



## OPEN ACCESS

## EDITED BY

Andreas Karavos,  
Ionian University, Greece

## REVIEWED BY

Pruthwik Mishra,  
Sardar Vallabhbhai National Institute of  
Technology Surat, India  
Gunjan Kumar,  
BBS Group of Institutions, India  
Varun Dogra,  
Lovely Professional University, India  
Poornima Shetty,  
Manipal Academy of Higher Education, India  
Alexandre De Spindler,  
Zurich University of Applied Sciences,  
Switzerland

## \*CORRESPONDENCE

Stefan Haufe  
✉ haufe@tu-berlin.de

RECEIVED 28 August 2025  
ACCEPTED 31 October 2025  
PUBLISHED 05 January 2026

## CITATION

Wilming R, Dox A, Schulz H, Oliveira M, Clark B and Haufe S (2026) GECOBench: a gender-controlled text dataset and benchmark for quantifying biases in explanations. *Front. Artif. Intell.* 8:1694388. doi: 10.3389/frai.2025.1694388

## COPYRIGHT

© 2026 Wilming, Dox, Schulz, Oliveira, Clark and Haufe. This is an open-access article distributed under the terms of the [Creative Commons Attribution License \(CC BY\)](#). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

# GECOBench: a gender-controlled text dataset and benchmark for quantifying biases in explanations

Rick Wilming<sup>1</sup>, Artur Dox<sup>2</sup>, Hjalmar Schulz<sup>2</sup>, Marta Oliveira<sup>1</sup>,  
Benedict Clark<sup>1</sup> and Stefan Haufe<sup>1,2\*</sup>

<sup>1</sup>Physikalisch-Technische Bundesanstalt, Berlin, Germany, <sup>2</sup>Technische Universität Berlin, Berlin, Germany

Large pre-trained language models have become a crucial backbone for many downstream tasks in natural language processing (NLP), and while they are trained on a plethora of data containing a variety of biases, such as gender biases, it has been shown that they can also inherit such biases in their weights, potentially affecting their prediction behavior. However, it is unclear to what extent these biases also affect feature attributions generated by applying “explainable artificial intelligence” (XAI) techniques, possibly in unfavorable ways. To systematically study this question, we create a gender-controlled text dataset, GECO, in which the alteration of grammatical gender forms induces class-specific words and provides ground truth feature attributions for gender classification tasks. This enables an objective evaluation of the correctness of XAI methods. We apply this dataset to the pre-trained BERT model, which we fine-tune to different degrees, to quantitatively measure how pre-training induces undesirable bias in feature attributions and to what extent fine-tuning can mitigate such explanation bias. To this extent, we provide GECOBench, a rigorous quantitative evaluation framework for benchmarking popular XAI methods. We show a clear dependency between explanation performance and the number of fine-tuned layers, where XAI methods are observed to benefit particularly from fine-tuning or complete retraining of embedding layers.

## KEYWORDS

XAI, NLP, benchmark, dataset, explainability, interpretability, language models

## 1 Introduction

Large neural network architectures are often complex, making it difficult to understand the mechanisms by which model outputs are generated. This has led to the development of dedicated *post-hoc* analysis tools, commonly referred to as “explainable artificial intelligence” (XAI). In many cases, XAI methods provide so-called feature attributions, which assign an “importance” score to each feature of a given input (e.g. [Ribeiro et al., 2016](#); [Lundberg and Lee, 2017](#); [Sundararajan et al., 2017](#)). In the Natural Language Processing (NLP) domain, feature attribution methods in supervised learning settings are expected to highlight parts of an input text (e.g., words or sentences) that are related to the predicted target, such as a sentiment score.

However, it remains unclear to what extent feature attribution methods help answer specific explanation goals, such as model debugging (Haufe et al., 2024). With it, questions arise about the correctness of feature attributions. One reason why it is challenging to determine the correctness of attribution methods is the tension between model-centric and data-centric explanations (e.g. Murdoch et al., 2019; Chen et al., 2020; Fryer et al., 2020; Haufe et al., 2024). In these scenarios, it is unknown how to define notions of correctness in a principled manner; thus, the extent to which feature attribution methods provide explanations that are purely model-centric or data-centric is unknown. Empirical studies on synthetic datasets have demonstrated that numerous feature attribution methods fail to fulfill basic data-centric requirements, such as highlighting features that have a statistical association with the prediction target (also referred to as the *Statistical Association Property* (SAP)) (e.g. Wilming et al., 2022; Oliveira et al., 2024; Clark et al., 2024). Here, we adopt this data-centric view of assessing the correctness of feature attributions and apply it to the NLP domain.

Furthermore, within the NLP domain attribution methods are typically applied to large pre-trained language models, which are adapted to downstream tasks through transfer learning [e.g., BERT (Devlin et al., 2019) and its variants (Devlin et al., 2019; Liu et al., 2019; Radford et al., 2018; Bugliarello et al., 2021; Dodge et al., 2020; OpenAI, 2023)].

Pre-trained language models are commonly trained on large corpora of text scraped from public and non-public sources, including Wikipedia, Project Gutenberg,<sup>1</sup> or OpenWebText.<sup>2</sup> These large corpora contain a variety of biases, such as biases against demographic groups (Beukeboom, 2014; Graells-Garrido et al., 2015; Reagle and Rhue, 2011). It has been shown that such biases affect model weights (Mitchell, 2007; Montañez et al., 2019) and that text corpora exhibiting problematic biases are amplified in large language models, such as BERT (e.g. Bordia and Bowman, 2019; Gonen and Goldberg, 2019; Blodgett et al., 2020; Nadeem et al., 2021).

However, it remains unclear to what extent biases contained in pre-training corpora are reflected in explanations provided by feature attribution methods, potentially hindering them from meeting specific correctness requirements such as the SAP with respect to the target data distribution and prediction task. Using the example of grammatical gender, we can imagine one particular way in which pre-training biases might lead to incorrect feature attributions or point to residual bias in fine-tuned models. In a gender classification task, asymmetries in the frequencies of specific words may be present in a pre-training corpus but not in the target domain. For example, historical novels may be biased toward male protagonists and depict women less frequently and in more narrowly defined roles, often adhering to historical gender norms. However, the association between, for example, role-specific words and gender in these texts is irrelevant when it comes to distinguishing grammatical gender (as well as for many other tasks). A feature attribution method that highlights respective words, thus suggesting the influence of pre-training biases.

To study and quantify the data-centric correctness of feature attribution methods and the influence of biases, we make two key contributions: (1) *GECO*—a gender-controlled dataset and (2) *GECOBench*—a quantitative benchmarking framework to assess the correctness of feature attributions for language models on gender classification tasks. Both contribute to the future development of novel XAI methods, helping with their evaluation and correctness assessment. An overview is shown in Figure 1.

*GECO*<sup>3</sup> is a gender-controlled dataset in which each sentence  $x$  appears in three grammatically gendered variants: male  $x^M$ , female  $x^F$ , and non-binary  $x^{NB}$ . The three variants are identical apart from gender-specific words such as pronouns. For example, consider the sentence “She loves to spend time with her favorite cat.” We label this sentence as “female (F)” because it entails the pronouns “she” and “her.” By replacing the pronouns with “he” and “his,” we define the “male (M)” counterpart of this sentence. Our approach to creating sentences with minimal changes can be seen as similar to counterfactual data augmentation (Kaushik et al., 2019; Liu et al., 2021).

*GECOBench*<sup>4</sup> is a workflow to quantitatively benchmark the correctness of feature attributions, specifically evaluating XAI methods for NLP classification tasks induced by *GECO* or similar datasets. Here, we showcase the use of *GECOBench*, where BERT (Devlin et al., 2019), a language model pre-trained on Wikipedia data, serves as an exemplary language model. While this benchmark can be extended to include more models, our primary focus is on benchmarking explanation methods rather than the language models themselves.

