

Blesch, Kristin; Watson, David S.; Wright, Marvin N.

**Article — Published Version**

## Conditional feature importance for mixed data

AStA Advances in Statistical Analysis

**Provided in Cooperation with:**

Springer Nature

*Suggested Citation:* Blesch, Kristin; Watson, David S.; Wright, Marvin N. (2023) : Conditional feature importance for mixed data, AStA Advances in Statistical Analysis, ISSN 1863-818X, Springer, Berlin, Heidelberg, Vol. 108, Iss. 2, pp. 259-278,  
<https://doi.org/10.1007/s10182-023-00477-9>

This Version is available at:

<https://hdl.handle.net/10419/308997>

### Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

### Terms of use:

*Documents in EconStor may be saved and copied for your personal and scholarly purposes.*

*You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.*

*If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.*



<https://creativecommons.org/licenses/by/4.0/>



# Conditional feature importance for mixed data

Kristin Blesch<sup>1,2</sup> · David S. Watson<sup>3</sup> · Marvin N. Wright<sup>1,2,4</sup>

Received: 6 October 2022 / Accepted: 28 March 2023 / Published online: 29 April 2023  
© The Author(s) 2023

## Abstract

Despite the popularity of feature importance (FI) measures in interpretable machine learning, the statistical adequacy of these methods is rarely discussed. From a statistical perspective, a major distinction is between analysing a variable's importance before and after adjusting for covariates—i.e., between *marginal* and *conditional* measures. Our work draws attention to this rarely acknowledged, yet crucial distinction and showcases its implications. We find that few methods are available for testing conditional FI and practitioners have hitherto been severely restricted in method application due to mismatched data requirements. Most real-world data exhibits complex feature dependencies and incorporates both continuous and categorical features (i.e., mixed data). Both properties are oftentimes neglected by conditional FI measures. To fill this gap, we propose to combine the conditional predictive impact (CPI) framework with sequential knockoff sampling. The CPI enables conditional FI measurement that controls for any feature dependencies by sampling valid knockoffs—hence, generating synthetic data with similar statistical properties—for the data to be analysed. Sequential knockoffs were deliberately designed to handle mixed data and thus allow us to extend the CPI approach to such datasets. We demonstrate through numerous simulations and a real-world example that our proposed workflow controls type I error, achieves high power, and is in-line with results given by other conditional FI measures, whereas marginal FI metrics can result in misleading interpretations. Our findings highlight the necessity of developing statistically adequate, specialized methods for mixed data.

**Keywords** Interpretable machine learning · Feature importance · Knockoffs · Explainable artificial intelligence

---

✉ Kristin Blesch  
[blesch@leibniz-bips.de](mailto:blesch@leibniz-bips.de)

Extended author information available on the last page of the article

## 1 Introduction

Interpretable machine learning is on the rise as practitioners become interested in not only achieving high prediction accuracy in supervised learning tasks, but also understanding why certain predictions were made. Evaluating the importance of input variables (features) to the target prediction plays a crucial role in facilitating such endeavours. Several feature importance (FI) measures have been proposed by the machine learning community, but differing conceptualizations are spread across the literature.

We identify at least five dichotomies that orient FI methods: (1) global vs. local; (2) model-agnostic vs. model-specific; (3) testing vs. scoring; (4) methods that do and do not accommodate mixed tabular data; and (5) conditional vs. marginal measures. This defines a grid with  $2^5 = 32$  cells that helps categorize FI measures. For example, the popular SHAP algorithm (Lundberg and Lee 2017) produces local, model-agnostic FI scores that can accommodate mixed data and measures marginal FI. We emphasize that there is no “ideal” configuration of these five options—each is the right answer to a different question that is irreducibly context-dependent. However, this grid helps identify a notable lacuna: There are few global, model-agnostic FI methods that accommodate mixed data with error control for conditional FI measurement.

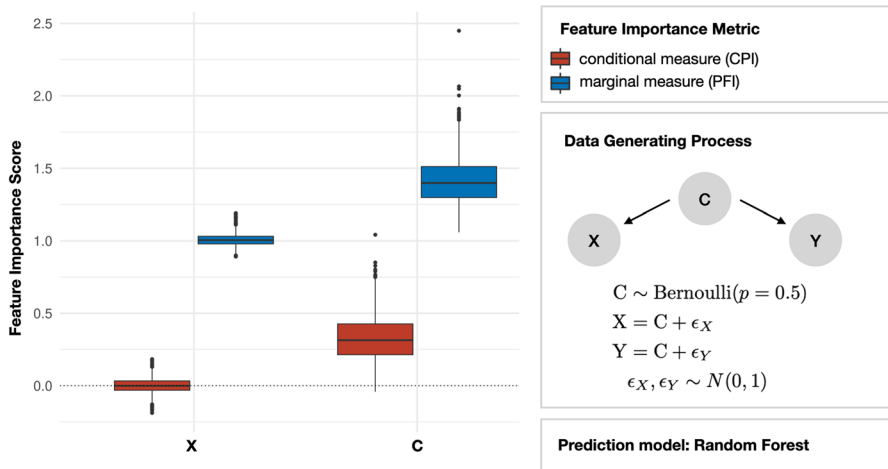
Explaining the dichotomies in more detail, local FI measures (Lundberg and Lee 2017; Ribeiro et al. 2016) are optimized for a particular point or region of the feature space, e.g., a single observation, while global FI scores (Fisher et al. 2019; Friedman 2001) measure a variable’s overall importance. Model-specific measures (Breiman 2001; Kursa and Rudnicki 2010; Shrikumar et al. 2017) exploit the properties of a particular function class for more efficient or precise FI calculation, while model-agnostic measures (Apley and Zhu 2020; Ribeiro et al. 2018) treat the underlying model as a black box. Testing methods include some inference procedure for error control (Lei et al. 2018), while scoring methods (Covert et al. 2020) do not. Some methods are proposed with limited applicability to certain data types, e.g. only continuous inputs (Watson and Wright 2021), while others are more flexible (Molnar et al. 2023). We discuss a selection of FI methods briefly in Sect. 2, but refer readers to review papers on FI interpretability methods, e.g. Linardatos et al. (2021), for a wider discussion on the topic.

Through the lens of statistics, the division (5), conditional vs. marginal measures, is particularly important, yet insufficiently acknowledged in both literature and practice (Apley and Zhu 2020; Hooker et al. 2021; Molnar et al. 2023; Watson and Wright 2021). The complementary concepts become evident when relating the statistical conception of independence testing to the machine learning view on FI measurement. We can think of the marginal null hypothesis as testing whether the input feature  $X_j$  is independent of other covariates  $X_{-j}$  or the target variable  $Y$ :

$$H_0^M : X_j \perp\!\!\!\perp \{Y, X_{-j}\} \quad (1)$$

On the other hand, testing against (2) accounts for the covariates  $X_{-j}$  and hence corresponds to conditional FI:

## Marginal versus conditional feature importance (FI) measurement



**Fig. 1** Boxplots contrasting marginal and conditional FI metrics for a prediction of  $Y$  with  $C$  and  $X$  ( $N = 200$ ) through a random forest prediction model across 1000 replicates. The conditional FI measure attributes no importance to  $X$ , whereas the marginal measure attributes nonzero importance to  $X$  because (due to induced correlation between  $X$  and  $Y$  by  $C$ ) it is predictive of  $Y$

$$H_0^C : X_j \perp\!\!\!\perp Y \mid X_{-j} \quad (2)$$

These tests clearly target different objectives. In this setup, we have  $H_0^M$  entailing  $H_0^C$ , but not the other way around. However, this strength comes with a certain loss of specificity, because rejecting  $H_0^M$  leaves it unclear whether  $X_j$  is correlated with  $Y$ ,  $X_{-j}$ , or both.

