



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Explanations for automatic speech recognition

Citation for published version:

Wu, X, Bell, P & Rajan, A 2023, Explanations for automatic speech recognition. in *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. International Conference on Acoustics, Speech, and Signal Processing (ICASSP), IEEE, 2023 IEEE International Conference on Acoustics, Speech and Signal Processing, Rhodes Island, Greece, 4/06/23.
<https://doi.org/10.1109/ICASSP49357.2023.10094635>

Digital Object Identifier (DOI):

[10.1109/ICASSP49357.2023.10094635](https://doi.org/10.1109/ICASSP49357.2023.10094635)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



EXPLANATIONS FOR AUTOMATIC SPEECH RECOGNITION

Xiaoliang Wu, Peter Bell, Ajitha Rajan

School of Informatics, University of Edinburgh

ABSTRACT

We address quality assessment for neural network based ASR by providing explanations that help increase our understanding of the system and ultimately help build trust in the system. Compared to simple classification labels, explaining transcriptions is more challenging as judging their correctness is not straightforward and transcriptions as a variable-length sequence is not handled by existing interpretable machine learning models.

We provide an explanation for an ASR transcription as a subset of audio frames that is both a minimal and sufficient cause of the transcription. To do this, we adapt existing explainable AI (XAI) techniques from image classification - (1) Statistical Fault Localisation (SFL) [1] and (2) Causal [2]. Additionally, we use an adapted version of Local Interpretable Model-Agnostic Explanations (LIME) [3] for ASR as a baseline in our experiments. We evaluate the quality of the explanations generated by the proposed techniques over three different ASR – Google API [4], the baseline model of Sphinx [5], DeepSpeech [6] – and 100 audio samples from the Commonvoice dataset [7].

Index Terms— Explanation, Automatic Speech Recognition

1. INTRODUCTION

Owing to the prevalence of ASR in our daily lives, concerns such as their quality and accountability have become particularly important. The complex and often blackbox [8] (inaccessible internal state) nature of neural network-based ASR makes it hard to ensure their quality. We focus on addressing quality assessment of ASRs by providing explanations for a given transcription that can help increase our understanding of the ASR and have the potential to help fix the failure causing faults and facilitate the accountability process. This is supported by Madsen et al. [9] who state, “*Quality assessment should be done through model explanations.*”

XAI techniques¹ have emerged rapidly in the past five years. One of the first XAI techniques was developed for explaining image classification labels. It was a local perturbation-based technique called LIME [3] that randomly samples inputs around the instance being explained using perturbations of the original input instance. By feeding these samples into the model, LIME creates a new, local and linear interpretable model around the original input that can help

explain the input/output instance. Other popular explanation techniques include some variants of LIME [10, 11, 12]; SHAP [13], a game-theory based input feature ranking based on its contribution to the output prediction; DeepLIFT [14] and Integrated Gradients [15], gradient-based methods using a reference input (chosen by the user) and difference in original and reference activations backpropagated to the input layer for contribution scores of input features. Minimal Contrastive Editing (MICE) [37] is a perturbation and example-based explanation method especially developed for NLP models. MICE describes contrastive explanations as answers to questions in the form “Why p and not q ?”, why a prediction event p happened instead of contrast event q . The majority of these XAI techniques attempt to explain prediction labels for image or natural language classification models.

There are a few recent techniques that interpret speech-input models from a white-box perspective. Specifically, [16] introduced time-independent Neuron Activation Profiles (NAPs) for each activation neuron of a given ASR for certain groups of inputs. They discovered different layers will be activated by different types of inputs by clustering NAPs. Later, [17] found that averaging the feature maps after the ReLU activation in each transposed convolutional layer produces interpretable time series data that summarizes encodings of features of speech at the corresponding layer. This also allowed them to successfully understand what properties of speech will be encoded in each layer. These techniques have two limitations: (1) They are proposed for a specific white-box model and neural network architecture, making it difficult to generalize to other models; and (2) They do not explain the transcription of ASR, making it difficult to assess its quality with these techniques. Compared to simple classification labels, transcriptions – as a variable-length output sequence – are more challenging for two primary reasons: (1) existing interpretable machine learning models, like linear regression model that LIME is based on, cannot handle variable length output sequences, s ; and (2) Unlike classification labels that involve simple comparisons for determining their correctness, judging transcription correctness involves semantic comparisons with room for small deviations.

