

# XPROAX-Local explanations for text classification with progressive neighborhood approximation

Yi Cai  
L3S Research Center  
Leibniz Universität Hannover  
Hannover, Germany  
cai@l3s.de

Arthur Zimek  
Dept. of Math. and Comp. Science  
University of Southern Denmark  
Odense, Denmark  
zimek@imada.sdu.dk

Eirini Ntoutsi  
Dept. of Math. and Comp. Science  
Freie Universität Berlin  
Berlin, Germany  
eirini.ntoutsi@fu-berlin.de

**Abstract**—The importance of the neighborhood for training a local surrogate model to approximate the local decision boundary of a black box classifier has been already highlighted in the literature. Several attempts have been made to construct a better neighborhood for high dimensional data, like texts, by using generative autoencoders. However, existing approaches mainly generate neighbors by selecting purely at random from the latent space and struggle under the curse of dimensionality to learn a good local decision boundary. To overcome this problem, we propose a progressive approximation of the neighborhood using counterfactual instances as initial landmarks and a careful 2-stage sampling approach to refine counterfactuals and generate factuals in the neighborhood of the input instance to be explained. Our work focuses on textual data and our explanations consist of both word-level explanations from the original instance (intrinsic) and the neighborhood (extrinsic) and factual- and counterfactual-instances discovered during the neighborhood generation process that further reveal the effect of altering certain parts in the input text. Our experiments on real-world datasets demonstrate that our method outperforms the competitors in terms of usefulness and stability (for the qualitative part) and completeness, compactness and correctness (for the quantitative part).

**Index Terms**—Explainable AI, Local explanations, Counterfactuals, Neighborhood approximation, Text classification

## I. INTRODUCTION

With wide applications of machine learning models in real-world scenarios, concerns on transparency of decision making process keep growing [11]. Explainable artificial intelligence hence becomes a popular aspect that helps human users to understand behaviors of a black box, which hides its internal logic and inner working from users. An enormous number of efforts have been put into explaining decision systems working with tabular [9], [24] and image data [6], [23], [28]. Unfortunately, most of these methods cannot be directly applied to text classifiers because of the nature of textual data, thus this field is left underdeveloped.

The main focus of this paper is a local model-agnostic explanation method for text classifiers. Explanation methods in this category firstly train a surrogate model with generated neighbors to mimic the local behavior of the black box waiting to be explained, and then extract knowledge from such a transparent local predictor as the explanation for an inquired black box decision over an instance. Under this circumstance, the quality of generated neighbors dominates the upper-bound

of the fidelity that a surrogate model can achieve. To generate neighborhood points, perturbation of inputs [33] is a common choice. It modifies the values of numeric features within a small range or switches the values of categorical features into some others. Again, this becomes a much more challenging task while dealing with textual data due to the lack of a formal definition of the neighborhood of a text. Some existing works [25], [26] generate neighboring texts by randomly dropping words from the given text. XSPELLS [16] proposes to generate neighboring texts in a latent space learned through a generative autoencoder. This method can generate more realistic (semantically meaningful and grammatically correct) texts than the former choice. However, both methods use a random sampling strategy no matter in which space the perturbation is applied, which might not be the optimal option.

In this work, we propose XPROAX<sup>1</sup> (local eXplainer with PROgressive neighborhood ApproXimation), a novel explanation method with progressive neighborhood approximation for text classifiers. Main contributions of this work are summarized as follows:

- We propose a neighborhood generation method in a latent space for explaining text classifiers with two-staged progressive approximation of neighborhood.
- The output of the explanation method is composed of word-level and factual-level explanations for comprehensive understandings.
- We design an automatic evaluation method that quantitatively evaluates explanations for text classifiers regarding completeness, compactness, and correctness.

The rest of the paper is organized as follows. In Section II, we discuss state-of-the-art works in the field of black box explanation. Section III introduces the problem settings and preliminaries of explaining a text classifier. Afterwards, Section IV details the proposed method XPROAX. To qualitatively and quantitatively evaluate the proposed method, exhaustive experiments are designed and performed in Section V. Finally, we discuss future research directions and conclude the paper in Section VI.

<sup>1</sup>The source code is available at <https://github.com/caiy0220/XPROAX>

## II. RELATED WORK

There have been many studies on black box explanation over the past few years [3], [8]. Depending on whether the explanation is input-specific, we can categorize them as: *global* or *local* explanation. These explanation methods can be further categorized by the requirements on black box, namely model-agnostic or model-specific. As aforementioned, this work falls into the category of local model-agnostic explanation method.