The relationship between FI and independence testing sheds light on another aspect, which may even be considered another dichotomy: does the FI measure aim for investigating the model behaviour or the underlying data structure (Chen et al. 2020)? For example, conditional independence tests that are part of some conditional FI measures (Watson and Wright 2021) may be used for causal structure learning, which often is based on repeated conditional independence testing (Glymour et al. 2019). Therefore, conditional FI measures can help explain the underlying data structure, whereas marginal FI measures differentiate between variables the predictive model relies on, which can be used to evaluate the fairness of a model. This does not preclude practitioners from using marginal and conditional FI measures in conjunction, and since marginal measures are often faster to compute, they might be preferable for quick assessments in large pipelines with many iterations. However, practitioners must be careful to interpret these measures properly and not infer a conditional signal from a marginal test.

In Fig. 1, we illustrate the difference between marginal (permutation feature importance (PFI), Fisher et al. 2019, Breiman 2001) and conditional (conditional predictive impact with Gaussian knockoffs (CPIgauss), Watson and Wright 2021) FI

measures. In this example, the confounding variable  $C$  is a common cause of both  $X$  and  $Y$ . This causal structure induces spurious correlation between  $X$  and  $Y$ , leading the marginal FI measure to attribute nonzero importance values to both  $C$  and  $X$  in predicting  $Y$ . On the contrary, the conditional FI measure attributes nonzero FI only to  $C$ , since  $X$  has no *additional* predictive value for  $Y$  above  $C$ .

This paper explores global, model-agnostic FI methods that accommodate mixed data with error control for conditional FI measurement. This is not a niche problem: mixed tabular data is the norm in many important areas such as health care, economics, and industry, and inference procedures are essential for decision-making in high risk domains to minimize costly errors. With the proliferation of machine learning algorithms, model-agnostic approaches can help standardize FI tasks without recalibrating to a particular function class for each new application. Conditional, global measures are valuable when practitioners seek mechanistic understanding that takes data covariance into account and go beyond individual model outputs.

Even though the empirical relevance of this kind of FI measurement is eminent, specialized methods are lacking. Some FI methods have yet to be evaluated in mixed data settings (Covert et al. 2020; Molnar et al. 2023; Lei et al. 2018), while others are currently inapplicable to mixed data (Watson and Wright 2021). The consequences of neglecting the special nature of mixed data for conditional FI measurement remain unexplored, and therefore practitioners currently have no guidance on how to proceed with conditional FI measurement in such cases, which proves a severe limitation in real-world applications.

We propose to combine the *conditional predictive impact* (CPI) testing framework proposed by Watson and Wright (2021) with the use of *sequential knockoffs* (Kormaksson et al. 2021) in order to enable conditional, global, model-agnostic FI testing for mixed data. CPI is a flexible, model-agnostic tool that relies on the usage of so-called knockoffs (Candès et al. 2018). In short, knockoffs are synthetic variables that carry over the major statistical properties of the original variables, such as the correlation structure among covariates. While Watson and Wright (2021) claim that the CPI should in principle work with any valid set of knockoffs, it has thus far only been applied and evaluated with Gaussian knockoffs (Candès et al. 2018). This currently limits practitioners to using the CPI method only with continuous variables or to disregard the specialities of mixed data. We analyse consequences of such a disregard when using CPI with Gaussian knockoffs (Candès et al. 2018) (CPIgauss) and deep knockoffs (Romano et al. 2020) (CPIdeep) and propose a specialized solution strategy to tackle the mixed data case: using sequential knockoffs (Kormaksson et al. 2021)—a knockoff sampling algorithm explicitly developed for mixed data—within the CPI framework (CPIseq).

The paper will be structured as follows. We present relevant methodology and FI measures in Sect. 2. Section 2.2 reviews several knockoff sampling algorithms, demonstrating the need for specialized procedures with mixed data and motivating our proposed solution CPIseq. Through simulation studies in Sects. 3.1 and 3.2, we will evaluate our newly proposed workflow in more depth and further compare it to other methods. Finally, we illustrate method application to a real-world dataset in Sect. 3.3 before concluding and discussing our findings in Sect. 4.

## 2 Methods

With a focus on the measurement of model-agnostic, global, conditional FI, this section presents related measures proposed by previous literature and discusses their applicability to mixed data. We acknowledge that methods from the statistical literature on conditional independence testing (Shah and Peters 2020; Williamson et al. 2021) might also be utilized for conditional FI measurement; however, a full comparison of such methods is beyond the scope of this paper. Further, it is worth clarifying at this point that we understand FI here as a concept that is tied to the variable's effect on the predictive performance in a supervised learning task.

### 2.1 Feature importance measures

#### 2.1.1 Conditional subgroup approach (CS)

A global, model-agnostic FI measure that acknowledges the crucial distinction between conditional and marginal measures of importance is the *conditional subgroup* (CS) approach proposed by Molnar et al. (2023). CS partitions the data into interpretable subgroups, i.e., groups whose feature distributions are homogeneous within but heterogeneous between groups. The method is promising, as it explicitly specifies the conditioning between subgroups and further allows for an unconditional interpretation within subgroups. This means the method provides both a global conditional and a within-group unconditional interpretation, which sheds light on feature dependence structures.

To determine FI, CS evaluates the change in loss when the variable of interest is permuted within subgroups, which lowers extrapolation to low-density regions of the feature space, thereby mitigating a common problem with permutation-based approaches (Hooker et al. 2021). To decide on a suitable partition, the authors suggest determining subgroups via transformation trees. Using a pre-specified loss function, the average increase in loss is reported for multiple permutations versus the original ordering of variables.

CS is not affected by mixed data other than through the choice of an appropriate prediction algorithm, which is why this method is suspected to work equally well with mixed data. However, for this approach to work, researchers must assume that the data are separable into subgroups. Further, for testing FI, the method would need to rely on computationally expensive permutation tests as no inherent testing procedure is provided.

#### 2.1.2 Leave-one-covariate-out (LOCO)

*Leave-one-covariate-out* (LOCO) is a fairly simple approach to measuring FI, which, as the name suggests, evaluates the change in predictive performance of a model when leaving out a covariate of interest (Lei et al. 2018). This means, FI is

determined by comparing the loss of the model fitted including or excluding the covariate of interest.