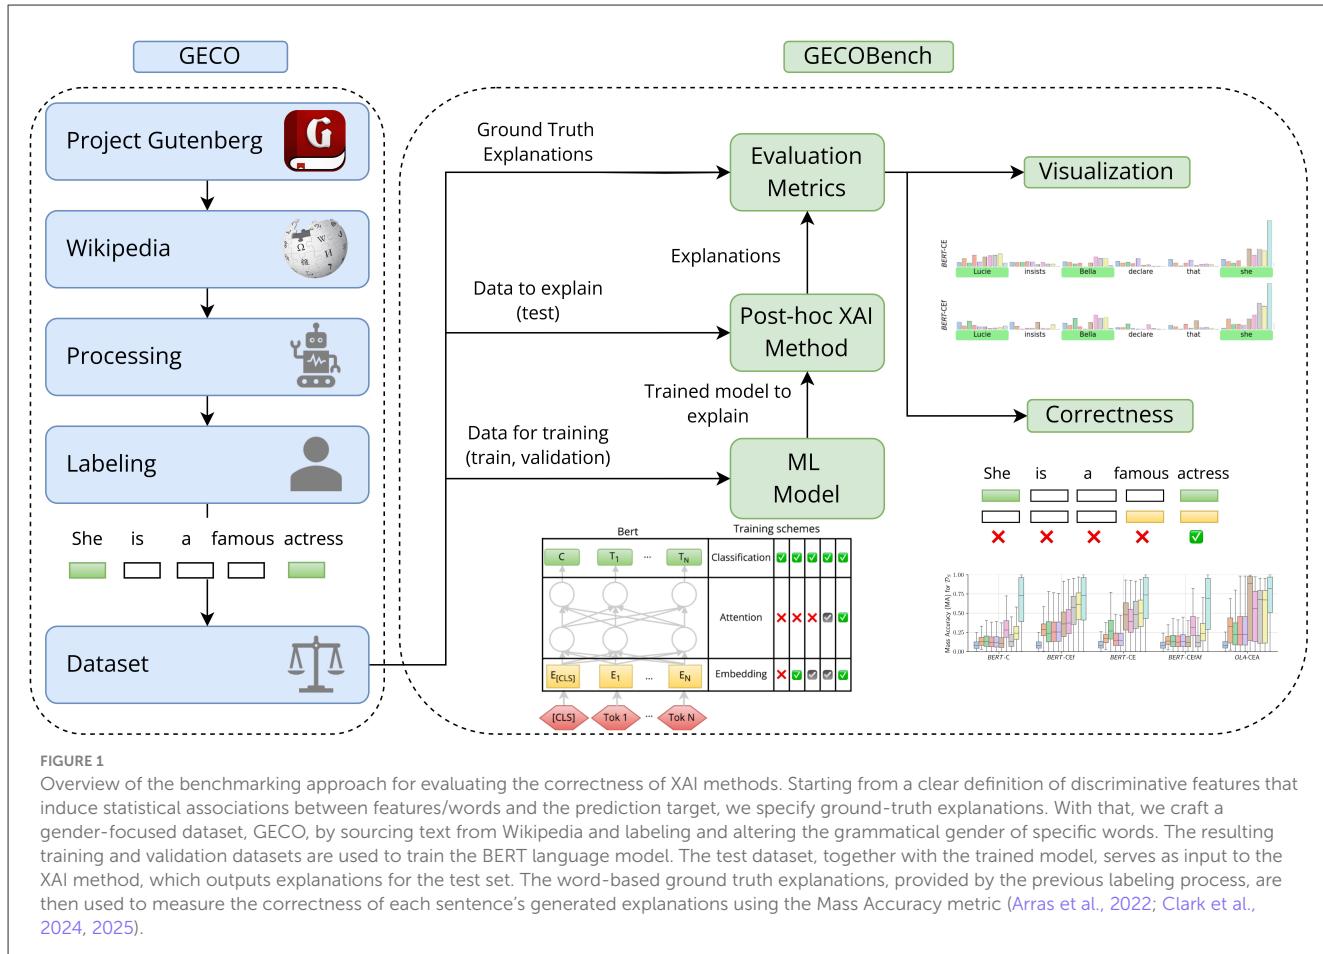
With the gender-controlled dataset *GECO*, we aim to construct sentences composed of discriminative features (or words) and, therefore, ground truth feature attributions that are gender-balanced concerning the classification task. Further, it is known that BERT suffers from gender biases (Nadeem et al., 2021; Ahn and Oh, 2021). Thus, when using *GECO* as a test set, any residual asymmetry in feature attributions can be traced back to biases induced by pre-training. Via the data-centric notion of correctness, we quantify this effect for different stages of retraining or fine-tuning distinct layers of BERT’s architecture to investigate to what extent retraining or fine-tuning BERT on gender-controlled data can mitigate gender bias in feature attributions. In other words, we analyze to what extent the correctness of feature attribution is indicative of biases in models. By ensuring that the distinctly trained models have equivalent classification accuracy throughout the considered fine-tuning stages, we can assess how these training regimes impact correctness performance with the proposed dataset. Generally, we do not expect any XAI method to perform perfectly, as data-centric correctness is only one goal of interpreting machine learning models and not necessarily the primary purpose of each explanation approach. Using *GECO* and *GECOBench*, we aim to answer the following two main research questions:

<sup>1</sup> <https://www.gutenberg.org>

<sup>2</sup> <https://github.com/jcpeterson/openwebtext>

<sup>3</sup> Available on OSF: [https://osf.io/74j9s/?view\\_only=8f80e68d2bba42258da325fa47b9010f](https://osf.io/74j9s/?view_only=8f80e68d2bba42258da325fa47b9010f).

<sup>4</sup> All code, including dataset generation, model training, evaluation, and visualization, is available at: <https://github.com/braindatalab/gecobench>.



RQ1: What is the performance of widely adopted XAI methods in the regime of data-centric feature importance given word/token-level ground truth feature attributions?

RQ2: Does gender bias contained in pre-trained language models affect data-centric explanation performance of feature attribution methods, and if so, does this effect depend on the selection of layers that are fine-tuned or re-trained?

## 2 Related research

Although the applications of XAI have increased in the past years (e.g. Lundberg et al., 2018; Jiménez-Luna et al., 2020; Tran et al., 2021; Zhang et al., 2022), the problems to be addressed by XAI have rarely been formally defined (Murdoch et al., 2019). In particular, the widely used metaphor of identifying features “used” by a model, measured through “faithfulness” or “fidelity” metrics (e.g. Jacovi and Goldberg, 2020; Hooker et al., 2019; Rong et al., 2022), can lead to fundamental misinterpretations, as such a notion depends strongly on the structure of the underlying data generative model and the resulting distribution of the (training) data (Haufe et al., 2014; Wilming et al., 2023; Haufe et al., 2024). Wilming et al. (2023) investigate such metrics, showing that many perturbation and pixel-flipping methods fail to detect statistical dependencies or

other feature effects like suppressor variables (Friedman and Wall, 2005; Haufe et al., 2014), and are therefore unsuitable to directly measure certain meaningful notions of explanation “correctness.” To objectively evaluate whether a feature attribution method possesses this property, the availability of ground truth data is instrumental. Ground truth data for feature attributions in domains such as image, tabular, and time series data have been developed in the last few years (e.g. Kim et al., 2018; Ismail et al., 2019, 2020; Tjoa and Guan, 2023; Agarwal et al., 2022; Arras et al., 2022). However, most of these benchmarks do not present realistic correlations between class-dependent and class-agnostic features (e.g., the foreground or object of an image vs. the background) (Clark et al., 2024), and often use surrogate metrics, such as faithfulness, instead of directly measuring explanation performance. Other works discuss the need for normative frameworks (Sullivan, 2024) or studying data manipulation and its impact on XAI methods’ output (Mhasawade et al., 2024), rather than focusing on ground truth feature attributions. Several NLP-related benchmarks have been presented (DeYoung et al., 2020; Rychener et al., 2020); however, they also have certain limitations. In the case of DeYoung et al. (2020), faithfulness of the model is measured in alignment with human-annotated rationales, which do not necessarily align with statistical association—opening the door to cognitive biases. Rychener et al. (2020) presents a benchmark dataset consisting of a question-answering task, where the ground truth feature attributions originate from a text context providing the answer.

However, as the authors emphasize, defining a ground truth for question-answering cannot depend on a single word but rather on a context of words that provides the prediction models with sufficient information. This work, therefore, does not provide word-level ground truth feature attributions in the sense of statistical association. [Balagopalan et al. \(2022\)](#) and [Dai et al. \(2022\)](#) analyze the fairness behavior of XAI methods, focusing on model fidelity, highlighting disparities between social groups, rather than considering the correctness aspect of feature attributions.

Moreover, [Joshi et al. \(2024\)](#) propose a mitigation technique for gender bias in natural language generation based on feature attribution methods' output. However, since no token-level ground truth is provided, neither the correctness nor the selection-ability of biased tokens by feature attributions can be verified. [Gamboa and Lee \(2024\)](#) introduce the bias attribution score, an information-theoretic metric for quantifying token-level contributions to biased behavior in multilingual pre-trained language models, demonstrating the presence of sexist and homophobic biases in these models. Unlike GECO, neither a controlled counterfactual dataset nor ground truth attributions for evaluating the correctness of feature attribution methods are provided. [Dehdarirad \(2025\)](#) propose a unified framework for evaluating feature attribution methods in language classification models, comparing SHAP, LIME, Integrated Gradients, and interaction-based approaches across classical and transformer architectures to assess their faithfulness ([Samek et al., 2019](#); [Jacovi and Goldberg, 2020](#)) under different datasets. However, a controlled dataset with ground truth attributions is not provided. Given the tension between model-centric and data-centric feature attributions, ground-truth-based evaluations, as in GECOBench, enable a more principled study of feature attribution methods in the context of language models.

In the NLP research community, the development of datasets for bias detection, metrics for fairness and bias assessment, and methods for bias mitigation is an active field of research. For example, [Bolukbasi et al. \(2016\)](#) demonstrated that word embeddings encode gender stereotypes and proposed subspace-based debiasing, specifically learning a “gender direction” and projecting it out from gender-neutral words. This was refined by [Prost et al. \(2019\)](#), [Dev et al. \(2020\)](#) utilizes natural language inference as a surrogate to study and mitigate biased inferences arising from embeddings systematically. Further benchmarks for social bias analysis in large pre-trained language models have been proposed, for example, in [Manzini et al. \(2019\)](#), [Nangia et al. \(2020\)](#), [Costa-jussà et al. \(2020\)](#), [Nadeem et al. \(2021\)](#), [Parrish et al. \(2022\)](#), [Jentsch and Turan \(2022\)](#), [Zakizadeh et al. \(2023\)](#), [Navigli et al. \(2023\)](#), and [Cimitan et al. \(2024\)](#). Nevertheless, identical sentences for each grammatical gender, which differ only in specific and controlled positions across grammatical gender forms, providing a ground truth for feature attribution benchmarking, are not currently possible based on these datasets and benchmarks.