A local model-agnostic method generates explanations for specific instances. The instance, which is classified by a black box, is given as the input of the explanation method along with the black box. To investigate the decision making process, neighbors around the given instance are generated in the first place by altering feature values [18], [27]. Closer (more similar) neighbors are assigned with higher weights as they are considered to be more representative of the locality and vice versa. After that, a transparent model (so-called surrogate model) is trained on these generated neighbors with the intention to simulate exact local behaviors of the black box [13], [25]. An explanation can then be constructed by analysing the local predictor.

As the first step, the quality of generated neighbors can make a huge difference to the final explanation [17]. Some methods like [25] and [33] perform a simple random perturbation. Basically, random perturbation modifies the values of several randomly selected features within a small range, it iterates the slight fluctuation in the vicinity of given input until sufficiently many neighbors have been generated. Although this simple approach generates satisfactory neighbors for most data representations, it requires fine-tuning for more complicated forms of data, which is textual data in our case. One difficulty of applying perturbation to text is how to modify the value of each feature (word). LIME [25] and SHAP [20] suggested *word dropping* as a solution – randomly dropping words from the input text. Such an approach is truly intuitive and has been shown to be effective. However, there still remain two problems:

- Without considering the internal relationships between words, most of the perturbed instances are incomplete and meaningless.
- The number of possible mutations is limited by the length of the input text. This becomes crucial when the input text is extremely short since the surrogate model will not have enough samples for approximating the black box.

Alternatively, XSPELLS [16] proposes to generate neighbors of text in a latent space learned through a variational autoencoder [4]. Prior to XSPELLS, ALIME [29] and ABELE [12] make use of autoencoders to construct a latent space for explaining decisions over tabular and image data respectively. ALIME still generates neighbors in the original feature space of an input, the autoencoder only serves as a weighing function for computing distances in the latent space. ABELE deploys an adversarial autoencoder [22] for generating synthetic images. XSPELLS fine-tunes the framework of ABELE so that it can deal with text classifiers. It constructs the neighborhood of an

input text in a latent space. This way, it tackles the limitation of *word dropping* by generating realistic instances with a generative model. However, the choice of the generative model impacts the final explanations. Ability of reconstruction and a properly organized latent space are highly desired features from the generative model in order to produce neighbors that satisfy the locality constraint. Also, a random sampling approach will not be able to efficiently and effectively explore a high-dimensional latent space, so the generation strategy in a latent space is worth further discussing. Besides, both ABELE and XSPELLS train the local predictor in the latent space. We argue that building the local predictor in a latent space can be a choice only if there are sufficient investigations on corresponding latent features.

The proposed method extends XSPELLS by addressing the issues mentioned above. It differs from other approaches in the literature because of the progressive approximation of the neighborhood and the diversity-based (counter-)factual selection.

## III. DEFINITIONS AND PRELIMINARIES

Let  $b(\cdot)$  be a black box model like a Deep Neural Network or a Random Forest. We have no access to the internal mechanism of the black box except its decisions for given instances. Specifically, for a query text instance  $x$ , the black box returns the confidence score vector<sup>2</sup>  $b(x) = \langle p_+, p_- \rangle$ . The class with the maximum probability is returned to the user, i.e.,  $\hat{y} = \arg \max(b(x))$ . The goal is to explain the decision of the black box for the given instance  $x$ , i.e.,  $b(x)$ . Our method uses the probability vector  $b(x)$  rather than only the final decision  $\hat{y}$  to comprehensively study the behavior of the black box.

Our approach falls into the category of local model explanation methods which use local surrogate models to approximate the predictions of the underlying black box model. These models are built in the “neighborhood” of the instance  $x$  to be explained. Let  $N(x)$  denote the neighborhood<sup>3</sup> of  $x$  and let  $m(\cdot)$  be the local surrogate model trained upon  $N(x)$ . For the neighborhood generation, we assume a generative model  $G$  that approximates the underlying data distribution (Section IV-A). Contrary to existing approaches that generate neighborhoods by sampling randomly in the latent space [16], [35], we propose to *progressively* approximate the neighborhood of  $x$  using *counterfactual instances-as-landmarks* (Section IV-B). The generated instances  $N(x)$  are then labeled by the black box model  $b(\cdot)$  and comprise the training set for learning the local surrogate model  $m(\cdot)$ . The progressive neighborhood approximation helps us to identify the *turning points* for the class decision of  $x$  and explain its classification (Section IV-C2) together with information inferred from the neighborhood (Section IV-C1).