While this is a very intuitive approach, it does involve several drawbacks. First, the model has to be retrained with a different set of variables, which not only incurs high computational cost, but also yields an entirely different model, raising concerns about comparability in general. Further, if correlations or other complex dependencies are present within the data, LOCO might give misleading results if only one covariate at a time is excluded, as this neglects potential interaction effects between groups of variables. In the presence of such group-wise structures, the exclusion of multiple covariates at a time is advisable (Au et al. 2022; Rinaldo et al. 2016).

For the speciality of mixed data, we can again see that all reliance is on the level of model choice, hence, as long as the prediction model is able to process mixed data, LOCO is not affected by different data types.

### 2.1.3 Shapley additive global importance (SAGE)

*Shapley additive global importance (SAGE)* (Covert et al. 2020) is a model-agnostic FI measure that aims to take into account feature interactions on a global level. The method is based on Shapley values (Shapley 1953), which have received much attention in interpretable machine learning recently. While Shapley values are widespread in their use for giving local explanations, i.e. explaining the role of features in individual predictions made by the model, Covert et al. (2020) propose a global extension such that the role of features can be understood on a model-wide level. SAGE values are Shapley values for the features with regard to the predictive power of the model. Therefore, SAGE values can also be calculated by directly calculating Shapley values for the model loss, e.g. as proposed in LossSHAP (Lundberg et al. 2020), and then average across all instances to achieve a global measure. However, Covert et al. (2020) propose a fast approximation algorithm.

The SAGE methodology allows for taking feature interaction effects into account, however, in practice, implementations typically use marginal sampling as an approximation to the conditional densities when sampling to replace the respective feature in various coalitions. This results in explanations that are comparable to marginal measures of FI when applied to real-world data.

Mixed data affect SAGE at the variable sampling step to build the coalitions and through the choice of the predictive model. With the use of marginal imputation and a model that is able to process mixed data, SAGE should not be affected by mixed data types.

### 2.1.4 Conditional predictive impact (CPI)

A fairly general approach to tackle conditional FI measurement is the *conditional predictive impact (CPI)* proposed by Watson and Wright (2021). To capture conditional FI, a flexible conditional independence test is introduced that works with any supervised learning algorithm, valid knockoff sampler and well-defined loss function. CPI ties FI to predictive performance, arguing that the inclusion of a relevant

variable in the model should improve its predictive performance. Building on this idea, first, a supervised learning algorithm is trained to predict the outcome from given input variables. Then, using a knockoff sampling algorithm, so-called knockoff copies of the input features are generated. These knockoffs retain the covariance structure of the input features,<sup>1</sup> but are (conditional on the input features) independent of the response variable. They therefore serve as a set of negative controls against which to compare the original data. In detail, to compute the CPI statistic, the trained model from the first step is used to predict the target twice: first using the original test data, and again after replacing one or several features of interest in the test data by their knockoff copies. The change in loss is then averaged across samples. Finally, the authors propose to apply inference procedures, such as a paired  $t$  test, to get valid  $p$ -values and confidence intervals for the FI scores.

Given that the prediction algorithm works with mixed data, sampling valid knockoffs for mixed data is the sticking point. As Watson and Wright (2021) claim, the CPI setup is knockoff-agnostic and hence works for any knockoff sampler. However, their simulations are limited to settings of continuous data and Gaussian knockoff sampling, i.e., using `CPIgauss`, only. Resulting from this, practitioners facing mixed data cannot use `CPIgauss` directly and are forced to use workarounds that may perform poorly in practice, e.g. dummy encoding variables and treating them as continuous, of which the effects on the method are thus far unknown. The present work sheds light on the consequences of such procedures, see further Sect. 3.1. To propose an efficient way of making CPI applicable to mixed data, we will now delve into the methodology of knockoffs in greater depth.

## 2.2 Model-X knockoffs

The model-X knockoff framework (Candès et al. 2018) was proposed for variable selection while controlling the false discovery rate (FDR). The idea is to use knockoffs as negative controls in the model, which prevents spuriously correlated variables from being detected as important. These knockoffs are a set of variables  $\tilde{\mathbf{X}}$  that mimic the correlational structure between the original input variables  $\mathbf{X}$ , but crucially are known to be irrelevant to the target variable  $Y$ , conditional on the input data. Intuitively, if  $X_j$  does not significantly outperform  $\tilde{X}_j$  by some importance measure, then  $X_j$  can be removed from the model (Candès et al. 2018).

More formally, to construct a valid knockoff matrix  $\tilde{\mathbf{X}}$  for the  $p$ -dimensional feature matrix  $\mathbf{X}$ , two conditions have to be met. The first is pairwise exchangeability, i.e. for any proper subset  $S \subset \{1, \dots, p\}$ :

$$(\mathbf{X}, \tilde{\mathbf{X}})_{\text{swap}(S)} \stackrel{d}{=} (\mathbf{X}, \tilde{\mathbf{X}}), \quad (3)$$

where  $\stackrel{d}{=}$  represents equality in distribution and  $\text{swap}(S)$  indicates swapping the respective variables in  $S$  with their knockoff counterparts. The second condition is conditional independence, i.e.

<sup>1</sup> This holds true for knockoffs that are at least of second-order, i.e. exhibit the same first two moments as the original data.



$$\tilde{\mathbf{X}} \perp\!\!\!\perp \mathbf{Y} \mid \mathbf{X}. \quad (4)$$

Knockoff methodology is an active field of research. Numerous approaches to knockoff sampling have been proposed, for example, methods based on distributional assumptions (Bates et al. 2021; Candès et al. 2018; Sesia et al. 2018), Bayesian frameworks (Gu and Yin 2021) or deep learning (Jordon et al. 2019; Liu and Zheng 2018; Romano et al. 2020; Sudarshan et al. 2020). While a comprehensive review of knockoff samplers is beyond the scope of this paper, we will present a selection of knockoff samplers that is particularly interesting for applications on mixed data. Namely, we will investigate Gaussian knockoffs (Candès et al. 2018) because of their widespread use, deep knockoffs (Romano et al. 2020) as a representative of deep learning based knockoff generation, and sequential knockoffs as a specialized approach to tackle mixed data.

### 2.2.1 Gaussian knockoffs

As the name suggests, the *Gaussian knockoff* sampler (Candès et al. 2018) is based on the assumption that the input data matrix  $\mathbf{X} \in \mathbb{R}^{N \times p}$  is multivariate Gaussian, i.e.  $\mathbf{X} \sim N(\mu, \Sigma)$ . For simplicity, we assume  $\mu = 0$  and get for the joint distribution which satisfies Eq. (3)

$$(\mathbf{X}, \tilde{\mathbf{X}}) \sim N(0, \mathbf{G}), \quad \text{where } \mathbf{G} = \begin{bmatrix} \Sigma & \Sigma - \text{diag}\{s\} \\ \Sigma - \text{diag}\{s\} & \Sigma \end{bmatrix}$$

with diagonal matrix  $\text{diag}\{s\}$  to ensure positive semi-definiteness of the joint covariance matrix  $\mathbf{G}$ . Knockoffs can then be sampled from the conditional distribution  $\tilde{\mathbf{X}} \mid \mathbf{X} \stackrel{d}{=} N(\mu, \mathbf{V})$ , where  $\mu, \mathbf{V}$  can be calculated from regular regression formulas. For details see Candès et al. (2018).