To explain ASR transcriptions, we first come up with a means to label transcriptions as Correct or Incorrect based on the degree of similarity to the desired transcription. We then aim to provide an explanation for an ASR transcription from a given input audio as a subset of audio frames. We refer to *frames* in explanations as raw data bins in time dimension, and they are not to be confused with frames represented by features like filterbanks or MFCCs in the frequency

¹Our focus is restricted to *post-hoc* explainable methods that provide their explanation *after* a model is trained.

or cepstral domains. To provide explanations, we adapt existing XAI techniques from image classification - (1) Statistical Fault Localisation (SFL) [1], (2) Causal [2]. We choose to adapt SFL and Causal since they are causality-based XAI techniques that produce minimal and sufficient explanations outperforming other state-of-the-art techniques in the image classification domain. Additionally, we use LIME, a popular XAI technique in the literature, as a baseline in our comparison. It is worth noting that we choose to adapt XAI techniques from image classification rather than NLP classification tasks as segmentation of input text often produces words with specific meanings. However, in image- and speech-based tasks, input segmentation results in pixels or frames that have no meaning independently and only make sense when they are considered in logical groups. This similarity motivates us to choose image-based explanation techniques over NLP.

We evaluate quality of explanations from different techniques in terms of their size and consistency across different ASR. We found SFL and Causal explanations did well on all three ASR systems with respect to size and consistency. Source code for X-ASR and examples from our experiment are available at <https://anonymous.4open.science/r/Xasr-6E11>.

2. METHODOLOGY

Our framework, X-ASR, generates explanations for ASR transcriptions in two steps, (1) *Classify* ASR transcriptions with respect to reference and (2) *Adapt* existing image explanation techniques – SFL, Causal and LIME – to work over audio input. We need the classification means in Step 1 because perturbation-based explanations, considered in Step 2, rely on classifying outputs from input perturbations and use the changes in outputs to assign importance to perturbed input regions. The classification step relies on existing similarity metrics in the literature. Adapting image-based explanations for ASR in step 2 has not been explored previously.

2.1. Classifying ASR Transcriptions

Unlike classification tasks where it is straightforward to judge whether the output label is Correct (matches with expected label) or Incorrect, transcriptions from ASR are harder to assign a binary label as it may differ from expected transcription but still be Correct. For example, the expected transcription for an audio input may be “I’d like an apple”, while the actual ASR output is “I like apple”. We might still consider the ASR transcription as acceptable. Assessing the correctness of an ASR transcription is subject to human judgement and allowances for differences from expected transcription.

To address this, we attach Correct or Incorrect labels to transcriptions based on the extent of similarity to expected transcriptions, using a user-defined threshold, T . If the similarity to original transcription is higher than the threshold T , then the perturbed audio transcription is marked as Correct, and Incorrect otherwise. We support *two similarity metrics* in X-ASR: (1) WER, the widely used Word Error Rate metric that scores similarity between two sentences based on number of insertions, deletions and substitutions; and (2) Bert, where

Algorithm 1 Explanation using SFL for ASR

Input: F : ASR model; x : original audio; M : SFL technique measurement
 M ; $metric$: similarity metric including Bert and WER
Output: a subset of frames P_{exp} ;
1: $T(x)$ = mutants from x
2: **for** each frame f_i in P **do**
3: calculate the $\langle a_{ep}^i, a_{ef}^i, a_{np}^i, a_{np}^i \rangle$ from $T(x)$
4: $value = M(\langle a_{ep}^i, a_{ef}^i, a_{np}^i, a_{np}^i \rangle)$
5: $ranking$ = frames in P from high $value$ to low
6: $P_{exp} = \phi$
7: **for** each frame f_i in $ranking$ **do**
8: $P_{exp} = P_{exp} \cup \{f_i\}$
9: x_{exp} = mask frames of x that are not in P_{exp}
10: $Flag = Similarity(F(x), F(x_{exp}), metric)^2$
11: **if** $Flag = Correct$ **then**
12: **return** P_{exp}
13: **end if**
14: **end for**
15: **end for**

we compute the cosine similarity between the semantic vectors for the original and perturbed audio transcriptions using Sentence Bert [18].