### 3 Materials and methods

To enable correctness evaluations for explanation methods, we introduce the GECO dataset, which comprises a set of manipulated sentences  $x$  in which grammatical subjects and objects assume either their male  $x^M$ , female  $x^F$ , or non-binary

$x^{NB}$  forms. These three grammatically gendered variants give rise to the downstream task of gender classification, with labels “M,” “F,” and “NB,” which involves discriminating between the variants of sentences and is represented by the dataset  $D := \{(x^{(i,M)}, “M”), (x^{(i,F)}, “F”), (x^{(i,NB)}, “NB”) \}_{i=1}^n$ . Importantly, in all cases, ground truth feature attributions on a word-level basis are available by construction.

#### 3.1 Data sourcing and generation

For the dataset, we restrict ourselves to source sentences with a human subject, such that each sentence of our manipulated dataset is guaranteed to have a well-defined gender label. This type of sentence naturally occurs in books and novels. The Gutenberg archive offers a vast collection of classical titles, enabling users to identify relevant text content from well-known novels and nonfiction works. To comply with licensing requirements surrounding the listed books, we collect the content of their corresponding Wikipedia pages and use only text pieces related to the plot of the story. We query the list of the top 100 popular books on the Gutenberg project and obtain their corresponding Wikipedia pages. More details on data licensing are provided in [Supplementary material 1.1.1](#).

We create two ground truth data sets  $D_S$  and  $D_A$ . Each contains 1,610 sentences in a male, a female, and a non-binary version, comprising 4830 sentences in total (see [Table 1](#)).  $D_S$  contains sentences in which *only* words specifying the gender of the grammatical subject are manipulated to be either in male  $x_S^M$ , female  $x_S^F$ , or non-binary  $x_S^{NB}$  form, while  $D_A$  contains sentences in which all gender-related words are manipulated, which we denote as  $x_A^M$ ,  $x_A^F$ , or  $x_A^{NB}$ , respectively. [Table 2](#) shows an exemplary sentence and the resulting manipulations employing this labeling scheme. The dataset  $D_S$  instantiates a substantially more challenging task compared to dataset  $D_A$  due to the reduction of discriminative features. In this scenario, the model is required to differentiate between subject and object when they have different grammatical genders, necessitating a deeper understanding of the sentence's context and structure to address the task effectively. Thus, employing both types of datasets allows us to investigate whether the model inadvertently focuses on irrelevant parts of the sentence for the prediction task, potentially introducing bias that could impact explanation performance. The process for creating these datasets consists of two consecutive steps: (i) Preprocessing of scraped Wikipedia pages. (ii) Manual labeling is used to detect and adapt relevant subjects and objects in a sentence. More details on labeling and format are provided in [Supplementary material 1.1.2](#). Further details on hosting and future maintenance are provided in the [Supplementary material](#).

#### 3.2 Bias assessment of GECO

We employ a co-occurrence metric ([Zhao et al., 2017](#)), a rudimentary bias measure, highlighting the unbiasedness of GECO. Specifically, we adopt the co-occurrence metric proposed by [Cabello et al. \(2023\)](#) to measure gender bias in the datasets

TABLE 1 Overview of the GECO dataset properties.

Dataset	Manipulation scope	# Base sent.	# Variants	Total Sent.	# Train sent.	# Test sent.
$D_S$	Subject words only	1,610	3 (M, F, NB)	4,830	3,864	966
$D_A$	All gender-related words	1,610	3 (M, F, NB)	4,830	3,864	966

Each dataset variant contains sentences in three gendered forms: male (M), female (F), and non-binary (NB).

TABLE 2 Example of the labeling and alteration scheme of sentences, showing the original sentence and the six manipulated versions.

Version	Sentence						
Original	She	touches	the	heart	of	her	Aunt.
Subj. Female ( $x_S^F$ )	She	touches	the	heart	of	her	Aunt.
Subj. Male ( $x_S^M$ )	He	touches	the	heart	of	his	Aunt.
Subj. Non-binary ( $x_S^{NB}$ )	They	touch	the	heart	of	their	Aunt.
All Female ( $x_A^F$ )	She	touches	the	heart	of	her	Aunt.
All Male ( $x_A^M$ )	He	touches	the	heart	of	his	Uncle.
All Non-binary ( $x_A^{NB}$ )	They	touch	the	heart	of	their	Parent's Sibling.

Words marked as the ground truth for feature attributions are written in bold and color-coded depending on the grammatical gender.

$D_S$  and  $D_A$ . For a given sentence  $x \in D$ , we approximately measure the bias induced by grammatical gender when considering the co-occurrence between the sentence's gender terms and the remaining words. First, we define a set of grammatical gender terms  $A := \{\text{"she"}, \text{"her"}, \text{"he"}, \text{"they"}, \dots\}$  and second, a word vocabulary without grammatical gender terms  $V := W \setminus A$ , where the vocabulary  $W$  contains all words available in a corpus of  $D$ ; then, the co-occurrence metric  $C$  is defined as follows:

$$C(D) = \sum_{x \in D} \sum_{w \in V} \sum_{a \in A} \mathbf{1}_x(a, w). \quad (1)$$

Furthermore, we consider the decomposition  $D = D^M \cup D^F \cup D^{NB}$  with male  $D^M$ , female  $D^F$  and non-binary  $D^{NB}$  sentences, respectively. Then we define the bias, exemplarily for  $D^F$ , according to the co-occurrence metric  $C$  as

$$bias_C(D^F) = \frac{C(D^F)}{C(D^F) + C(D^M) + C(D^{NB})}. \quad (2)$$

A perfect balance is achieved for  $bias_C(D^F) = 1/3$ , indicating that the dataset is evenly distributed among men, women, and non-binary individuals. Deviations from this value indicate the presence of biases: values closer to 0 suggest a male or non-binary bias, while those approaching 1 indicate a female bias.

### 3.3 Explanation benchmarking

The alteration of sentences induces discriminative features by construction, and their uniqueness automatically renders them the only viable ground truth feature attributions, defining two different gender classification tasks represented by the two datasets  $D_S$  and  $D_A$ . By extension, every word that is not grammatically gender-related, and therefore not altered, becomes a non-discriminative feature. We train machine learning models on these tasks and apply *post-hoc* feature attribution methods to the trained models to obtain

explanations expressing the importance of features according to each XAI method's intrinsic criteria. When evaluating the XAI methods, ground truth feature attributions are adduced to measure if their output highlights the correct features. An overview is shown in Figure 1.

#### 3.3.1 Ground truth feature attributions

We consider a supervised learning task, where a model  $f: R^d \rightarrow R$  learns a function between an input  $x^{(i)} \in R^d$  and a target  $y^{(i)} \in \{-1, 1\}$ , based on training data  $D = \{(x^{(i)}, y^{(i)})\}_{i=1}^N$ . Here,  $x^{(i)}$  and  $y^{(i)}$  are realizations of the random variables  $\mathbf{X} = (X_1, X_2, \dots, X_d)^\top$  and  $Y$ , with joint probability density function  $p_{\mathbf{X}, Y}(x, y)$ , and  $[d] := \{1, \dots, d\}$  is the set of feature indices for a vector-wise feature representation  $(X_i | i \in [d])$ . We formally cast the problem of finding an explanation or important features as a decision problem  $([d], F, f)$  where  $F \subseteq [d]$  is the set of *important features*. Moreover, an explanation or saliency map  $s: R^d \rightarrow R^d$  should assign a numerical value reflecting the significance of each feature. Then we are interested in finding a test  $h: R^d \rightarrow \{0, 1\}^d$ , which one can use to define the set of important features  $F := \{j | h_j(x) = 1, \text{ for } j \in [d]\}$ . For the concrete definition of the test and the resulting set of important features, we adopt the approach of Wilming et al. (2022) and Wilming et al. (2023) and give the following definition.

**Definition 3.1 (Statistical Association Property (SAP)).** Given the supervised learning task from above, we say that an XAI method has the Statistical Association Property (SAP) if for any feature  $X_j$  with non-zero (or, significantly larger than zero) importance, there also exists a statistical dependency between  $X_j$  and the target  $Y$ , i.e.,  $X_j \not\perp\!\!\!\perp Y$ .

This definition is based on the discussion that most feature attribution methods implicitly or explicitly assume that such a statistical association exists (Wilming et al., 2022). Now, defining a test via  $h_j(x) = 1$  if  $X_j \not\perp\!\!\!\perp Y$  and  $h_j(x) = 0$  otherwise, we can summarize the set of *potentially* important features via their

univariate statistical dependence with the target  $F = \{j \mid X_j \not\perp\!\!\!\perp Y, \text{ for } j \in [d]\}$ . Thus, each sentence of the GECO corpus and its corresponding token sequence  $x^{(i)}$  has a matching ground truth map  $h(x^{(i)}) \in \{0, 1\}^d$ , representing corresponding important tokens.

### 3.3.2 Feature attribution methods

Here, we focus on *post-hoc* attribution methods, which can be broadly divided into gradient-based methods and local sampling or surrogate approaches. Generally, these methods produce an explanation  $s: R^d \rightarrow R^d$ , which is a mapping that depends on the model  $f$  and an instance  $x^*$  to be explained. Gradient-based methods locally approximate a differentiable model  $f$  around a given input sequence  $x^*$ . From this class, we consider Saliency (Simonyan et al., 2013), InputXGradient (Shrikumar et al., 2017), DeepLift (Shrikumar et al., 2017), Guided Backpropagation (Springenberg et al., 2015), and Integrated Gradients (Sundararajan et al., 2017). Surrogate models, on the other hand, sample around the input  $x^*$  and use a model's output  $f(x)$  to train a simple, usually linear, model and interpret  $f$  through this local approximation. In this work, we consider the surrogate methods LIME (Ribeiro et al., 2016) and Kernel SHAP (Lundberg and Lee, 2017). Additionally, our study includes Gradient SHAP (Lundberg and Lee, 2017), an approximation of the Shapley value sampling method.