Clearly, it is reasonable to suspect this knockoff sampler to work well with Gaussian data. However, with mixed data types, discrete values can only be handled after encoding, e.g. introducing dummy variables, which are evidently non-Gaussian. The consequences of such transformations, i.e. neglecting the special nature of mixed data, have not yet been evaluated for the Gaussian knockoff sampler. In an attempt to quantify such implications to some extent, we will include this knockoff sampler in our analysis in Sect. 3.1 and compare it to more well-suited alternatives.

### 2.2.2 Deep knockoffs

*Deep knockoffs* as proposed by Romano et al. (2020) rely on a random generator, consisting of a deep neural network, to sample valid knockoffs. For variables  $\mathbf{X}$  sampled independently from an unknown distribution  $P_{\mathbf{X}}$ , the random generator is trained such that the joint distribution of  $(\mathbf{X}, \tilde{\mathbf{X}})$  is invariant under swapping, such that Eq. (3) is satisfied. In detail, the neural network takes variables  $\mathbf{X}$  and i.i.d. sampled noise  $\mathcal{E}$  as input to optimize a scoring function that quantifies the extent to which  $\tilde{\mathbf{X}}$  is a good knockoff copy for  $\mathbf{X}$  by evaluating how well Eq. (3) is approximated. Considering

the neural network architecture, the authors suggest using a width  $h$  that is ten times the dimensionality of the input feature space, i.e.  $h = 10p$  and six hidden layers which they claim should work well for a “wide range of scenarios”, but acknowledge that “more effective designs” might be found (Romano et al. 2020).

Making use of recent deep learning advances, deep knockoffs should—according to the authors—generalize well to the mixed data case. Romano et al. (2020) claim that this framework samples approximate knockoffs for arbitrary distributions. However, it is worth noting that there is little explicit methodology available to the user beyond making general claims about the generalizability of the method. Therefore, an applied user is again left with a knockoff sampler that does not return valid mixed data knockoffs.<sup>2</sup>

### 2.2.3 Sequential knockoffs

*Sequential knockoff* (Kormaksson et al. 2021) sampling is based on the conditional independent pairs algorithm (Candès et al. 2018) given in Supplementary Information A with a specialized strategy to model the conditional distribution  $P(X_j | X_{-j}, \tilde{X}_{1:j-1})$  and sample knockoffs for mixed data.

Sequential knockoffs are synthesized by sampling continuous knockoffs from a Gaussian distribution and categorical knockoffs from a multinomial distribution with distribution parameters that have been sequentially estimated through penalized<sup>3</sup> linear or multinomial logistic regression models. The procedure is given in more detail in Algorithm 1, where  $X_{-j} := (X_1, \dots, X_{j-1}, X_{j+1}, \dots, X_p)$  and  $\tilde{X}_{1:j-1} := (\tilde{X}_1, \dots, \tilde{X}_{j-1})$ .

---

**Algorithm 1** Sequential knockoffs through conditional independent pairs

---

```

j = 1
while j ≤ p do
  if Xj continuous then
    sample  $\tilde{X}_j$  from  $N(\hat{\mu}, \hat{\sigma})$  with  $\hat{\mu}, \hat{\sigma}$  estimated from
    penalized linear regression  $\tilde{X}_j \sim X_{-j}, \tilde{X}_{1:j-1}$ 
  else if Xj categorical then
    sample  $\tilde{X}_j$  from  $Multinom(\hat{\pi})$  with  $\hat{\pi}$  estimated from
    penalized multinomial logistic regression  $\tilde{X}_j \sim X_{-j}, \tilde{X}_{1:j-1}$ 
  end if
  j = j + 1
end while

```

---

Algorithm 1 yields valid knockoff copies for data that may consist of both categorical and continuous covariates. Hence, the present paper puts a special focus on


<sup>2</sup> Deep generative models for mixed data is an active and promising area of research (Xu et al. 2019; Watson et al. 2022), although we are unaware of any implementation for knockoff sampling.

<sup>3</sup> In our experiments, we follow the advice of Kormaksson et al. (2021) to use an elastic net (Zou and Hastie 2005). Note that the ordering of variables might be of relevance in finite samples and that the procedure requires the various levels of the categorical variable to occur sufficiently often.

this method and evaluates its suitability for conditional FI measurement with mixed data.

### 2.3 CPI with sequential knockoffs: CPIseq

We propose to combine two frameworks that have, thus far, not been analysed in conjunction, the CPI (Watson and Wright 2021) and sequential knockoffs (Kor-maksson et al. 2021), as a viable solution for conditional FI measurement with mixed data. Section 2 reveals that amongst the limited number of conditional FI measurement methods available, CPI is one of the few conditional FI methods that allows for the direct application of statistical testing procedures. Further, we have seen that the major obstacle of CPI with mixed data is the knockoff generation step. When surveying the literature on knockoffs in Sect. 2.2, the sequential knockoff sampler stands out as a solution that tackles the special nature of mixed data. Algorithm 2 presents details on the procedure we propose here. Note that for calculating CPIseq for several features (or groups)  $j$ , steps 1 and 2 of the algorithm do not have to be recalculated for each  $j$ .



**Fig. 2** Rejection rates of one-sided paired  $t$  tests at  $\alpha = 0.05$  to detect relevant variables, i.e. power and type I error rates, for CPI with various knockoff samplers across 500 simulation runs.  $X_1, X_3$  are 10-level categoricals,  $X_2, X_4$  are Gaussian. Effect size  $\beta = 0.5$  and random forest prediction model

**Algorithm 2** CPI with sequential knockoffs (CPIseq)

**Input:**  $(X^{train}, Y^{train}), (X^{test}, Y^{test})$ , supervised learner  $f$ , feature (or group) of interest  $j$ , sequential knockoff sampler  $s$  (Alg. 1), loss function  $L$ , inference procedure  $h$

- 1: learn  $\hat{f} \leftarrow f(X^{train}, Y^{train})$
- 2: sample knockoffs  $\tilde{X}^{test} \leftarrow s(X^{test})$
- 3: for feature (or group)  $j$  calculate instance-wise loss difference
 
$$\Delta^{(i)} \leftarrow L(\hat{f}, X^{test(i)}) - L(\hat{f}, \{X_{-j}^{test(i)}, \tilde{X}_j^{test(i)}\})$$
- 4: calculate conditional predictive impact
 
$$\widehat{\text{CPI}} \leftarrow \frac{1}{N} \sum_{i=1}^N \Delta^{(i)}$$
- 5: apply inference procedure for  $p$ -value  $p$  and confidence interval  $ci$ 

$$p, ci \leftarrow h(\Delta)$$

**Output:**  $\widehat{\text{CPI}}, p, ci$


The CPIseq we propose here combines the features of the CPI methodology with ease of applicability to real data, which often consists of mixed data types. Providing frequentist inference procedures without model refitting is the major advantage over other conditional FI methods, such as CS and LOCO. To ensure high power for these testing procedures, adequate handling of mixed data is a prerequisite and CPIseq assures this through the flexible sequential knockoff subroutine.

