inc.
- Karimi AH, Barthe G, Schölkopf B, Valera I (2022). A survey of algorithmic recourse: Contrastive explanations and consequential recommendations. *ACM Computing Surveys*, 55(5): 1–29.
- Kim B (2022). Beyond Interpretability: Developing a Language to Shape Our Relationships with AI.
- Kodikara S, Ellul S, Cao KA L (2022). Statistical challenges in longitudinal microbiome data analysis. *Briefings in Bioinformatics*, 23(4). <https://doi.org/10.1093/bib/bbac273>
- Koh PW, Liang P (2017). Understanding black-box predictions via influence functions. In: *Proceedings of the 34th International Conference on Machine Learning - Volume 70, ICML'17*, (D Precup, Y.W Teh, eds.), 1885–1894. [JMLR.org](http://jmlr.org).
- Koh PW, Nguyen T, Tang YS, Mussmann S, Pierson E, Kim B, et al. (2020). Concept bottleneck models.
- Kolesnikov A, Dosovitskiy A, Weissenborn D, Heigold G, Uszkoreit J, Beyer L, et al. (2021). An image is worth 16x16 words: Transformers for image recognition at scale.
- Kostic AD, Gevers D, Siljander H, Vatanen T, Hyötyläinen T, Hämäläinen AM, et al. (2015). The dynamics of the human infant gut microbiome in development and in progression toward type 1 diabetes. *Cell Host and Microbe*, 17(2): 260–273. <https://doi.org/10.1016/j.chom.2015.01.001>
- Krishnan M (2019). Against interpretability: A critical examination of the interpretability problem in machine learning. *Philosophy and Technology*, 33(3): 487–502. <https://doi.org/10.1007/s13347-019-00372-9>
- Kundaliya D (2023). Computing - incisive media: Google AI chatbot bard gives wrong answer in its first demo. *Computing*. Nom - OpenAI; Copyright - Copyright Newstex Feb 9, 2023; Dernière mise à jour - 2023-11-30.
- Lee JS (2021). Transformers: a Primer.
- Lipton ZC (2018). The mythos of model interpretability: In machine learning, the concept of interpretability is both important and slippery. *Queue*, 16(3): 31–57. <https://doi.org/10.1145/3236386.3241340>
- Liu Y, Khandagale S, White C, Neiswanger W (2021). Synthetic benchmarks for scientific research in explainable machine learning.
- Loh W (2014). Fifty years of classification and regression trees. *International Statistical Review*, 82(3): 329–348. <https://doi.org/10.1111/insr.12016>
- Lundberg SM, Erion G, Chen H, DeGrave A, Prutkin JM, Nair B, et al. (2020). From local explanations to global understanding with explainable AI for trees. *Nature Machine Intelligence*, 2(1): 56–67. <https://doi.org/10.1038/s42256-019-0138-9>
- Lundberg SM, Lee SI (2017). A unified approach to interpreting model predictions. In: *Proceed-*

- ings of the 31st International Conference on Neural Information Processing Systems, NIPS'17*, (U. von Luxburg, I Guyon, S Bengio, H Wallach, R Fergus, eds.), 4768–4777. Curran Associates Inc., Red, Hook, NY, USA.
- Lundstrom DD, Huang T, Razaviyayn M (2022). A rigorous study of integrated gradients method and extensions to internal neuron attributions. In: *Proceedings of the 39th International Conference on Machine Learning* (K Chaudhuri, S Jegelka, L Song, C Szepesvari, G Niu, S Sabato, eds.), volume 162 of *Proceedings of Machine Learning Research*, 14485–14508. PMLR.
- Ma J, Lai V, Zhang Y, Chen C, Hamilton P, Ljubenkov D, et al. (2024). OpenHEXAI: An open-source framework for human-centered evaluation of explainable machine learning.
- Mehrabi N, Morstatter F, Saxena N, Lerman K, Galstyan A (2021). A survey on bias and fairness in machine learning. *ACM Computing Surveys*, 54(6): 1–35. <https://doi.org/10.1145/3457607>
- Mothilal RK, Sharma A, Tan C (2020). Explaining machine learning classifiers through diverse counterfactual explanations. In: *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, (M Hildebrandt, C Castillo, E Celis, S Ruggieri, L Taylor, G Zanfir-Fortuna, eds.), ACM, New York, NY, USA.