We also consider two baselines. Firstly, we set the explanation for a particular input sequence  $x^*$  to uniformly distributed noise  $s(x^*) \sim \mathcal{U}(0, 1)^d$ . This serves as a null model corresponding to the hypothesis that the XAI method has no knowledge of the informative features  $h(x^*)$ . Secondly, we employ the Pattern approach (Haufe et al., 2014; Wilming et al., 2022). We apply a variant of it by employing the covariance between input features and target  $s_j(x^*) = \text{Cov}(x_j^*, y)$ . We call this the Pattern Variant, for which we utilized the tf-idf (Sparck Jones, 1972) representation of each input sequence  $x^{(i)}$ . Clearly, the explanation  $s$  is independent of both the model  $f$  and an instance  $x^*$ ; therefore, it yields the same feature attributions for all input sequences.

We apply these XAI methods to all fine-tuning variants of the BERT model and compute explanations on all test data sentences using the default parameters of each method. For all XAI methods except LIME, we use their Captum (Kokhlikyan et al., 2020) implementation. For LIME, we use the author's original code.<sup>5</sup>

### 3.3.3 Explanation performance quantification

For a given instance  $x^* \in D^{test}$  we aim to quantitatively assess the correctness of its explanation  $s(x^*)$ . The corresponding ground truth  $h(x^*)$  defines a set of potentially important tokens based on alteration of words; however, a model  $f$  might only use a subset of such tokens for its predictions. Hence, an explanation method that only highlights a subset of tokens that correspond to ground truth, compared to all tokens of the ground truth, must be considered equally correct. Expressed in information retrieval terms, we are interested in mitigating the impact of false negatives and emphasizing the impact of false positives on explanation performance. False negatives occur when a token flagged as part

of the ground truth receives a low importance score, and false positives occur when a feature flagged as not part of the ground truth receives a high importance score. The Mass Accuracy metric (*MA*) (Arras et al., 2022; Clark et al., 2024) provides such properties and is defined as

$$MA(h(x^*), s(x^*)) = \sum_{j=1}^d s_j(x^*) h_j(x^*). \quad (3)$$

Here the feature attributions  $s$  are normalized, such that  $\sum_j^d s_j = 1$  and  $s(x^*) \in [0, 1]^d$ . The score  $MA(h(x^*), s(x^*)) = 1$  shows a perfect explanation, marking only ground truth tokens as important. For instance, a sentence with only two ground truth tokens  $h(x^*) = (1, 1, 0, 0)^\top$ , where the attribution for only one ground truth token is high, say  $s(x^*) = (0.9, 0, 0, 0.1)$ , the *MA* metric still produces a high score of  $MA(h(x^*), s(x^*)) = 0.9$ , de-emphasizing false negatives. With respect to false positives, high attributions to non-ground-truth tokens do not directly contribute to the *MA*, yet, through the normalization of  $s$ , all other tokens get assigned a relatively low (non-zero) importance, leading to an overall low *MA* score, effectively penalizing false-positive attributions.

Note, feature attributions  $s$  are calculated at the sub-word level. To align them with word-level ground truth, we normalize attribution scores across a sentence and then aggregate sub-word contributions back to the word level. For example, the word "benchmark" may be split into "bench" and "mark" by the BERT tokenizer, with attributions  $s_{\text{bench}}$  and  $s_{\text{mark}}$ , which are combined to  $s_{\text{benchmark}} = s_{\text{bench}} + s_{\text{mark}}$ .

### 3.3.4 Explanation bias quantification

The change in mass accuracy *MA* serves two purposes: (i) It assesses the correctness of XAI methods' output with respect to the ground truth, and (ii) with deviations from the ground truth, depending on which layer is fine-tuned or retrained, we can define a notion of what we call *explanation bias*. Explanation bias is defined via the relative mass accuracy (*RMA*)

$$RMA_{\text{baseline}} := \frac{MA}{MA_{\text{baseline}}}, \quad (4)$$

where the deviation of explanation performance with regard to a baseline model is quantified, a zero-shot BERT model is used in this work (see next Section 3.3.5).

### 3.3.5 Classifiers

In our analysis, we focus on the popular BERT model (Devlin et al., 2019), though one can expand this work using other common language models such as RoBERTa (Liu et al., 2019), XLNet (Yang et al., 2019), or GPT models (Radford et al., 2018, 2019). For all experiments, we use the pre-trained uncased BERT model (Devlin et al., 2019).<sup>6</sup> To investigate the impact that fine-tuning or retraining of different parts of BERT's architecture

<sup>5</sup> <https://github.com/marcotcr/lime>

<sup>6</sup> Hosted by Hugging Face: <https://huggingface.co/google-bert/bert-base-uncased>.

can have on explanation performance, we consider four different training paradigms: (i) We roughly split BERT’s architecture into three parts: *Embedding*, *Attention*, and *Classification*. The standard approach to adopting BERT for a new downstream task is to train the last classification layer, which we refer to as *Classification*, while fixing the weights for all remaining parts of the model, specifically *Embedding* and *Attention*. We thereby only train a newly initialized classification layer and refer to the resulting model as *BERT-C*. (ii) We additionally train the embedding layer from scratch, resulting in a model called *BERT-CE*. (iii) In the third model, *BERT-CEf*, the embeddings are fine-tuned as opposed to newly initialized. In training paradigm (iv), we fine-tune the *Embedding* and *Attention* parts of BERT’s architecture, resulting in model *BERT-CEfAf*. Moreover, a zero-shot model *BERT-ZS*, which experienced no gradient updates, was applied. Lastly, a vanilla one-layer attention model, *OLA-CEA*, comprising a lower-dimensional embedding layer, one attention layer, and a classification layer, was trained from scratch only on the GECO dataset. Therefore, without pre-training on external corpora, it represents the simplest attention-based model free from residual biases, providing a clean reference point against which more complex, pre-trained models like BERT can be compared. All models achieve an accuracy above or close to 80% on the test set. Previous works on classification problems involving BERT suggest that accuracy results ranging from 60 to 90% are standard (Gao et al., 2019; Zheng and Yang, 2019; Yu and Jiang, 2019). Therefore, we consider our results as evidence that the model has successfully generalized to the given downstream task. Table 3 summarizes model performance with average accuracy and standard deviation over five models trained with different seeds. More details are given in *Supplementary material 1.2.1* and experiments’ configuration file.<sup>7</sup>

## 4 Experiments and results

Our bias analysis on GECO shows that there is no gender bias present in the  $D_A$  dataset with  $bias_C(D_A^F) = bias_C(D_A^M) = bias_C(D_A^{NB}) \approx 0.33$ . For dataset  $D_S$  we achieve the scores  $bias_C(D_S^{NB}) \approx 0.3$ ,  $bias_C(D_S^M) \approx 0.36$ , and  $bias_C(D_S^F) \approx 0.33$ . The small difference from a perfect score can be attributed to labeling errors but is also expected for the dataset  $D_S$  due to its construction, as we only change the human subject of the sentence; other gender terms, referring to other protagonists in the sentence, are kept unchanged.

Using the unbiased dataset GECO, we conduct experiments to study the influence of biased models on explanation performance. After fine-tuning and re-training the models (see Table 3), we apply feature attribution methods. Figure 2 shows the explanation performance of sample-based attribution maps  $s(x^*)$  produced by the selected feature attribution and baseline methods.

In the following, we present the results toward the research questions RQ1 and RQ2. Firstly, we present the results for explanation performance, addressing questions about the data-centric correctness of the analyzed feature attribution methods.

Secondly, we present bias-related results, focusing on how the correctness of feature attribution indicates biases in language models.

**Regarding RQ1:** We observe a general difference in *MA* between datasets  $D_A$  and  $D_S$ . While for the majority of attribution methods, the performance for experiments on dataset  $D_S$  stays at a level lower than 0.25, experiments on dataset  $D_A$  are often able to offset these results into levels above 0.25 (see Figure 2b). However, dataset  $D_S$  has fewer altered gender words, thus fewer discriminative tokens, leading to an overall degradation of classification accuracy across all models (see Table 3), which also impacts explanation performance. For all BERT models and both datasets, Integrated Gradients consistently outperforms other methods compared to the uniform random baseline. Though, LIME and Gradient SHAP, are highest-performing methods as well, compared to the Pattern Variant baseline, and with respect to the data-centric SAP criterion (see Definition 3.1).

Comparing the *OLA-CEA* model to all BERT models, we observe a stark contrast in explanation performance. Recall that the *OLA-CEA* model was purely trained from scratch on the gender-controlled dataset GECO; hence, it does not suffer from any gender bias. The mass accuracy for the *OLA-CEA* model is similar between the two datasets, with higher variance for the dataset  $D_S$ . In addition to relatively well-performing methods such as Integrated Gradients, LIME, and Gradient SHAP, the *MA* of InputXGradient comes very close to the Pattern Variant baseline, making it the best-performing method.

Although no explanation method achieves the correctness score of Pattern Variant, fine-tuning a biased embedding layer for a downstream task has a high impact on the output for some methods. The Pattern Variant is a model-independent global explanation method that relies solely on the intrinsic structure of the data itself. It performs optimally when the feature and target relation is governed by a linear relationship, which is mainly the case for GECO. This can be seen in Figure 4 of the *Supplementary material*, where we visualize the Pearson correlation of the term frequency-inverse document frequency (tf-idf) (Sparck Jones, 1972) representation of words and the target, clearly showing how we infuse dependency through the word alteration procedure.