### 3 Experiments

In this section, we analyse the performance of various FI measures on both simulated and empirical data. Through simulation studies, we evaluate the performance of our newly proposed workflow in comparison to other approaches. First, we investigate how CPIseq compares to CPI with other knockoff samplers, namely CPIgauss and CPIdeep (Sect. 3.1) in terms of power and effective FDR control. Further, we compare feature rankings given by our proposed approach and other conditional FI-related measures that do not use knockoffs (Sect. 3.2). Finally, we use a real-world data example to illustrate method application (Sect. 3.3).

#### 3.1 Comparing knockoffs

Major differences in the performance of CPIgauss, CPIdeep, and CPIseq on mixed data are illustrated using the following simulation setup. Consider a linear system of input variables  $S = \{X_1, X_2, X_3, X_4\}$  and target variable  $Y$ , visualized by the directed acyclic graph (DAG)  $\mathcal{G}$  in Fig. 2. Since the joint distribution is Markov with respect to  $\mathcal{G}$ , it follows by  $d$ -separation (Pearl 2009) that  $X_1 \perp\!\!\!\perp Y \mid S \setminus \{X_1\}$  and  $X_2 \perp\!\!\!\perp Y \mid S \setminus \{X_2\}$ , whereas  $X_3 \not\perp\!\!\!\perp Y \mid S \setminus \{X_3\}$  and  $X_4 \not\perp\!\!\!\perp Y \mid S \setminus \{X_4\}$ . Therefore, a



**Fig. 3** Mean AUC value with  $\pm$  one standard deviation across 500 simulation runs. Categorical variables with  $c = 5$  levels, pairwise correlation  $\rho = 0.8$  and a random forest prediction model for continuous target  $Y$

conditional FI measure should only attribute nonzero importance to variables  $X_3, X_4$ , but not to  $X_1, X_2$ . We consider three scenarios to track consequences of mixed data closely. For the baseline scenario (I),  $S$  will be Gaussian; for scenario (II),  $X_1$  or  $X_3$  will be binary; and in scenario (III),  $X_1$  and/or  $X_3$  will be categorical with  $c \in \{4, 10\}$  levels. Scenarios (II) and (III) further include an all categorical setting, i.e.  $S$  will be categorical, as a point of reference. We carefully select relevant combinations of category levels (2, 4 or 10), type of the target variable (continuous or binary) and fitted model (generalized linear model or random forest). See Supplementary Information B.1 to B.4 for further details on the experimental setup, including details on the prediction models and their validation.

### 3.1.1 Results

For scenario (I), we find CPI achieving high power and effective type I error control with every knockoff sampling algorithm. Naturally, as the data is Gaussian, we see CPIgauss achieving high power in this setting, see Supplementary Information Fig. 3. When transforming  $X_1$  and  $X_3$  into binary variables, (scenario (II)), we still observe high power and type I error control.

For input data consisting of mixed data types where the categorical variables are of high-cardinality (scenario (III)), we can see from Fig. 2 that the sequential knockoff sampler provides greater sensitivity than the deep or Gaussian alternatives across all tested sample sizes. Rejection rates for CPIseq grow quickly with sample size, reaching about 90% power around  $N = 2000$ . By contrast, CPIgauss only reaches about 50% and the deep knockoff sampler about 70% power at the maximal  $N = 7000$ . In terms of type I error control, all methods seem to be robust against the categorical nature of the irrelevant variable  $X_1$ , as the rejection rate in Fig. 2 is kept close to  $\alpha = 0.05$  for all knockoff samplers.

A full presentation of results is given in Supplementary Information B.5, including Figures for the all categorical cases, for which we find similar results as in mixed data settings.


This simulation study demonstrates that the power of CPIgauss and CPIdeep might be severely affected by high-cardinality features. We find CPIseq to provide a powerful solution to conditional FI measurement, i.e. to detect conditionally important categorical features, whereas CPIgauss and CPIdeep are less sensitive with such data. It is worth noting that CPIgauss and CPIdeep perform surprisingly well when mixed data is limited to continuous and binary data types, even though Gaussian and deep knockoffs inevitably generate data outside the support of Boolean variables. Nevertheless, CPIseq appears to be the most powerful solution for conditional FI measurement with high-cardinality categorical data.

### 3.2 Comparing feature importance measures

Through a simulation study, our newly proposed workflow CPIseq will now be set in comparison with LOCO (Lei et al. 2018), CS (Molnar et al. 2023), SAGE (Covert et al. 2020), and permutation feature (PFI) importance (Breiman 2001; Fisher et al. 2019). Even though CPIgauss and CPIdeep have been shown to be outperformed by CPIseq in Sect. 3.1, we add these two methods to the simulation in order to provide a complete picture on how they relate to other measures of FI. Further enriching the picture of FI measure comparison, we discuss a random forest model-specific FI procedure (Kursa and Rudnicki 2010) and its performance in comparison to the other FI measures in the Supplementary Information C.6.

We simulate multivariate normal data with a pre-specified correlation structure to ensure a simple setup while incorporating a larger number of variables than in our toy example in Sect. 3.1. Again, we transform several variables into categorical, such that we end up with mixed data. We distinguish between variables having zero, weak, or strong effect on the outcome  $Y$ , and for the continuous variables we further separate variables with a linear or nonlinear effect on  $Y$ . Further, we ensure that there is an equal number of relevant and irrelevant variables, such that each relevant variable is correlated with exactly one irrelevant variable of the same type, yielding a total of  $p = 12$  variables. In sum, we analyse a total of 24 settings by varying the correlation strength ( $\rho = 0.5$  or  $0.8$ ), type of target variable  $Y$  (continuous or binary), varying number of category levels ( $c = 2$  or  $5$ ) and fitting various machine learning prediction models (generalized linear model, random forest or neural network), see Supplementary Information C.1 and C.2 for further details.