<sup>2</sup>For simplicity, we assume a binary classification problem:  $y \in \{+, -\}$ .

<sup>3</sup>When it is clear from the context, we denote the neighborhood of  $x$  by  $N$  instead of  $N(x)$ .

#### IV. XPROAX: LOCAL EXPLANATIONS WITH PROGRESSIVE NEIGHBORHOOD APPROXIMATION

A surrogate model approximates the predictions of the underlying black box model locally. In order for the limited capacity model  $m(\cdot)$  to be on par with the complex model  $b(\cdot)$ , the training instances  $N(x)$  for building  $m(\cdot)$  should be carefully generated locally in the neighborhood of  $x$ .

To this end, we propose a local model-agnostic explanation method called XPROAX. The method approximates the neighborhood of the instance to be explained in the latent space (Section IV-B), which is constructed with the locality requirement being taken into account (Section IV-A). Subsequently, we derive the explanations from the constructed neighborhood (Section IV-C).

##### A. Locality-preserving neighborhood generation

To build the neighborhood  $N$  of the instance  $x$  to be explained, we need a mechanism that generates “meaningful” neighbors around  $x$ . A simple idea would be to use corruption [15], which removes or replaces words at random. Nevertheless, this process can result in semantically and/or syntactically invalid texts. Generative autoencoders have been successfully used for text generation, however not all models map similar texts to nearby latent vectors [38]. Therefore, for the generative model we use DAAE [31], an approach that maps similar texts to similar latent representations by augmenting the adversarial autoencoder with a denoising objective where original texts are reconstructed from their perturbed versions. DAAE learns neighborhood-preserving text representations while producing high quality texts with properly organized latent space geometry. Moreover, DAAE retains the ability to reconstruct input texts in contrast to VAE [4], a popular text generation method, that loses its reconstruction ability due to the KL regularization term [7], [34].

The DAAE model,  $G$ , is trained in an unsupervised manner upon data from the same domain, this dataset is denoted by  $X_G$ ; please note that  $X_G$  does not need to include the data which are used for training the black box model  $b(\cdot)$ . A randomly sampled subset  $X'_G \subseteq X_G$  is retained as a *corpus* for selecting the *initial landmarks*; the corpus is down-sampled for efficiency but due to random sampling it should still reflect the distribution of the dataset. DAAE consists of: A deterministic encoder  $E : \mathcal{X} \rightarrow \mathcal{Z}$  that maps from the text space to the latent space and a probabilistic decoder  $D : \mathcal{Z} \rightarrow \mathcal{X}$  that reconstructs word sequences from latent representations.

To generate the neighbors of a given input  $x \in \mathcal{X}$ , the generative model  $G$  will firstly map  $x$  to the latent space and will represent it as a latent vector  $z = E(x)$ . A set of neighboring vectors can then be created in the latent space by manipulating the latent vector  $z \in \mathcal{Z}$ . The manipulation includes pure random perturbation [16], [20], [25], locality-based sampling [17] and neighborhood generation using genetic algorithms [13]. Clearly, a simple random perturbation cannot lead to a good neighborhood  $N(x)$  and consequently, to a good surrogate model  $m(\cdot)$ . Reference [17] addresses

this problem by sampling instances within a fixed-size hypersphere as specified by the closest counterfactual to the query instance  $x$ . However, within this sphere instances are sampled uniformly. Reference [13] uses a genetic algorithm to create a synthetic neighborhood around  $x$ ; however, the employed fitness function only focuses on instances similar to  $x$  and moreover, the solution instances might converge to the same or very similar instances over the epochs.