- Murdoch WJ, Singh C, Kumbier K, Abbasi-Asl R, Yu B (2019). Definitions, methods, and applications in interpretable machine learning. *Proceedings of the National Academy of Sciences*, 116(44): 22071–22080. <https://doi.org/10.1073/pnas.1900654116>
- Nannini L, Balayn A, Smith AL (2023). Explainability in AI policies: A critical review of communications, reports, regulations, and standards in the EU, US, and UK. In: *2023 ACM Conference on Fairness, Accountability, and Transparency*. ACM, New York, NY, USA.
- Ngiam J, Khosla A, Kim M, Nam J, Lee H, Ng AY (2011). Multimodal deep learning. In: *Proceedings of the 28th International Conference on International Conference on Machine Learning, ICML'11*, (L Getoor, T Scheffer, eds.), 689–696. Omnipress, Madison, WI, USA.
- Nguyen LH, Holmes S (2019). Ten quick tips for effective dimensionality reduction. *PLoS Computational Biology*, 15(6): e1006907. <https://doi.org/10.1371/journal.pcbi.1006907>
- Nussberger AM, Luo L, Celis LE, Crockett MJ (2022). Public attitudes value interpretability but prioritize accuracy in artificial intelligence. *Nature Communications*, 13(1): 5821.
- Oppermann M, Munzner T (2022). Vizsnippets: Compressing visualization bundles into representative previews for browsing visualization collections. *IEEE Transactions on Visualization and Computer Graphics*, 28(1): 747–757. <https://doi.org/10.1109/TVCG.2021.3114841>
- Pennington J, Socher R, Manning C (2014). Glove: Global vectors for word representation. In: *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, (A Moschitti, B Pang, W Daelemans, eds.), Association for Computational Linguistics.
- Poursabzi-Sangdeh F, Goldstein DG, Hofman JM, Wortman Vaughan JW, Wallach H (2021). Manipulating and measuring model interpretability. In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, (Y Kitamura, A Quigley, K Isbister, T Igarashi, P Bjørn, S Drucker, eds.), ACM, New York, NY, USA.
- Raghu M, Unterthiner T, Kornblith S, Zhang C, Dosovitskiy A (2021). Do vision transformers see like convolutional neural networks? In: *Advances in Neural Information Processing Systems* (M Ranzato, A Beygelzimer, Y Dauphin, PS Liang, JW Vaughan, eds.), volume 34, 12116–12128. Curran Associates, Inc.
- Raghu M, Zhang C, Kleinberg J, Bengio S (2019). *Transfusion: Understanding Transfer Learning for Medical Imaging*. Curran Associates Inc., Red, Hook, NY, USA.
- Resnick P, Varian HR (1997). Recommender systems. *Communications of the ACM*, 40(3):

- 56–58. <https://doi.org/10.1145/245108.245121>
- Ribeiro MT, Singh S, Guestrin C (2016). “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, KDD '16*, (B Krishnapuram, M Shah, A Smola, C Aggarwal, D Shen, R Rastogi, eds.), 1135–1144. Association for Computing Machinery, New York, NY, USA.
- Rong Y, Leemann T, Nguyen TT, Fiedler L, Qian P, Unhelkar V, et al. (2024). Towards human-centered explainable AI: A survey of user studies for model explanations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(4): 2104–2122. <https://doi.org/10.1109/TPAMI.2023.3331846>
- Ross AS, Hughes MC, Doshi-Velez F (2017). Right for the right reasons: Training differentiable models by constraining their explanations. In: *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence. International Joint Conferences on Artificial Intelligence Organization*, (C Sierra, ed.), California.
- Rudin C (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5): 206–215. <https://doi.org/10.1038/s42256-019-0048-x>
- Sankaran K, Holmes SP (2023). Generative models: An interdisciplinary perspective. *Annual Review of Statistics and Its Application*, 10(1): 325–352. <https://doi.org/10.1146/annurev-statistics-033121-110134>
- Sedlmair M, Meyer M, Munzner T (2012). Design study methodology: Reflections from the trenches and the stacks. *IEEE Transactions on Visualization and Computer Graphics*, 18(12): 2431–2440. <https://doi.org/10.1109/TVCG.2012.213>
- Selvaraju RR, Cogswell M, Das A, Vedantam R, Parikh D, Batra D (2017). Grad-cam: Visual explanations from deep networks via gradient-based localization, (K Ikeuchi, G Medioni, M Pelillo, eds.), In: *2017 IEEE International Conference on Computer Vision (ICCV)*, 618–626.
- Silverman JD, Shenhav L, Halperin E, Mukherjee S, David LA (2018). Statistical considerations in the design and analysis of longitudinal microbiome studies.
- Simonyan K, Vedaldi A, Zisserman A (2014). Deep inside convolutional networks: Visualising image classification models and saliency maps. In: *Workshop at International Conference on Learning Representations*.