Exemplarily, Figure 6 in the *Supplementary material* highlights a sentence labeled as “female,” which is shown together with its word-based feature attributions as bar plots for each fine-tuning stage. We observe high variability in token attribution between differently fine-tuned BERT models, and the pronoun “she” receives relatively high importance compared to other words. However, not all XAI methods agree on the importance of the token “she”; for example, for model *BERT-CE* InputXGradient attributes high importance to it, yet for model *BERT-CEfAf*, it attributes rather high importance to the word “Bella.”

**Regarding RQ2:** In Figure 2a we can observe a consistent pattern with respect to the fine-tuning stages across both data scenarios  $D_S$  and  $D_A$  in terms of *RMA<sub>BERT-ZS</sub>* performance. Here, *BERT-ZS* is utilized as a baseline model, as it represents the “untouched” pre-trained model without gradient updates. Using it as a baseline allows us to quantify how the performance of feature attribution correctness evolves as models become increasingly specialized in the gender-classification task.

<sup>7</sup> [https://osf.io/74j9s/files/p23yh?view\\_only=8f80e68d2bba42258da325fa47b9010f](https://osf.io/74j9s/files/p23yh?view_only=8f80e68d2bba42258da325fa47b9010f)

TABLE 3 Overview of BERT transfer learning paradigms and the performance of the resulting models on the test datasets  $D_A^{test}$  and  $D_S^{test}$ .

Models	Embedding	Attention	Classification	Acc. (%) $D_A^{test}$	Acc. (%) $D_S^{test}$
<i>BERT</i> -ZS	Fix	Fix	Fix	$65.2 \pm 0.0$	$56.6 \pm 0.0$
<i>BERT</i> -C	Fix	Fix	Re-trained	$98.7 \pm 0.1$	$90.2 \pm 0.4$
<i>BERT</i> -CE	Re-trained	Fix	Re-trained	$99.8 \pm 0.1$	$96.6 \pm 0.2$
<i>BERT</i> -CEf	Fine-tuned	Fix	Re-trained	$99.8 \pm 0.1$	$97.4 \pm 0.3$
<i>BERT</i> -CEfAf	Fine-tuned	Fine-tuned	Re-trained	$99.7 \pm 0.1$	$98.8 \pm 0.2$
OLA-CEA	Re-trained	Re-trained	Re-trained	$99.5 \pm 0.2$	$85.1 \pm 14.6$

Specifically, it highlights how fine-tuned models focus more on the discriminatory tokens compared to the zero-shot model, thereby providing a relative quantification of residual biases. For the scenario  $D_A$ , it is clear that the models *BERT*-CE and *BERT*-CEf, where the embedding layer was trained or fine-tuned, respectively, outperform *BERT*-C and *BERT*-CEfAf (see also Figure 2b). This shows that the embeddings encode numerous bias information and, indeed, influence data-centric explanation performance.

## 5 Discussion

With GECO and GECOBench, we propose an open framework for benchmarking the correctness of feature attributions of pre-trained language models as well as aspects of fairness. Our initial results demonstrate (a) differences in explanation performance between feature attribution methods, (b) a general dependency of explanation performance on the amount of re-training/fine-tuning of BERT models, and (c) residual gender biases as contributors to sub-par explanation performance.

More generally, the proposed gender classification problem is a simplification that does not reflect the complexity and diversity of gender identification in our world today; however, by providing non-binary gendered sentences, we attempt to counteract historical gender norms and provide a more inclusive basis for gender-bias research. We view the gender classification task as a minimal proxy for the gender bias issue, modeling all necessary properties to analyze bias propagation into feature attribution methods. We also view predicting a sentence's gender as an auxiliary task, which we consider more of an academic problem that naturally arises from how we construct sentences but has, as we see it, no immediate application or societal impact. While GECO provides a controlled environment to study how gender bias influences feature attributions, we acknowledge that our design inevitably oversimplifies gender by restricting it to pronoun alterations (e.g., he/she/they). This simplification risks reinforcing notions of gender that fail to represent the full spectrum of identities. In addition, gender-controlled datasets such as GECO could be misused, for example, to build or evaluate models explicitly aimed at gender classification rather than for bias analysis. To mitigate this risk, we emphasize that GECO is intended solely for studying the correctness assessment of feature attribution methods under controlled bias conditions, not for downstream

applications involving sensitive demographic prediction. We view GECO as a first step toward systematic evaluation of bias in feature attributions, with the expectation that future work will extend its coverage to more diverse text sources, richer notions of gender, and broader fairness concepts. While direct extrapolation to more complex applications is challenging, our results indicate that feature attributions are indeed affected by gender bias. This motivates caution in downstream tasks, such as sentiment analysis or toxic language detection, where attribution methods might incorrectly highlight some gender-related tokens due to bias rather than semantic relevance. For example, this renders model debugging tasks challenging, as developers and researchers cannot distinguish between feature attributions that suggest “flaws” in the model, arising from genuine model bias or from other artifacts of the data.

When it comes to data selection, we are convinced that if biased models are applied to semantically gender-neutral sentences, no reliance on words representing classical gender roles can be expected; thus, the potential impact of biases on feature attributions cannot be measured. Thus, the sentences selected to create GECO were intentionally taken from Wikipedia articles outlining the storylines of classic novels, as they were likely to employ historical gender norms. These are, indeed, prerequisites for assessing the behavior of XAI methods' output applied to language models experiencing various levels of biases.

By creating grammatical female, male, and non-binary versions of a particular sentence based on pronouns, we are aiming to break such historical gender associations with respect to the *classification task* represented by the datasets  $D_S$  and  $D_A$ . Single sentences may still entail historical gender associations, which can be utilized by biased machine learning models. However, the classification task arising from GECO is gender-balanced, and models specialized in that task, through successive fine-tuning and retraining of an increasing number of layers, learn not to rely on such historical gender associations, as altered pronouns are, by construction, the only words associated with the prediction target. It can then be shown that models basing their decisions on words other than the words altered by us have learned stereotypical associations. For example, consider the sentence “She prepares dinner in the kitchen while he is outside fixing the car.” This sentence illustrates “traditional” gender roles, where the woman is associated with domestic tasks and the man with mechanical or manual labor, reinforcing stereotypes. By altering the pronouns to use only “she” or “he” pronouns, we break these stereotypical gender roles.

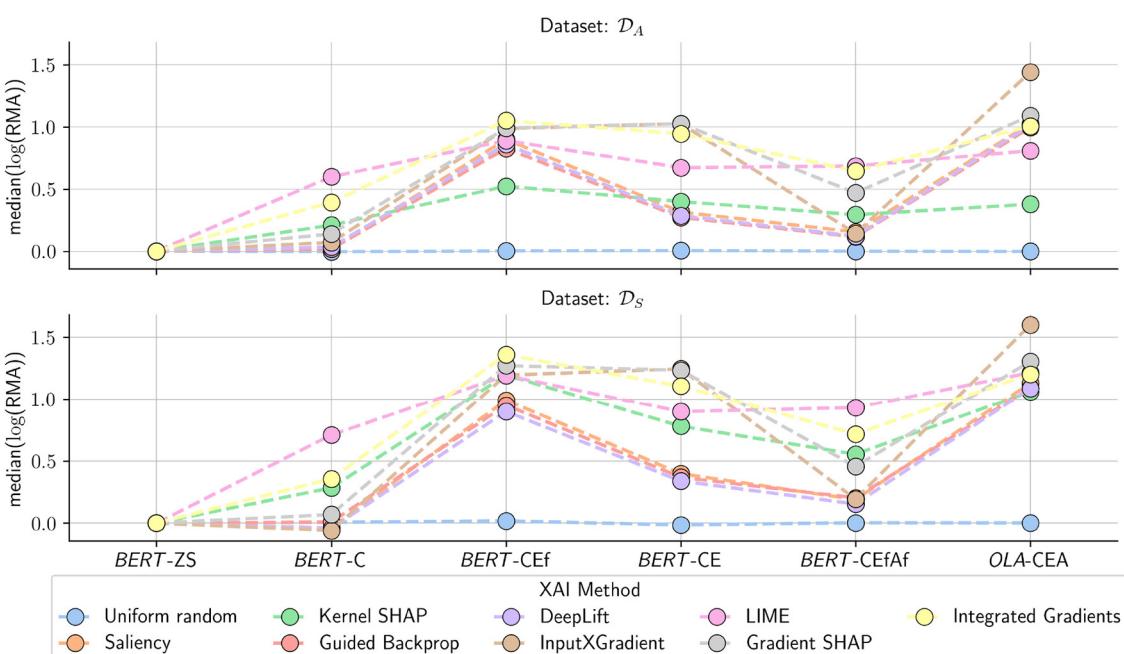
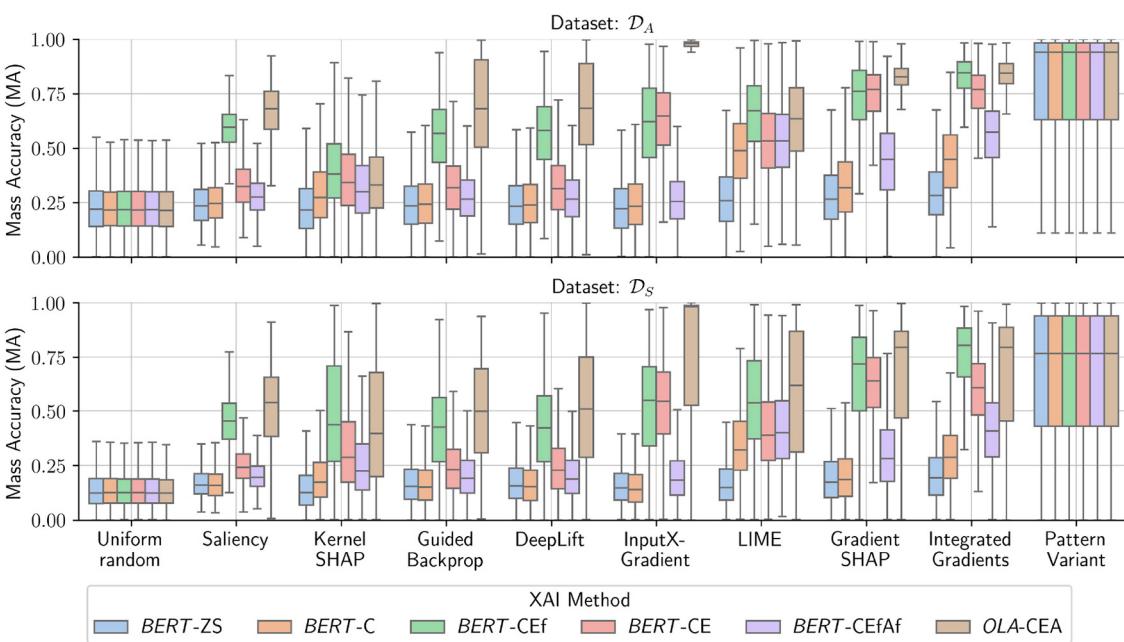
**a.****b.**