2.2. Adapting Image Explanations for ASR

We discuss three image-based explanation techniques that we adapt for ASR. All three techniques rely on input perturbations. We classify outputs from perturbed inputs by measuring similarity to the original transcription.

Adapting SFL Explanations: SFL [1] adapted statistical fault localization from software testing literature, to rank importance of pixels for a given classification label and built minimal explanations from this ranking. Statistical fault localization is traditionally used to rank program elements such as statements based on their likelihood of being failure causing faults.

We describe step-wise how this idea is used in image-based SFL explanations to understand how we use it for speech.

Step 1: Given an input image x classified as y by a DNN, a series of mutants is generated by randomly setting a certain proportion of pixels to the background color (masked pixels) and classify each of these mutants. If the mutant classification label is y (matching original). Then the mutant is a passing test, while a mutant whose classification is not y is considered a failing test. To keep the number of passing and failing tests balanced, once a failing test is discovered, the number of pixels to be masked for the next mutant is reduced to make it more likely to be a passing test. Conversely, increase the number of pixels to be masked if a passing test is discovered.
Step 2: Within each mutant, some pixels are consistent with the original input while others are masked. For each pixel in the input sample, count the number of mutants in which the pixels are masked versus unmasked and the output is changed versus unchanged. Based on the count, compute the importance score for every pixel using the chosen fault localisation measure, like Tarantula [19].

Step 3: Starting from the top ranked pixel, add pixels in

²The input to the *Similarity* function is the original transcription, the transcription of the mutant and the similarity metric.

Algorithm 2 Responsibility

Input: P_i : a partition; F : PR model; x : an original audio; *metric*: similarity metric including Bert and WER

Output: a responsibility map of P_i ;

1: $\tilde{X}_i =$ the set of mutants obtained from x by masking subsets of P_i

2: **for** each superframe $P_{i,j}$ in P_i **do**

3: $\tilde{X}_i^j = \{x_m : P_{i,j} \text{ is not masked in } x_m \text{ and } x_m \text{ is Correct}\}$

4: $k = \min \{diff(x_m, x) | x_m \in \tilde{X}_i^j\}$

5: $r_{i,j} = \frac{1}{k+1}$

6: **end for**

7: **return** $r_{i,0}, \dots, r_{i,|P_i|-1}$

the rank list progressively to the explanation set, until classification of the explanation set matches the original image classification, y . This is returned as the explanation for an input image x being classified as y .

When applied to the audio signal, we use frames instead of pixels in all steps. We generate mutant audios by randomly selecting frames and setting the data points in them to zero. We use similarity metrics BERT or WER for mutant classification in Step 1 and explanation classification in Step 3. We show SFL explanations adapted for ASR in Algorithm 1. In Lines 4 and 5 in Algorithm 1, for each frame f_i , we construct a parameter vector $\langle a_{ep}^i, a_{ef}^i, a_{np}^i, a_{nf}^i \rangle$ where a_{ep}^i is the number of mutants in the mutants set labeled Correct in which f_i is not masked; a_{ef}^i is the number of mutants in the mutant set labeled Incorrect in which f_i is not masked; a_{np}^i is the number of mutants in the mutant set labeled Correct in which f_i is masked; a_{nf}^i is the number of mutants in the mutant set labeled Incorrect in which f_i is masked. After that, we use the ranking measurement in [1] to compute the importance and rank frames within the given audio. Finally, as shown in Line 6 to Line 14 in Algorithm 1, we create the explanation for a given audio following Step 3 of the method.

Adapting Causal for ASR: In 2021, a causal theory-based approach, *Causal*, was proposed for computing explanations of image classifier [2]. Explanations from *Causal* were robust in the presence of image occlusions and outperformed SOTA. Causation is a relation between two events A (the cause) and B (the effect) when A causes B. Counterfactual theories [20] to define causation in terms of a counterfactual relation state that “An event A causally depends on B if, and only if: 1. If B had occurred, then A would have occurred. 2. If B had not occurred, then A would not have occurred.”