To overcome the limitations of the above methods and better approximate the local data structure around  $x$ , we propose a progressive approximation of the neighborhood using landmarks to delimit the neighborhood boundary and by carefully interpolating within this boundary. The initial landmarks are counterfactual instances from the corpus  $X'_G$ , which are *iteratively* replaced by interpolated counterfactual instances<sup>4</sup> thus allowing for a refinement of the neighborhood. The interpolation is two-staged: i) between counterfactual instances to allow for the exploration of the neighborhood and a larger variety of counterfactuals; ii) between a counterfactual and the pivot point<sup>5</sup>  $z$  to allow for a more localized approximation of the decision boundary, in contrast to existing approaches that mainly sample uniformly within the neighborhood [16], [20], [25]. Once the neighborhood in the latent space  $Z$  is refined, the input instances for building the surrogate model  $m(\cdot)$  can be derived using the decoder as:  $N = D(Z)$ .

##### B. Neighborhood generation with progressive boundary approximation

A straightforward approach to local boundary approximation in the latent space would be to use the nearest neighbors with the opposite label of  $x$  as being evaluated by the cosine distance, since the distance of vectors in the latent space indicates the dissimilarity of texts [31]. Let  $\text{dist}(z_i, z_j)$  be the cosine distance between two instances  $x_i, x_j$  with corresponding latent vectors  $z_i, z_j$ . While texts are not evenly mapped to the latent space (typically following a Gaussian distribution [4]), a fix range that constraints the neighborhood within a certain hyper-sphere is not optimal for texts that locate in various manifolds of the latent space. Therefore, we propose to use some of the preserved corpus  $X'_G$  instances as initial landmarks to constrain the locality in the latent space. In particular, we use the *k-closest counterfactuals* in the corpus as the *seed landmarks*. To find these landmarks, the counterfactuals (the instances of  $X'_G$  predicted by  $b(\cdot)$  as of the opposite label of  $x$ ) are mapped to the latent space via the encoder. The *k*-closest to  $x$  counterfactuals in the latent space comprise the set  $C$  of landmark points. Intuitively, the landmarks reflect the local distribution of vectors around the input  $x$ , the usage of landmarks maintains the exploration within the local manifold and results in producing high quality neighbors.

Still  $C$  depends on the available corpus  $X'_G$  and further refinement is possible. We further refine the local neigh-

<sup>4</sup>Counterfactual instances can be either real instances from the corpus or synthetic instances generated during the latent space exploration.

<sup>5</sup>That is, the latent space representation of the input  $x$ .

---

**Algorithm 1** Neighborhood approximation

---

**Input:**  $x$ : query instance;  $C$ : set of counterfactual landmarks**Output:**  $N_{new}$ : generated neighbors;  $C_{new}$ : updated  $C$ 

```

1:  $C_{new} = \emptyset$ 
2: repeat
3:    $z_p, z_q = \text{RandomlyDraw}(C, 2)$ 
     # draw randomly 2 vectors from the landmark set
4:    $Z' = I(z_p, z_q)$ 
     # 1st interpolation: between counterfactuals
5:    $Z = \emptyset$ 
6:   for  $z_i \in Z'$  do # 2nd interpolation: between polarities
7:     if  $b(D(z_i)) \neq b(x)$  then
8:        $Z \leftarrow Z + I(z_i, E(x))$ 
9:     end if
10:  end for
11:   $z_c = \text{ClosestCounterfactual}(Z, 1)$ 
12:   $C_{new}.insert(z_c)$ 
13:   $N_{new} \leftarrow N_{new} + D(Z)$ 
14: until  $k$  times
15: return  $N_{new}, C_{new}$ 

```

---

neighborhood constrained by the landmarks  $C$ , using interpolation, a commonly used technique in autoencoder-based text generation [4], [34], [38] which implements successive and meaningful text manipulations based on two given prototypes through latent vector operations. The result of the interpolation between two latent vectors  $z_p, z_q$  is given by (1):

$$I(z_p, z_q) = \{z_i \mid z_i = z_p + i \cdot \frac{(z_q - z_p)}{s}, 0 \leq i \leq s\} \quad (1)$$

$s$  is the number of steps to be taken during the interpolation. For  $s = 1$ , only one instance is interpolated between  $z_p, z_q$ .

The neighborhood construction process (Algorithm 2) is an iterative process, in each iteration a better neighborhood approximation is achieved as closer counterfactuals to  $x$  are discovered and exploited. And more neighboring instances are interpolated between  $x$  and the discovered counterfactuals (Algorithm 1).