FIGURE 2

Explanation performance of different *post-hoc* XAI methods applied to language models that were adapted from BERT using five different transfer learning schemes. XAI evaluations were conducted on classified sentences in two gender classification tasks, represented by datasets  $D_S$  and  $D_A$ . The baseline performance for uniformly drawn random feature attributions is denoted by *Uniform Random*. *Pattern Variant* denotes a model- and pretraining-agnostic global explanation method. In (a), the relative change in explanation performance with respect to a zero-shot BERT model shows consistent changes for models with fine-tuned embeddings. In (b), fine-tuning or retraining of the embedding layers of BERT leads to consistent improvements in explanation correctness even when model performance is held constant for all models. Applying XAI methods to the *OLA* model leads to overall higher explanation performance, with *InputXGradient* becoming on par with *Pattern Variant*.

However, only partially in this instance will one part of the sentence always reflect traditional gender roles, e.g., “She prepares dinner in the kitchen... .” Nevertheless, for the classification task

represented by the datasets  $D_S$  and  $D_A$ , only the altered words represent a relationship with the prediction target. Biased language models might then leverage words like “kitchen” or “car” for their

decision, and unbiased models must only rely on altered words, and historical gender norms embedded in sentences become irrelevant. In future research, co-reference resolution could be an immediate extension of GECO because it takes the same gender-manipulated sentences and asks the model not just to classify gender per sentence but to resolve references consistently across discourse, thereby testing explanation correctness under contextual and bias-sensitive conditions.

In terms of data-centric correctness assessments of feature attribution methods, Pattern Variant indeed offers strong theoretical justification for detecting important features according to statistical associations (Haufe et al., 2014), establishing a solid baseline for the upper bound of explanation performance in our benchmark. Compared to the random baseline, we observe two further high-performing attribution methods in the transfer learning regime [in terms of SAP (see Definition 3.1)]: Integrated Gradients and Gradient SHAP. Nevertheless, these methods still do not achieve the same level of accuracy as Pattern Variant. The reasons can be two-fold: (i) As shown by Clark et al. (2024) and Wilming et al. (2023), feature attribution methods consistently attribute importance to suppressor variables, features not statistically associated with the target but utilized by machine learning models to increase accuracy. And (ii), model bias impacts feature attributions. We show that the gender bias in BERT leads to residual asymmetries in feature attributions and forms a consistent pattern of deviation in correctness, depending on which layer of BERT was fine-tuned or re-trained, while still achieving equivalent classification accuracy. As a result, updating embedding layers has the strongest impact on feature attributions. These findings indicate that embeddings contain significant bias affecting feature attribution methods, and that the proposed data-centric notion of correctness of feature importance is indicative of model bias.

While this is, to the best of our knowledge, the first XAI benchmark addressing a well-defined notion of data-centric correctness of feature importance in the NLP domain, we do not consider it an exhaustive evaluation of feature attribution methods but rather a first step toward this. A possible limitation of our approach is that the criterion of univariate statistical association used here to define important features or tokens does not account for nonlinear feature interactions that are prevalent in many real-world applications. However, for analyzing the fundamental behaviors of feature attribution methods, this characteristic allows for straightforward evaluation strategies, permitting us to embed these statistical properties into the proposed corpus and establish a ground truth of word relevance. Designing metrics for evaluating explanation performance, particularly for measuring correctness, is another area that warrants further research.

## 6 Conclusion

We have introduced GECO—a novel gender-controlled ground truth text dataset designed for the development and evaluation of feature attribution methods—and GECOBench—a quantitative benchmarking framework to perform objective assessments of explanation performance for language models. We demonstrated the use of GECO and GECOBench by applying them to the pre-trained language model BERT, a model known to exhibit

gender biases. With this analysis, we showed that the SAP criterion is an effective condition to quantify the data-centric correctness of feature attribution methods applied to the language model BERT, and that residual biases contained in BERT affect feature attributions and can be mitigated through fine-tuning and retraining of different layers of BERT, positively impacting explanation performance.

## Data availability statement

The datasets presented in this study can be found in online repositories. The names of the repository/repositories and accession number(s) can be found in the article/[Supplementary material](#).

## Author contributions

RW: Data curation, Formal analysis, Methodology, Software, Validation, Visualization, Writing – original draft, Writing – review & editing. AD: Data curation, Formal analysis, Methodology, Software, Validation, Visualization, Writing – original draft, Writing – review & editing. HS: Data curation, Formal analysis, Methodology, Software, Validation, Visualization, Writing – original draft, Writing – review & editing. MO: Conceptualization, Data curation, Software, Validation, Writing – original draft, Writing – review & editing. BC: Formal analysis, Methodology, Software, Validation, Writing – original draft, Writing – review & editing. SH: Conceptualization, Funding acquisition, Investigation, Methodology, Project administration, Resources, Supervision, Writing – original draft, Writing – review & editing.

## Funding

The author(s) declare that financial support was received for the research and/or publication of this article. This result was part of a project that has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (Grant agreement No. 758985), and the German Federal Ministry for Economy and Climate Action (BMWK) within the framework of the QI-Digital Initiative.

## Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

## Generative AI statement

The author(s) declare that no Gen AI was used in the creation of this manuscript.

Any alternative text (alt text) provided alongside figures in this article has been generated by Frontiers with the support of artificial intelligence and reasonable efforts have

been made to ensure accuracy, including review by the authors wherever possible. If you identify any issues, please contact us.

## Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the

reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

## Supplementary material

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/frai.2025.1694388/full#supplementary-material>

## References

Agarwal, C., Krishna, S., Saxena, E., Pawelczyk, M., Johnson, N., Puri, I., et al. (2022). "OpenXAI: towards a transparent evaluation of model explanations," in *Thirty-sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.

Ahn, J., and Oh, A. (2021). "Mitigating language-dependent ethnic bias in BERT," in *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, Online and Punta Cana, Dominican Republic*, eds. M.-F. Moens, X. Huang, L. Specia, and S. W.-t. Yih (Punta Cana: Association for Computational Linguistics), 533–549. doi: 10.18653/v1/2021.emnlp-main.42

Arras, L., Osman, A., and Samek, W. (2022). CLEVR-XAI: a benchmark dataset for the ground truth evaluation of neural network explanations. *Inf. Fusion* 81, 14–40. doi: 10.1016/j.inffus.2021.11.008

Balagopalan, A., Zhang, H., Hamidieh, K., Hartvigsen, T., Rudzicz, F., Ghassemi, M., et al. (2022). "The road to explainability is paved with bias: measuring the fairness of explanations," in *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency, FAccT '22* (New York, NY: Association for Computing Machinery), 1194–1206. doi: 10.1145/3531146.3533179

Beukeboom, C. J. (2014). "Mechanisms of linguistic bias: how words reflect and maintain stereotypic expectancies," in *Social Cognition and Communication*, eds. J. P. Forgas, O. Vincze, and J. László (New York, NY: Psychology Press), 313–330.

Blodgett, S. L., Barocas, S., Daumé, I. I. I., and Wallach, H. H. (2020). "Language (technology) is power: a critical survey of 'bias' in NLP," in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, eds. D. Jurafsky, J. Chai, N. Schluter, and J. Tetreault (Association for Computational Linguistics), 5454–5476. doi: 10.18653/v1/2020.acl-main.485

Bolukbasi, T., Chang, K.-W., Zou, J., Saligrama, V., and Kalai, A. (2016). "Man is to computer programmer as woman is to homemaker? Debiasing word embeddings," in *Proceedings of the 30th International Conference on Neural Information Processing Systems, NIPS'16* (Red Hook, NY: Curran Associates Inc), 4356–4364.