Some of the methods included in the comparison do not provide statistical testing procedures. Therefore, we will compare methods by their tendency to rank relevant features higher than irrelevant alternatives. By construction,  $p = 6$  variables are relevant to the outcome, whereas the other  $p = 6$  variables are not. Hence, when we ask the methods to rank the variables according to their importance, ideally, the 6 relevant variables are ranked amongst the top 6. We will use the area under the receiver operating characteristic curve (AUC) as a measure of



**Fig. 4** Proportion of features ranked amongst the top 6 of 12 by variable type across 500 simulation runs. Solid lines (relevant variables) correspond to sensitivity, dashed lines (irrelevant variables) correspond to 1-specificity. Categorical variables with  $c = 5$  levels, pairwise correlation  $\rho = 0.8$  and a random forest prediction model for continuous target  $Y$

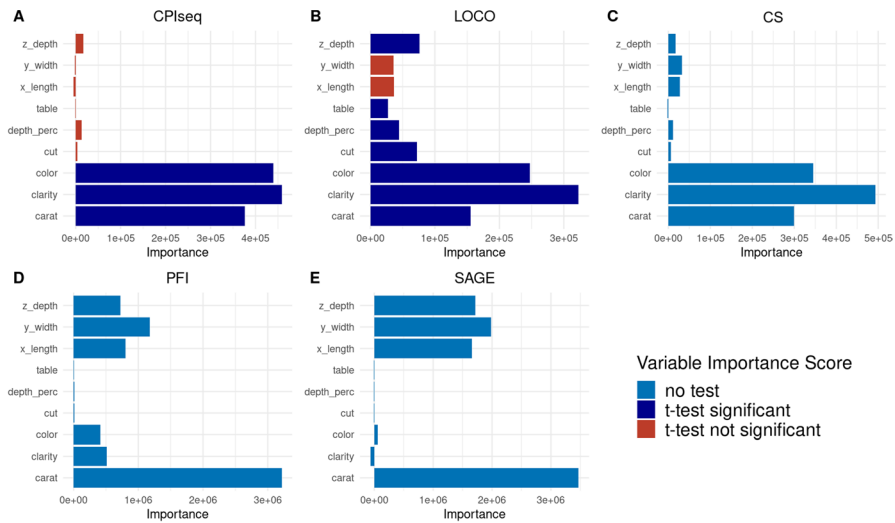
performance and will further report sensitivity and 1-specificity for each of the methods. See further Supplementary Information C.3.

### 3.2.1 Results

We find CS, CPIseq, CPIgauss, CPIdeep and LOCO outperforming PFI and SAGE in ranking the relevant variables amongst the top 6 variables in terms of AUC scores (Fig. 3). AUC scores rise with increasing sample size, however, while the conditional measures form a group that gets close to the optimal score of 1, the performance of marginal measures<sup>4</sup> flattens out. This behaviour stems from the phenomenon of marginal methods to attribute nonzero importance to correlated, but irrelevant variables, affecting the methods ability to separate the top 6 from the bottom 6 variables, as can be further investigated from Fig. 4.

Figure 4 depicts the proportion of the respective variable types being ranked amongst the top 6 variables. Ideally, this proportion should be high for relevant variables (solid lines) and low for irrelevant variables (dashed lines). Panel (B) shows that both PFI and SAGE mistakenly rank the irrelevant continuous variables with a linear effect, which are correlated to the relevant continuous linear variables, amongst the top 6 variables. This is unsurprising, because relevant continuous variables with a linear effect on the target are the easiest to detect, and hence, irrelevant variables correlated to these variables are most likely to be mixed up by marginal measures in the full ranking. Note that because each of the methods has to assign ranks 1–12, an irrelevant variable being mistakenly ranked amongst the

<sup>4</sup> Note that SAGE here is closer to a marginal measure because of the marginal imputation subroutine.



**Fig. 5** Feature importance scores for predicting the selling price of diamonds using a random forest model. For the CPIseq and LOCO,  $t$ -tests are at  $\alpha = 5\%$ , using the Holm procedure to adjust for multiple testing

top 6 variables in return leads to a relevant variable being ranked within the bottom 6 ranks. For example, due to the marginal measurement of FI, the PFI measure is ranking correlated yet irrelevant variables amongst the most important predictors (dashed line in Fig. 4, Panel B), which in turn forces PFI to mistakenly rank some relevant variables low (solid line in Fig. 4, Panel A).

Regarding the comparison of CPI-based methods, we find CPIseq outperforming CPIgauss and CPIdeep in detecting relevant categorical variables in the mixed data setting, see Fig. 4, Panel A, which underpins the findings of simulations in Sect. 3.1.

To check for robustness, we used several predictive models (generalized linear model, random forest, and neural network), varied the type of the target variable (regression or classification task) and the number of categories for the categorical variables (2, 5), and found similar results. Further, we analysed the fit of the prediction models on test data to ensure reliable FI measurement. See Supplementary Information C.4 and C.5 for details on the robustness analyses.

In sum, this simulation demonstrates both that CPIseq is competitive with other conditional FI measures, and illustrates the importance of distinguishing between marginal and conditional measures. It is worth emphasizing again that the CPIseq workflow not only ranks features, but also enables powerful conditional FI testing. We will see the practical relevance of this in the following section.

### 3.3 Real-world data

We conclude the section on experiments with a real-world data application to illustrate our proposed workflow on empirical mixed data. As an example, we use the



diamonds dataset which is publicly available on OpenML<sup>5</sup> (Vanschoren et al. 2014). Consisting of 9 covariates (6 numerical, 3 categorical) which relate to characteristics of diamonds such as length, depth and colour. We predict the selling price of the diamond in USD (*price*) using a random forest prediction model. Similar to the experiments in Sect. 3.2, the importance of the covariates for the prediction model will be determined by CPIseq, CS, LOCO, PFI and SAGE. For further details on the dataset and the procedure, as well as a comparison to results given by another prediction model (neural network), see Supplementary Information D.

Figure 5 illustrates the difference between conditional and marginal measures of feature importance. The marginal measures (Fig. 5, Panels D, E) attribute high importance scores to the covariates *x\_length*, *y\_width*, *z\_depth* and *carat*, whereas the conditional measures (Fig. 5, Panels A, B, C) attribute high importance scores to the covariates *colour*, *clarity* and *carat*. Note that the scale of the FI measures in Fig. 5 differs, since marginal measures also incorporate the importances of correlated variables and hence, by construction, exhibit much larger values than conditional FI measures.

With some background knowledge on the physical characteristics of diamonds, we can understand the causal relationships that lead to this result. Carat is a measure of weight, and with round diamonds, this weight can be approximated by the formula  $carat = length \times width \times depth \times 0.0061$  (Miller 1988). Note that to ensure this formula holds, we only considered diamonds with a deviation  $< 0.02$  mm from a perfect round shape, yielding a subset of  $N = 4463$  observations. The covariates *x\_length*, *y\_width* and *z\_depth* therefore determine the weight (*carat*), which all the importance measures suggest as an important predictor variable for *price*. Conditional FI measures then suggest that *x\_length*, *y\_width* and *z\_depth* do not carry further information on the price, given the other covariates, including *carat*. Marginal measures, however, attribute importance irrespective of other covariates and hence do not condition on the information given by *carat*, which leads to high importance values for *x\_length*, *y\_width*, *z\_depth* as well as *carat*, even though it is reasonable to assume that *carat* absorbs all relevant information given by *x\_length*, *y\_width* and *z\_depth* on the price of diamonds.

The conditional FI measures further detect the variables *colour* and *clarity* to be relevant for the prediction of *price*. Note that we here again have to see this in a conditional sense. Given the other covariates, the variables *colour* and *clarity* do provide additional information on the *price*, whereas marginal measures estimate a rather low importance of these variables.