- Sokol K, Flach P (2020). One explanation does not fit all. *KI. Künstliche Intelligenz*, 34(2): 235–250. <https://doi.org/10.1007/s13218-020-00637-y>
- Sturmfels P, Lundberg S, Lee SI (2020). Visualizing the impact of feature attribution baselines. *Distill*, 5(1). <https://doi.org/10.23915/distill.00022>
- Sundararajan M, Taly A, Yan Q (2017). Axiomatic attribution for deep networks. In: *Proceedings of the 34th International Conference on Machine Learning - Volume 70, ICML'17*, (D Precup, YW Teh), 3319–3328. [JMLR.org](https://jmlr.org).
- Teh YW (2019). On Statistical Thinking in Deep Learning A Blog Post. *IMS Medallion Lecture*.
- Tufte ER (2001). *The Visual Display of Quantitative Information*, 2 edition. Graphics Press, Cheshire, CT.
- Tukey JW (1959). *A Survey of Sampling from Contaminated Distributions*. volume 33 of *STRG Technical report*, 57.
- Van Boxtel GJM, et al. (2021). *gsignal: Signal Processing*.
- Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, et al. (2017). Attention is all you need. In: *Proceedings of the 31st International Conference on Neural Information*

- Processing Systems, NIPS'17*, (U von Luxburg, I Guyon, S Bengio, H Wallach, R Fergus, eds.), 6000–6010. Curran Associates Inc., Red, Hook, NY, USA.
- Wang L, Shen Y, Peng S, Zhang S, Xiao X, Liu H, et al. (2022). A fine-grained interpretability evaluation benchmark for neural NLP. In: *Proceedings of the 26th Conference on Computational Natural Language Learning (CoNLL)*, (A Fokkens, V Srikumar, eds.), Association for Computational Linguistics, Stroudsburg, PA, USA.
- Williamson B, Feng J (2020). Efficient nonparametric statistical inference on population feature importance using Shapley values. In: *Proceedings of the 37th International Conference on Machine Learning (HD III (A Singh, ed.))*, volume 119 of *Proceedings of Machine Learning Research*, 10282–10291. PMLR.
- Wongsuphasawat K, Smilkov D, Wexler J, Wilson J, Mané D, Fritz D, et al. (2018). Visualizing dataflow graphs of deep learning models in tensorflow. *IEEE Transactions on Visualization and Computer Graphics*, 24: 1–12. <https://doi.org/10.1109/TVCG.2017.2744878>
- Wu T, Ribeiro MT, Heer J, Weld D (2021). Polyjuice: Generating counterfactuals for explaining, evaluating, and improving models. In: *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, (C Zong, F Xia, W Li, R Navigli, eds.), Association for Computational Linguistics.
- Xie SM, Raghunathan A, Liang P, Ma T (2022). An explanation of in-context learning as implicit Bayesian inference. In: *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. [OpenReview.net](https://openreview.net).
- Xin R, Zhong C, Chen Z, Takagi T, Seltzer M, Rudin C (2022). Exploring the whole rashomon set of sparse decision trees, arXiv preprint: <https://arxiv.org/abs/2209.08040>.
- Yosinski J, Clune J, Nguyen AM, Fuchs TJ, Lipson H (2015). Understanding neural networks through deep visualization, arXiv preprint: <https://arxiv.org/abs/1506.06579>.
- Yuksekgonul M, Wang M, Zou J (2023). Post-hoc concept bottleneck models. In: *The Eleventh International Conference on Learning Representations*, (H Wang, W Lin, H He, D Wang, C Mao, M Zhang, eds.).
- Zeiler MD, Fergus R (2013). Visualizing and understanding convolutional networks, arXiv preprint: <https://arxiv.org/abs/1311.2901>.
- Zeng J, Ustun B, Rudin C (2016). Interpretable classification models for recidivism prediction. *Journal of the Royal Statistical Society. Series A. Statistics in Society*, 180(3): 689–722. <https://doi.org/10.1111/rssa.12227>
- Zhong C, Chen Z, Liu J, Seltzer M, Rudin C (2023). Exploring and interacting with the set of good sparse generalized additive models.
- Zhou Y, Booth S, Ribeiro MT, Shah J (2022). Do feature attribution methods correctly attribute features? *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(9): 9623–9633.
- Zou X, Yang J, Zhang H, Li F, Li L, Wang J, et al. (2023). Segment everything everywhere all at once. In: *Advances in Neural Information Processing Systems* (A Oh, T Naumann, A Globerson, K Saenko, M Hardt, S Levine, eds.), volume 36, 19769–19782. Curran Associates, Inc.