Causal uses a framework of causality proposed by [21] which extends counterfactual reasoning by considering contingencies which are defined as changes in the current setting. *Causal* quantifies causality and expresses the degree of responsibility for any actual cause defined as the minimal change required to create a counterfactual dependence. The responsibility is defined as $1/(1+k)$ where k is the size of the smallest contingency. The responsibility measure for an actual cause can take any value between 0 and 1, where 0 means no causal dependency and 1 means very strong causal dependency. The algorithm for image-based explanations is briefly outlined below: First, partition the original image, without any overlap, into blocks referred to as superpixels. Since there

are several ways to partition an image into blocks, many such partitioning methods are considered and the resulting superpixels. A superpixel S is defined as the *cause* of an image x being labelled as y when there is a subset P of superpixels of x that satisfies the following three conditions, (1) the superpixel S does not belong to P , (2) masking any subset of P (setting pixels within to background color) will not affect the label y , and (3) masking both P and superpixel S at the same time will change the classification. P is referred to as the witness of superpixel S and the size of the smallest witness (contingency to create a counterfactual dependence) is used to assign a responsibility for each superpixel. When the size of the smallest witness is smaller, the actual cause S requires less changes in other superpixels to help it achieve a counterfactual dependence. That is, there is a stronger causal relationship between it and the image being correctly classified. The superpixels are then iteratively refined into smaller partitions, repeating the computation of responsibility. This process is repeated with several partitioning methods to come up with a ranked list of pixel responsibilities which is then used to construct explanations.

When applied to the audio signal, we partition the original audio, without any overlap, into superframes. We then compute responsibility of superframes, as shown in Algorithm 2, and iteratively partition them further, repeating responsibility computation. As shown in Lines 3, 4 and 5 in Algorithm 2, responsibility of a superframe, $P_{i,j}$ ³, is $r(i, j) = 1/(k + 1)$, where k is equal to the minimum difference between a mutant audio and the original audio over mutants x_m that do not mask $P_{i,j}$ and are labeled Correct. Masking a subset of superframes for computing witness in condition (2) involves setting a random selection of frames within to zero.

Adapting LIME Explanations: LIME, proposed in [3], is an XAI technique that can be applied to any model without needing any information about its structure. LIME provides a local explanation by replacing a complex neural network (NN) locally with something simpler, for example a linear regression model. LIME creates many perturbations of the original image by masking out random segments, and then weights these perturbations by their ‘closeness’ to the original image to ensure that drastic perturbations have little impact. It then uses the simpler model (for example, the linear model) to learn the mapping between the perturbations and any change in output label. This process allows LIME to determine which segments are most important to the classification decision.

When applied to the audio signal, we replace image segments with audio frames. It is worth noting that LIME is designed to output the importance ranking of pixels, or frames in our case, and this ranking is considered a LIME explanation. Inspired from minimal and sufficient explanations in SFL and *Causal*, we construct smaller LIME explanations using a greedy approach that starts from the top ranked frame, and then adds frames in the rank list iteratively to the explanation until classification of the explanation is correct with respect to the original audio transcription. We use these potentially smaller LIME explanations in our experiments in Section 4

³ $P_{i,j}$ means this is the j^{th} superframe in the i^{th} partition.

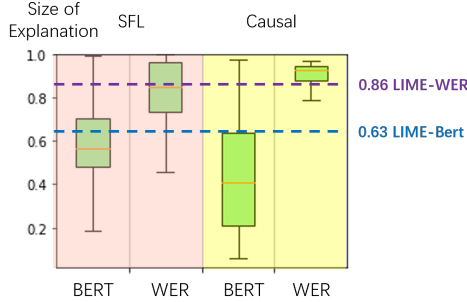


Fig. 1. Size of explanations using SFL and Causal against LIME using each of 2 similarities on Google ASR.