This real-world example emphasizes the difference between conditional and marginal FI measures and its implications. Again, it is worth repeating that out of the conditional measures, CPIseq facilitates the interpretation through inference procedures providing a clear indication of the relevant variables, whereas this indication is less clear with the LOCO testing procedure and CS not providing the user with testing procedures at all.

<sup>5</sup> [https://www.openml.org/search?type=data &sort=runs &id=42225](https://www.openml.org/search?type=data&sort=runs&id=42225).

## 4 Conclusion and discussion

In this work, we highlight the importance of taking statistical considerations into account when measuring FI in interpretable machine learning. Specifically, we focus on conditional versus marginal perspectives on FI measurement, and analyse conditional FI methods with special regard for mixed data. We introduce the combination of CPI and sequential knockoffs (CPIseq) as a strategy that enables testing of conditional, model-agnostic, global FI with mixed data. Through simulation studies, we show that CPIseq achieves high power, whereas CPIgauss and CPIdeep are less sensitive for categorical features. Further, we benchmark this method against other conditional FI measures, finding competitive performance, and use a real-world data example to illustrate empirical implications. In sum, we demonstrate that the CPIseq provides researchers with a powerful test for conditional FI while working on a global, model-agnostic level.

Our analyses are limited by the availability of specialized knockoff sampling algorithms for the generation of mixed data knockoffs. Astonishingly, the case of mixed data has not received much attention in the knockoff literature so far and even if some methods were claimed to generalize to the mixed data case (Romano et al. 2020), there is a lack of concrete methodology and software implementation. Also, the scarce availability of conditional FI measures that allow for effective statistical testing impedes efficient comparison between FI metrics, forcing the evaluation to rely on rankings. While rankings are oftentimes used in the literature on FI for illustrative purposes, a systematic gold standard for comparing rankings between methods has not emerged. We hypothesize that this might be due to the fact that in the machine learning community, simulation studies—a standard procedure in the statistics community—are relatively rare, and hence evaluations involving, e.g. ground truth variable rankings are not in the focus. In particular, with mixed data, a ground truth ranking of simulated variables is not straightforward since it is unclear how the categorical nature should be respected and challenging disagreements across methods are likely to occur (Krishna et al. 2022). Methodological development that bridges evaluation strategies commonly applied in statistics with the setting faced in interpretable machine learning, e.g. FI rankings, is highly desirable.

This work highlights the necessity for procedures that respect data-specific requirements, such as respecting the categorical nature of variables in mixed datasets. Our simulations show that a neglect of such requirements and the application of workarounds might lead to undesirable consequences. We encourage researchers to develop methods that are specifically designed for realistic (mixed) data, instead of leaving practitioners with broad claims of the generalizability of their method. While some generalizations are indeed effortless, e.g. for conditional independence testing with all categorical data exact  $p$ -values can be computed through permutations (Tsamardinos and Borboudakis 2010), whereas conditional independence testing in general, including mixed data cases, is severely more challenging (Shah and Peters 2020). Moreover, other data type specific adjustments such as the presence of ordinal data might be of interest for future research, for example, random forest regression models yield the

same results with ordinal as with numeric data (Hastie et al. 2009) and hence FI methods that exploit model-specific advantages for ordinal data might be proposed.

Further, the present work raises awareness of the fact that even though the concept of FI might sound intuitive at first, statistical perspectives on the problem reveal that, for example, the question of marginal in contrast to conditional measurement is of fundamental relevance. We hope this paper elucidates the potential of advancing interpretable machine learning methodology through statistical considerations, which might in turn be mutually beneficial for the future development of the field of explainable artificial intelligence and statistics.

**Supplementary Information** The online version contains supplementary material available at <https://doi.org/10.1007/s10182-023-00477-9>.

**Acknowledgements** MNW and KB received funding for this project from the German Research Foundation (DFG), Emmy Noether Grant 437611051.

**Funding** Open Access funding enabled and organized by Projekt DEAL.

**Data and code availability** Data and code to reproduce all the results presented in this paper is available at [https://github.com/bips-hb/CFI\\_mixedData](https://github.com/bips-hb/CFI_mixedData).

## Declarations

**Conflict of interest** MNW is associate editor of AStA—Advances in Statistical Analysis. KB and DSW declare no conflict of interest.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

- Apley, D.W., Zhu, J.: Visualizing the effects of predictor variables in black box supervised learning models. *J. R. Stat. Soc. Ser. B* **82**(4), 1059–1086 (2020). <https://doi.org/10.1111/rssb.12377>
- Au, Q., Herbringer, J., Stachl, C., Bischl, B., Casalicchio, G.: Grouped feature importance and combined features effect plot. *Data Min. Knowl. Disc.* **36**(4), 1401–1450 (2022). <https://doi.org/10.1007/s10618-022-00840-5>
- Bates, S., Candès, E., Janson, L., Wang, W.: Metropolized knockoff sampling. *J. Am. Stat. Assoc.* **116**(535), 1413–1427 (2021). <https://doi.org/10.1080/01621459.2020.1729163>
- Breiman, L.: Random forests. *Mach. Learn.* **45**(1), 5–32 (2001). <https://doi.org/10.1023/A:1010933404324>
- Candès, E., Fan, Y., Janson, L., Lv, J.: Panning for gold: Model-free knockoffs for high-dimensional controlled variable selection. *J. R. Stat. Soc. Ser. B (Stat. Methodol.)* **80**, 551–577 (2018). <https://doi.org/10.1111/rssb.12265>
- Chen, H., Janizek, J.D., Lundberg, S., Lee, S.-I.: True to the model or true to the data? ArXiv preprint (2020). <https://doi.org/10.48550/arXiv.2006.16234>