Bordia, S., and Bowman, S. R. (2019). "Identifying and reducing gender bias in word-level language models," in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Student Research Workshop*, eds. S. Kar, F. Nadeem, L. Burdick, G. Durrett, and N.-R. Han (Minneapolis: Association for Computational Linguistics), 7–15. doi: 10.18653/v1/N19-3002

Bugliarello, E., Cotterell, R., Okazaki, N., and Elliott, D. (2021). Multimodal pretraining unmasked: a meta-analysis and a unified framework of vision-and-language BERTs. *Trans. Assoc. Comput. Linguist.* 9, 978–994. doi: 10.1162/tacl\_a\_00408

Cabello, L., Bugliarello, E., Brandl, S., and Elliott, D. (2023). "Evaluating bias and fairness in gender-neutral pretrained vision-and-language models," in *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, eds. H. Bouamor, J. Pino, and K. Bali (Singapore: Association for Computational Linguistics), 8465–8483. doi: 10.18653/v1/2023.emnlp-main.525

Chen, H., Janizek, J. D., Lundberg, S., and Lee, S.-I. (2020). True to the model or true to the data? *arXiv [preprint]*. arXiv:2006.16234. doi: 10.48550/arXiv.2006.16234

Cimitan, A., Alves Pinto, A., and Geierhos, M. (2024). "Curation of benchmark templates for measuring gender bias in named entity recognition models," in *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, eds. N. Calzolari, M.-Y. Kan, V. Hoste, A. Lenci, S. Sakti, and N. Xue (Torino: ELRA and ICCL), 4238–4246.

Clark, B., Wilming, R., Dox, A., Eschenbach, P., Hached, S., Wodke, D. J., et al. (2025). EXACT: towards a platform for empirically benchmarking machine learning model explanation methods. *Meas. Sens.* 38:101794. doi: 10.1016/j.measen.2024.101794

Clark, B., Wilming, R., and Haufe, S. (2024). XAI-TRIS: non-linear image benchmarks to quantify false positive *post-hoc* attribution of feature importance. *Mach. Learn.* 113, 6871–6910. doi: 10.1007/s10994-024-06574-3

Costa-jussà, M. R., Li Lin, P., España-Bonet, C. (2020). "GeBioToolkit: automatic extraction of gender-balanced multilingual corpus of wikipedia biographies," in *Proceedings of the Twelfth Language Resources and Evaluation Conference*, eds. N. Calzolari, F. Béchet, P. Blache, K. Choukri, C. Cieri, T. Declerck, et al. (Marseille: European Language Resources Association), 4081–4088.

Dai, J., Upadhyay, S., Aivodji, U., Bach, S. H., and Lakkaraju, H. (2022). "Fairness via explanation quality: evaluating disparities in the quality of post hoc explanations," in *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society* (Oxford: ACM), 203–214. doi: 10.1145/3514094.3534159

Dehdarirad, T. (2025). Evaluating explainability in language classification models: a unified framework incorporating feature attribution methods and key factors affecting faithfulness. *Data Inf. Manag.* 9:100101. doi: 10.1016/j.dim.2025.100101

Dev, S., Li, T., Phillips, J. M., and Srikumar, V. (2020). On measuring and mitigating biased inferences of word embeddings. *Proc. AAAI Conf. Artif. Intell.* 34, 7659–7666. doi: 10.1609/aaai.v34i05.6267

Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2019). "Bert: pre-training of deep bidirectional transformers for language understanding," in *Proceedings of Naacl-HLT, Vol. 1* (Minneapolis, MN: ACL), 2.

DeYoung, J., Jain, S., Rajani, N. F., Lehman, E., Xiong, C., Socher, R., et al. (2020). "ERASER: a benchmark to evaluate rationalized NLP models," in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, eds. D. Jurafsky, J. Chai, N. Schluter, and J. Tetreault (Association for Computational Linguistics), 4443–4458. doi: 10.18653/v1/2020.acl-main.408

Dodge, J., Ilharco, G., Schwartz, R., Farhadi, A., Hajishirzi, H., Smith, N., et al. (2020). Fine-tuning pretrained language models: weight initializations, data orders, and early stopping. *arXiv [preprint]*. arXiv:2002.06305. doi: 10.48550/arXiv.2002.06305

Friedman, L., and Wall, M. (2005). Graphical views of suppression and multicollinearity in multiple linear regression. *Am. Stat.* 59, 127–136. doi: 10.1198/000313005X41337

Fryer, D. V., Strümke, I., and Nguyen, H. (2020). Explaining the data or explaining a model? Shapley values that uncover non-linear dependencies. *arXiv [preprint]*. arXiv:2007.06011. doi: 10.48550/arXiv.2007.06011

Gamboa, L. C. L., and Lee, M. (2024). "A novel interpretability metric for explaining bias in language models: applications on multilingual models from Southeast Asia," in *Proceedings of the 38th Pacific Asia Conference on Language, Information and Computation*, eds. N. Oco, S. N. Dita, A. M. Borlongan, and J.-B. Kim (Tokyo: Tokyo University of Foreign Studies), 296–305.

Gao, Z., Feng, A., Song, X., and Wu, X. (2019). Target-dependent sentiment classification with bert. *IEEE Access* 7, 154290–154299. doi: 10.1109/ACCESS.2019.2946594

Gonen, H., and Goldberg, Y. (2019). "Lipstick on a pig: debiasing methods cover up systematic gender biases in word embeddings but do not remove them," 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)*, eds. J. Burstein, C. Doran, and T. Solorio (Minneapolis, MN: Association for Computational Linguistics), 609–614.

Graells-Garrido, E., Lalmas, M., and Menczer, F. (2015). "First women, second sex: gender bias in Wikipedia," in *Proceedings of the 26th ACM Conference on Hypertext & Social Media* (New York, NY: ACM), 165–174. doi: 10.1145/2700171.2791036

Haufe, S., Meinecke, F., Görgen, K., Dähne, S., Haynes, J.-D., Blankertz, B., et al. (2014). On the interpretation of weight vectors of linear models in multivariate neuroimaging. *NeuroImage* 87, 96–110. doi: 10.1016/j.neuroimage.2013.10.067

Haufe, S., Wilming, R., Clark, B., Zhumagambetov, R., Panknin, D., Boubekki, A., et al. (2024). "Position: XAI needs formal notions of explanation correctness," in *Interpretable AI: Past, Present and Future*.

Hooker, S., Erhan, D., Kindermans, P.-J., and Kim, B. (2019). "A benchmark for interpretability methods in deep neural networks," in *Advances in Neural Information Processing Systems*, Vol. 32 (New York, NY: Curran Associates, Inc).

Ismail, A. A., Gunady, M., Corrada Bravo, H., and Feizi, S. (2020). Benchmarking deep learning interpretability in time series predictions. *Adv. Neural Inf. Process. Syst.* 33, 6441–6452.

Ismail, A. A., Gunady, M., Pessoa, L., Corrada Bravo, H., and Feizi, S. (2019). "Input-cell attention reduces vanishing saliency of recurrent neural networks," in *Advances in Neural Information Processing Systems*, Vol. 32 (New York, NY: Curran Associates, Inc).

Jacovi, A., and Goldberg, Y. (2020). "Towards faithfully interpretable NLP systems: how should we define and evaluate faithfulness?" in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics* (Association for Computational Linguistics), 4198–4205. doi: 10.18653/v1/2020.acl-main.386

Jentzsch, S., and Turan, C. (2022). "Gender bias in BERT - measuring and analysing biases through sentiment rating in a realistic downstream classification task," in *Proceedings of the 4th Workshop on Gender Bias in Natural Language Processing (GeBNLP)* (Seattle, WA: Association for Computational Linguistics), 184–199. doi: 10.18653/v1/2022.gebnlp-1.20

Jiménez-Luna, J., Grisoni, F., and Schneider, G. (2020). Drug discovery with explainable artificial intelligence. *Nat. Mach. Intell.* 2, 573–584. doi: 10.1038/s42256-020-00236-4

Joshi, R. K., Chatterjee, A., and Ekbal, A. (2024). Saliency guided debiasing: detecting and mitigating biases in LMs using feature attribution. *Neurocomputing* 563:126851. doi: 10.1016/j.neucom.2023.126851

Kaushik, D., Hovy, E., and Lipton, Z. C. (2019). Learning the difference that makes a difference with counterfactually-augmented data. *arXiv [preprint]*. arXiv:1909.12434. doi: 10.48550/arXiv.1909.12434

Kim, B., Wattenberg, M., Gilmer, J., Cai, C., Wexler, J., Viegas, F., et al. (2018). "Interpretability beyond feature attribution: quantitative testing with concept activation vectors (TCAV)," in *International Conference on Machine Learning (PMLR)*, 2668–2677.

Kokhlikyan, N., Miglani, V., Martin, M., Wang, E., Alsallakh, B., Reynolds, J., et al. (2020). Captum: a unified and generic model interpretability library for PyTorch. *arXiv [preprint]*. arxiv:2009.07896. doi: 10.48550/arxiv.2009.07896

Liu, Q., Kusner, M., and Blunsom, P. (2021). "Counterfactual data augmentation for neural machine translation," in *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, eds. K. Toutanova, A. Rumshisky, L. Zettlemoyer, D. Hakkani-Tur, I. Beltagy, S. Bethard, et al. (Association for Computational Linguistics), 187–197. doi: 10.18653/v1/2021.naacl-main.18

Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., et al. (2019). RoBERTa: a robustly optimized BERT pretraining approach. *arXiv [preprint]*. arXiv:1907.11692. doi: 10.48550/arXiv.1907.11692

Lundberg, S. M., and Lee, S.-I. (2017). "A unified approach to interpreting model predictions," in *Advances in Neural Information Processing Systems 30*, eds. I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (New York, NY: Curran Associates, Inc), 4765–4774.