3. EXPERIMENTS

We evaluate the explanations generated for ASR models using three quality metrics that are described below. We use three different ASR – Google API [4] (referenced as Google in the results), baseline model of Sphinx [5] and DeepSpeech [6]) 0.7.3 version – and 100 audio samples from Commonvoice dataset [7]. Within X-ASR, we evaluate three explanation techniques mentioned in Section 2, namely, SFL, Causal, LIME, with two similarity metrics (Bert and WER) used to classify transcriptions from perturbed inputs. We use the default setting for every ASR. We use Google Colab Pro with two NVIDIA Tesla T4 GPUs (16GB RAM, 2560 cores) to run our experiments. We use the following parameters in our experiments: an audio sampling rate of 16000Hz, frame length of 512. For SFL, Causal and LIME, the mutation factor is 0.05 which determines the proportion of frames to be randomly selected for each mutation. The size of the mutants set is 100. For Causal, we run the experiments 3 times to get a reliable ranking of frames. For similarity metrics, classification threshold for Bert is 0.5 and for WER is 0. We choose a zero threshold for WER to emulate strict classification. Additionally, we investigate the effect of choosing other classification threshold values for the similarity metrics and report our findings in Section 4.

Quality Metrics for Explanations: We use two previously used quality metrics from image classification explanations [22, 1] to compare and contrast explanations over three ASR, namely, *size* and *consistency*. We use size of the explanation as number of frames in the explanation versus the input. When size is smaller, the quality of the explanation is better as it indicates the technique is more selective in identify key frames important to the output. The consistency metric assesses the degree to which similar explanations are generated from different ASR for a given audio. We assess the consistency of explanation methods using Google API as the reference and calculate the fraction of frames that stay the same in explanations generated with the other systems, namely Sphinx and DeepSpeech.

4. RESULTS AND ANALYSIS

All three techniques in our method, SFL, Causal, LIME, were able to generate explanations for every input audio in our dataset and across all three ASR. Figure 1 shows size of

		SFL	Causal	LIME
Consistency (Google-Sphinx)	Bert	0.81	0.61	0.80
	WER	0.97	0.98	0.92
Consistency (Google-DeepSpeech)	Bert	0.67	0.60	0.67
	WER	0.96	0.98	0.89

Table 1. Consistency(with respect to Sphinx or DeepSpeech) of explanations generated by three explanation methods across two similarity metrics using Google ASR.

explanations generated by the different techniques with the two similarity metrics on Google ASR. Results on other two ASR had a similar trend and are not shown here due to space limitations. We find SFL and Causal outperform the LIME counterpart for Bert by generating significantly smaller explanations (statistically significant difference confirmed with one-way Anova followed by post-hoc Tukey’s test [23]). This is because both SFL and Causal not only pay attention to the changes in the output after masking a superframe but also take the number of other superframes affected by masking. Causal generates smaller explanations than SFL (statistically significant) owing to stricter conditions from causal-theory in assigning responsibility of a superframe and the iterative refinement of superframes.

Finally, there is a difference in size based on similarity metric. WER generates larger explanations than Bert. This is because WER has no tolerance for difference in transcription. As a result, more perturbations satisfy the condition for changing output classification. Bert accepts some changes in transcriptions with 0.5 threshold. Raising the threshold for WER would also result in smaller explanations. Experiments with other thresholds in both metrics followed a similar trend.

Table 1 shows consistency of explanations with Google ASR versus Sphinx (rows highlighted in pink) and Google versus DeepSpeech (rows highlighted in yellow). We find SFL is most consistent with Bert similarity on both ASR pairs. Causal is most consistent with WER similarity measure owing to its large explanation size (92% of input frames on average). This also follows for LIME explanations that are larger in size. Clearly, size and consistency are conflicting metrics and we believe SFL with Bert achieves a good compromise between the two – 56% size and 74% consistency, on average.

5. CONCLUSION

We propose the first framework, X-ASR, for generating explanations for ASR that supports three techniques: SFL, Causal and LIME. We found SFL and Causal explanations perform comparably while outperforming LIME on size and consistency. The study presented only investigates perturbation-based explanation techniques that are portable and do not consider ASR structure. We plan to investigate white-box ASR explanation techniques and their effectiveness in the future.