- Covert, I., Lundberg, S.M., Lee, S.-I.: Understanding global feature contributions with additive importance measures. *Adv. Neural Inf. Process. Syst.* **33**, 17212–17223 (2020)
- Fisher, A., Rudin, C., Dominici, F.: 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), 1–81 (2019)
- Friedman, J.H.: Greedy function approximation: a gradient boosting machine. *Ann. Stat.* 1189–1232 (2001)
- Glymour, C., Zhang, K., Spirtes, P.: Review of causal discovery methods based on graphical models. *Front. Genetics* (2019). <https://doi.org/10.3389/fgene.2019.00524>
- Gu, J., Yin, G.: Bayesian knockoff filter using gibbs sampler. *ArXiv preprint* (2021). <https://doi.org/10.48550/ArXiv.2102.05223>
- Hastie, T., Tibshirani, R., Friedman, J.H., Friedman, J.H.: *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, vol. 2. Springer, New York (2009)
- Hooker, G., Mentch, L., Zhou, S.: Unrestricted permutation forces extrapolation: variable importance requires at least one more model, or there is no free variable importance. *Stat. Comput.* **31**(6), 1–16 (2021). <https://doi.org/10.1007/s11222-021-10057-z>
- Jordon, J., Yoon, J., van der Schaar, M.: Knockoffgan: generating knockoffs for feature selection using generative adversarial networks. In: *International Conference on Learning Representations* (2019)
- Kormaksson, M., Kelly, L.J., Zhu, X., Haemmerle, S., Pricop, L., Ohlssen, D.: Sequential knockoffs for continuous and categorical predictors: with application to a large psoriatic arthritis clinical trial pool. *Stat. Med.* **40**(14), 3313–3328 (2021). <https://doi.org/10.1002/sim.8955>
- Krishna, S., Han, T., Gu, A., Pombra, J., Jabbari, S., Wu, S., Lakkaraju, H.: The disagreement problem in explainable machine learning: a practitioner's perspective. *ArXiv preprint* (2022). <https://doi.org/10.48550/ArXiv.2202.01602>
- Kursa, M.B., Rudnicki, W.R.: Feature selection with the Boruta package. *J. Stat. Softw.* **36**(11), 1–13 (2010). <https://doi.org/10.18637/jss.v036.i11>
- Lei, J., G'Sell, M., Rinaldo, A., Tibshirani, R.J., Wasserman, L.: Distribution-free predictive inference for regression. *J. Am. Stat. Assoc.* **113**(523), 1094–1111 (2018). <https://doi.org/10.1080/01621459.2017.1307116>
- Linaratos, P., Papastefanopoulos, V., Kotsiantis, S.: Explainable AI: a review of machine learning interpretability methods. *Entropy* (2021). <https://doi.org/10.3390/e23010018>
- Liu, Y., Zheng, C.: Auto-encoding knockoff generator for FDR controlled variable selection. *ArXiv preprint* (2018). <https://doi.org/10.48550/ARXIV.1809.10765>
- Lundberg, S.M., Lee, S.-I.: A unified approach to interpreting model predictions. *Adv. Neural Inf. Process. Syst.* **30** (2017)
- Lundberg, S.M., Erion, G., Chen, H., DeGrave, A., Prutkin, J.M., Nair, B., Katz, R., Himmelfarb, J., Bansal, N., Lee, S.-I.: From local explanations to global understanding with explainable AI for trees. *Nat. Mach. Intell.* **2**(1), 56–67 (2020). <https://doi.org/10.1038/s42256-019-0138-9>
- Miller, A.M.: *Jewelry from antiquity to the modern era*. In: *Gems and Jewelry Appraising*. Springer, Boston (1988). [https://doi.org/10.1007/978-1-4684-1404-2\\_5](https://doi.org/10.1007/978-1-4684-1404-2_5)
- Molnar, C., König, G., Bischl, B., Casalicchio, G.: Model-agnostic feature importance and effects with dependent features: a conditional subgroup approach. In: *Data Mining and Knowledge Discovery*, pp. 1–39 (2023). <https://doi.org/10.1007/s10618-022-00901-9>
- Pearl, J.: *Causality*. Cambridge University Press, Cambridge (2009). <https://doi.org/10.1017/CBO9780511803161>
- Ribeiro, M.T., Singh, S., Guestrin, C.: “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* (2016). <https://doi.org/10.1145/2939672.2939778>
- Ribeiro, M.T., Singh, S., Guestrin, C.: Anchors: high-precision model-agnostic explanations. In: *Proceedings of the AAAI Conference on Artificial Intelligence* (2018)
- Rinaldo, A., Wasserman, L., G'Sell, M., Lei, J.: Bootstrapping and sample splitting for high-dimensional, assumption-free inference (2016). <https://doi.org/10.48550/ARXIV.1611.05401>
- Romano, Y., Sesia, M., Candès, E.: Deep knockoffs. *J. Am. Stat. Assoc.* **115**(532), 1861–1872 (2020). <https://doi.org/10.1080/01621459.2019.1660174>
- Sesia, M., Sabatti, C., Candès, E.J.: Gene hunting with hidden Markov model knockoffs. *Biometrika* **106**(1), 1–18 (2018). <https://doi.org/10.1093/biomet/asy033>
- Shah, R.D., Peters, J.: The hardness of conditional independence testing and the generalised covariance measure. *Ann. Stat.* **48**(3), 1514–1538 (2020). <https://doi.org/10.1214/19-AOS1857>

- Shapley, L.: A value for  $n$ -Person games. In: Kuhn, H., Tucker, A. (eds.) *Contributions to the Theory of Games II*. Princeton University Press, Princeton (1953). <https://doi.org/10.1515/9781400881970-018>
- Shrikumar, A., Greenside, P., Kundaje, A.: Learning important features through propagating activation differences. In: *International Conference on Machine Learning*. PMLR (2017)
- Sudarshan, M., Tansey, W., Ranganath, R.: Deep direct likelihood knockoffs. *Adv. Neural Inf. Process. Syst.* **33** (2020)
- Tsamardinos, I., Borboudakis, G.: Permutation testing improves Bayesian network learning. In: *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD* (2010). [https://doi.org/10.1007/978-3-642-15939-8\\_21](https://doi.org/10.1007/978-3-642-15939-8_21)
- Vanschoren, J., Van Rijn, J.N., Bischl, B., Torgo, L.: Openml: networked science in machine learning. *ACM SIGKDD Explorations Newsl.* **15**(2), 49–60 (2014)
- Watson, D.S., Wright, M.N.: Testing conditional independence in supervised learning algorithms. *Mach. Learn.* **110**(8), 2107–2129 (2021). <https://doi.org/10.1007/s10994-021-06030-6>
- Watson, D.S., Blesch, K., Kapar, J., Wright, M. N.: Adversarial random forests for density estimation and generative modeling. In: *Proceedings of the 26th international conference on artificial intelligence and statistics*, PMLR **206** (2023)
- Williamson, B.D., Gilbert, P.B., Carone, M., Simon, N.: Nonparametric variable importance assessment using machine learning techniques. *Biometrics* **77**(1), 9–22 (2021). <https://doi.org/10.1111/biom.13392>
- Xu, L., Skoularidou, M., Cuesta-Infante, A., Veeramachaneni, K.: Modeling tabular data using conditional GAN. *Adv. Neural Inf. Process. Syst.* **32** (2019)
- Zou, H., Hastie, T.: Regularization and variable selection via the elastic net. *J. R. Stat. Soc. Ser. B (Stat. Methodol.)* **67**(2), 301–320 (2005). <https://doi.org/10.1111/j.1467-9868.2005.00503.x>

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

## Authors and Affiliations

Kristin Blesch<sup>1,2</sup>  · David S. Watson<sup>3</sup> · Marvin N. Wright<sup>1,2,4</sup>

David S. Watson  
david.watson@kcl.ac.uk

Marvin N. Wright  
wright@leibniz-bips.de

<sup>1</sup> Leibniz Institute for Prevention Research & Epidemiology - BIPS, Bremen, Germany

<sup>2</sup> Faculty of Mathematics and Computer Science, University of Bremen, Bremen, Germany

<sup>3</sup> Department of Informatics, King's College London, London, UK

<sup>4</sup> Department of Public Health, University of Copenhagen, Copenhagen, Denmark