Lundberg, S. M., Nair, B., Vavilala, M. S., Horibe, M., Eisses, M. J., Adams, T., et al. (2018). Explainable machine-learning predictions for the prevention of hypoxaemia during surgery. *Nat. Biomed. Eng.* 2, 749–760. doi: 10.1038/s41551-018-0304-0

Manzini, T., Yao Chong, L., Black, A. W., and Tsvetkov, Y. (2019). "Black is to criminal as caucasian is to police: detecting and removing multiclass bias in word embeddings," 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)*, pages eds. J. Burstein, C. Doran, and T. Solorio (Minneapolis, MN: Association for Computational Linguistics), 615–621. doi: 10.18653/v1/N19-1062

Mhasawade, V., Rahman, S., Haskell-Craig, Z., and Chunara, R. (2024). "Understanding disparities in post hoc machine learning explanation," in *The 2024 ACM Conference on Fairness, Accountability, and Transparency* (Rio de Janeiro Brazil: ACM), 2374–2388. doi: 10.1145/3630106.3659043

Mitchell, T. M. (2007). "The need for biases in learning generalizations," in *Rutgers CS tech report, CBM-TR-117*.

Montañez, G. D., Hayase, J., Lauw, J., Macias, D., Trikha, A., and Vendemiatti, J. (2019). "The futility of bias-free learning and search," in *Australasian Joint Conference on Artificial Intelligence* (Cham: Springer), 277–288. doi: 10.1007/978-3-030-35288-2\_23

Murdoch, W. J., Singh, C., Kumbier, K., Abbasi-Asl, R., and Yu, B. (2019). Definitions, methods, and applications in interpretable machine learning. *Proc. Nat. Acad. Sci.*, 116, 22071–22080. doi: 10.1073/pnas.1900654116

Nadeem, M., Bethke, A., and Reddy, S. (2021). "StereoSet: measuring stereotypical bias in pretrained language 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)*, eds. C. Zong, F. Xia, W. Li, and R. Navigli (Association for Computational Linguistics), 5356–5371. doi: 10.18653/v1/2021.acl-long.416

Nangia, N., Vania, C., Bhalerao, R., and Bowman, S. R. (2020). "CrowS-pairs: a challenge dataset for measuring social biases in masked language models," in *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, eds. B. Webber, T. Cohn, Y. He, and Y. Liu (Association for Computational Linguistics), 1953–1967. doi: 10.18653/v1/2020.emnlp-main.154

Navigli, R., Conia, S., and Ross, B. (2023). Biases in large language models: origins, inventory, and discussion. *J. Data Inf. Qual.* 15, 1–21. doi: 10.1145/3597307

Oliveira, M., Wilming, R., Clark, B., Budding, C., Eitel, F., Ritter, K., et al. (2024). Benchmarking the influence of pre-training on explanation performance in MR image classification. *Front. Artif. Intell.* 7:1330919. doi: 10.3389/frai.2024.1330919

OpenAI (2023). GPT-4 technical report. *arXiv [preprint]*. arXiv:2303.08774. doi: 10.48550/arXiv.2303.08774

Parrish, A., Chen, A., Nangia, N., Padmakumar, V., Phang, J., Thompson, J., et al. (2022). "BBQ: a hand-built bias benchmark for question answering," in *Findings of the Association for Computational Linguistics: ACL 2022*, eds. S. Muresan, P. Nakov, and A. Villavicencio (Association for Computational Linguistics), 2086–2105. doi: 10.18653/v1/2022.findings-acl.165

Prost, F., Thain, N., and Bolukbasi, T. (2019). "Debiasing embeddings for reduced gender bias in text classification," in *Proceedings of the First Workshop on Gender Bias in Natural Language Processing*, eds. M. R. Costa-jussà, C. Hardmeier, W. Radford, and K. Webster (Florence: Association for Computational Linguistics), 69–75. doi: 10.18653/v1/W19-3810

Radford, A., Narasimhan, K., Salimans, T., and Sutskever, I. (2018). *Improving Language Understanding by Generative Pre-training*. OpenAI.

Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., et al. (2019). Language models are unsupervised multitask learners. *OpenAI Blog* 1:9.

Reagle, J., and Rhue, L. (2011). Gender bias in wikipedia and britannica. *Int. J. Commun.* 5:21.

Ribeiro, M. T., Singh, S., and 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* (New York, NY: ACM), 1135–1144. doi: 10.18653/v1/N16-3020

Rong, Y., Leemann, T., Borisov, V., Kasneci, G., and Kasneci, E. (2022). "A consistent and efficient evaluation strategy for attribution methods," in *Proceedings of the 39th International Conference on Machine Learning*, Vol. 162, eds. K. Chaudhuri, S. Jegelka, L. Song, C. Szepesvari, G. Niu, and S. Sabato (PMLR), 18770–18795.

Rychener, Y., Renard, X., Seddah, D., Frossard, P., and Detyniecki, M. (2020). QUACKIE: a NLP classification task with ground truth explanations. *arXiv [preprint]*. arXiv:2012.13190. doi: 10.48550/arXiv.2012.13190

Samek, W., Montavon, G., Vedaldi, A., Hansen, L. K., and Müller, K.-R. (Eds.) (2019). *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning*, volume 11700 of *Lecture Notes in Computer Science*. Cham: Springer International Publishing. doi: 10.1007/978-3-030-28954-6

Shrikumar, A., Greenside, P., and Kundaje, A. (2017). "Learning important features through propagating activation differences," in *Proceedings of the 34th International Conference on Machine Learning - Volume 70, ICML'17 (JMLR)*, 3145–3153.

Simonyan, K., Vedaldi, A., and Zisserman, A. (2013). Deep inside convolutional networks: visualising image classification models and saliency maps. *arXiv [preprint]*. arXiv:1312.6034. doi: 10.48550/arXiv.1312.6034

Sparck Jones, K. (1972). A statistical interpretation of term specificity and its application in retrieval. *J. Doc.* 28, 11–21. doi: 10.1108/eb026526

Springenberg, J., Dosovitskiy, A., Brox, T., and Riedmiller, M. (2015). "Striving for simplicity: the all convolutional net," in *ICLR (Workshop Track)*.

Sullivan, E. (2024). "SIDES: separating idealization from deceptive "explanations" in xAI," in *The 2024 ACM Conference on Fairness, Accountability, and Transparency* (Rio de Janeiro: ACM), 1714–1724. doi: 10.1145/3630106.3658999

Sundararajan, M., Taly, A., and Yan, Q. (2017). "Axiomatic attribution for deep networks," in *Proceedings of the 34th International Conference on Machine Learning*, Volume 70 of *Proceedings of Machine Learning Research*, eds. D. Precup, and Y. W. Teh (PMLR), 3319–3328.

Tjoa, E., and Guan, C. (2023). Quantifying explainability of saliency methods in deep neural networks with a synthetic dataset. *IEEE Trans. Artif. Intell.* 4, 858–870. doi: 10.1109/TAI.2022.3228834

Tran, K. A., Kondrashova, O., Bradley, A., Williams, E. D., Pearson, J. V., Waddell, N., et al. (2021). Deep learning in cancer diagnosis, prognosis and treatment selection. *Genome Med.* 13:152. doi: 10.1186/s13073-021-00968-x

Wilming, R., Budding, C., Müller, K.-R., and Haufe, S. (2022). Scrutinizing xai using linear ground-truth data with suppressor variables. *Mach. Learn.* 111, 1903–1923. doi: 10.1007/s10994-022-06167-y

Wilming, R., Kieslich, L., Clark, B., and Haufe, S. (2023). "Theoretical behavior of XAI methods in the presence of suppressor variables," in *Proceedings of the 40th International Conference on Machine Learning, Volume 202 of Proceedings of Machine Learning Research*, eds. A. Krause, E. Brunskill, K. Cho, B. Engelhardt, S. Sabato, and J. Scarlett (PMLR), 37091–37107.

Yang, Z., Dai, Z., Yang, Y., Carbonell, J., Salakhutdinov, R. R., Le, Q. V., et al. (2019). "XLNet: generalized autoregressive pretraining for language understanding," in *Advances in Neural Information Processing Systems, Vol. 32* (New York, NY: Curran Associates, Inc).

Yu, J., and Jiang, J. (2019). "Adapting bert for target-oriented multimodal sentiment classification," in *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence*, 5408–5414. doi: 10.24963/ijcai.2019/751

Zakizadeh, M., Miandoab, K., and Pilehvar, M. (2023). "DiFair: a benchmark for disentangled assessment of gender knowledge and bias," in *Findings of the Association for Computational Linguistics: EMNLP 2023*, eds. H. Bouamor, J. Pino, and K. Bali (Singapore: Association for Computational Linguistics), 1897–1914. doi: 10.18653/v1/2023.findings-emnlp.127

Zhang, Y., Weng, Y., and Lund, J. (2022). Applications of explainable artificial intelligence in diagnosis and surgery. *Diagnostics* 12:237. doi: 10.3390/diagnostics12020237

Zhao, J., Wang, T., Yatskar, M., Ordonez, V., and Chang, K.-W. (2017). "Men also like shopping: reducing gender bias amplification using corpus-level constraints," in *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, eds. M. Palmer, R. Hwa, and S. Riedel (Copenhagen: Association for Computational Linguistics), 2979–2989. doi: 10.18653/v1/D17-1323

Zheng, S., and Yang, M. (2019). "A new method of improving bert for text classification," in *Intelligence Science and Big Data Engineering. Big Data and Machine Learning: 9th International Conference, IScIDE 2019, Nanjing, China, October 17–20, 2019, Proceedings, Part II* 9 (Cham: Springer), 442–452. doi: 10.1007/978-3-030-36204-1\_37