6. REFERENCES

- [1] Youcheng Sun, Hana Chockler, Xiaowei Huang, and Daniel Kroening, “Explaining deep neural networks using spectrum-based fault localization,” *CoRR*, vol. abs/1908.02374, 2019.
- [2] Hana Chockler, Daniel Kroening, and Youcheng Sun, “Compositional explanations for image classifiers,” *CoRR*, vol. abs/2103.03622, 2021.
- [3] Marco Túlio Ribeiro, Sameer Singh, and Carlos Guestrin, ““why should I trust you?”: Explaining the predictions of any classifier,” *CoRR*, vol. abs/1602.04938, 2016.
- [4] Google, ““google speech to text api”,” 2017.
- [5] Paul Lamere, Philip Kwok, William Walker, Evandro B Gouvêa, Rita Singh, Bhiksha Raj, and Peter Wolf, “Design of the cmu sphinx-4 decoder,” in *Interspeech*. Cite-seer, 2003.
- [6] Awni Hannun, Carl Case, Jared Casper, Bryan Catanzaro, Greg Diamos, Erich Elsen, Ryan Prenger, Sanjeev Satheesh, Shubho Sengupta, Adam Coates, and Andrew Y. Ng, “Deep speech: Scaling up end-to-end speech recognition,” 2014.
- [7] R. Ardila, M. Branson, K. Davis, M. Henretty, M. Kohler, J. Meyer, R. Morais, L. Saunders, F. M. Tyers, and G. Weber, “Common voice: A massively-multilingual speech corpus,” in *Proceedings of the 12th Conference on Language Resources and Evaluation (LREC 2020)*, 2020, pp. 4211–4215.
- [8] Alexandros Kastanos, Anton Ragni, and Mark JF Gales, “Confidence estimation for black box automatic speech recognition systems using lattice recurrent neural networks,” in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 6329–6333.
- [9] Andreas Madsen, Siva Reddy, and Sarath Chandar, “Post-hoc interpretability for neural nlp: A survey,” *arXiv preprint arXiv:2108.04840*, 2021.
- [10] Muhammad Rehman Zafar and Naimul Mefraz Khan, “DLIME: A deterministic local interpretable model-agnostic explanations approach for computer-aided diagnosis systems,” *CoRR*, vol. abs/1906.10263, 2019.
- [11] Kacper Sokol and Peter A. Flach, “Limetree: Interactively customisable explanations based on local surrogate multi-output regression trees,” *CoRR*, vol. abs/2005.01427, 2020.
- [12] Navdeep Gill, Megan Kurka, and Wen Phan, “Machine learning interpretability with h2o driverless ai,” 2019.
- [13] Scott M. Lundberg and Su-In Lee, “A unified approach to interpreting model predictions,” in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, Red Hook, NY, USA, 2017, NIPS’17, p. 4768–4777, Curran Associates Inc.
- [14] Avanti Shrikumar, Peyton Greenside, and Anshul Kundaje, “Learning important features through propagating activation differences,” *CoRR*, vol. abs/1704.02685, 2017.
- [15] Mukund Sundararajan, Ankur Taly, and Qiqi Yan, “Axiomatic attribution for deep networks,” *CoRR*, vol. abs/1703.01365, 2017.
- [16] Andreas Krug, René Knaebel, and Sebastian Stober, “Neuron activation profiles for interpreting convolutional speech recognition models,” 2018.
- [17] Gasper Begus and Alan Zhou, “Interpreting intermediate convolutional layers of cnns trained on raw speech,” *CoRR*, vol. abs/2104.09489, 2021.
- [18] Nils Reimers and Iryna Gurevych, “Sentence-bert: Sentence embeddings using siamese bert-networks,” *CoRR*, vol. abs/1908.10084, 2019.
- [19] James A Jones, Mary Jean Harrold, and John Stasko, “Visualization of test information to assist fault localization,” in *Proceedings of the 24th International Conference on Software Engineering. ICSE 2002*. IEEE, 2002, pp. 467–477.
- [20] David Hume, *A treatise of human nature*, Penguin Classics, 1985.
- [21] Joseph Y. Halpern and Judea Pearl, “Causes and explanations: A structural-model approach — part 1: Causes,” *CoRR*, vol. abs/1301.2275, 2013.
- [22] Marko Robnik-Šikonja and Marko Bohanec, *Perturbation-Based Explanations of Prediction Models*, pp. 159–175, Springer International Publishing, Cham, 2018.
- [23] John W Tukey et al., *Exploratory data analysis*, vol. 2, Reading, MA, 1977.