The neighborhood approximation based on a set of landmarks  $C$  is shown in Algorithm 1. It takes as input the set of current landmarks  $C$  and returns a new set of landmarks  $C_{new}$  and a new set of neighboring instances  $N_{new}$ . The *first-stage interpolation* (line 4) takes place between two counterfactuals that are randomly sampled from the landmark set  $C$  with replacement (line 3). For every point created by the first-stage interpolation (line 6), we reconstruct the text from the latent vector and check its label assigned by the black box  $b(\cdot)$  (line 7). The generated points  $Z'$  can have an opposite label (counterfactuals) or same label (factuals) as  $x$ . In the *second-stage interpolation*, we only interpolate between the input point  $x$  and the generated points  $Z'$  holding the opposite label to  $x$  (line 8). We keep all the neighbors generated through the second stage in a set  $Z$  (line 8). Subsequently, the new landmark set absorbs the closest counterfactual in  $Z$  (line 11, 12), the set of neighbors is expanded by neighboring texts

---

**Algorithm 2** Neighborhood construction

---

**Input:**  $x$ : query instance;  $C$ : initial set of landmarks**Output:**  $N$ : neighborhood of  $x$ 

```

1: while not terminate() do
2:    $N_{new}, C_{new} = \text{NeighborhoodApproximation}(C, x)$ 
3:    $C = C_{new}$ 
4:    $N \leftarrow N + N_{new}$ 
5: end while
6:  $N = \text{RemoveDuplicates}(N)$ 
7:  $N = \text{Closest}(N, n, n)$ 
   # output  $n$  closest instances for both classes
8: return  $N$ 

```

---

reconstructed from the latent vectors in  $Z$  (line 13). The process is repeated  $k$  times (line 2, 14), allowing a better exploration of the constrained local space.

The neighborhood construction process (Algorithm 2) runs iteratively until the number of iterations exceeds the limit or when no closer counterfactuals have been found in successive iterations. In each iteration, the method updates the set of counterfactual landmarks  $C_{new}$  and generates new neighbors  $N_{new}$  (line 3) based on the neighborhood approximation Algorithm 1 (line 2). Newly generated neighbors are stored as candidates for the final selection (line 4). After the completion of the construction, we eliminate duplicates contained in  $N$  (line 6) and output the  $n$  closest factuals and the  $n$  closest counterfactuals as the neighborhood  $N$  for explaining the decision of  $b(\cdot)$  over  $x$ .

Through this iterative process a better set of neighboring instances is discovered. The neighborhood achieves higher diversity due to the first stage interpolation and highlights the decision boundary between factuals and counterfactuals through the second stage interpolation.

### C. Extracting local explanations

The explanations are based on the generated neighborhood  $N(x)$ . In particular, we propose two explanation components: i) feature/word-level importance (Section IV-C1) and ii) factual and counterfactual instances identified during the neighborhood approximation process (Section IV-C2).

1) *Word-level explanation:* We train the surrogate model  $m(\cdot)$  upon the neighborhood  $N(x)$  generated in the previous step. In particular, we train a linear regression model in the original feature space. Unlike [16], we prefer the original feature space over the latent feature space because of the comprehensive understanding on the original features (words), this helps us in deriving knowledge from the surrogate model. We represent the texts using the bag-of-words model in order to simplify the complexity of the classification task. The sequential information vanishes as a trade-off between integrity and feasibility. The model is trained using the weighted square loss function  $\mathcal{L}$  [25] defined in (2). The first term in (2) indicates the weight of a generated neighbor  $x_i$  calculated from its latent space distance to the input  $x$  and  $\sigma$  is the kernel width controlling the influence of the weighting function. The

second term is the square loss of the surrogate model  $m(\cdot)$  on simulating local behavior of the black box  $b(\cdot)$ .

$$\mathcal{L}(N, x) = \sum_{x_i \in N} \exp\left(\frac{-\text{dist}(E(x_i), E(x))^2}{\sigma^2}\right) \cdot (b(x_i) - m(x_i))^2 \quad (2)$$

After training, the weights of features in the regression model can be derived as the importance values of the features/words. Along with words consisting in the input text, words appearing in the novel neighbors are also rated. Therefore, we split this explanation part into two subparts: i) words in the text  $x$  to be explained, we refer to them as *intrinsic features/words*; ii) words that only appeared in the neighborhood  $N(x)$ , we refer to them as *extrinsic features/words*.

The weights assigned to intrinsic features reveal the most important features in an input that cause the decision. Extrinsic features only appearing in the neighborhood are incorporated to facilitate further explanation of the decision [21], [36]. The relationships between the input and the extrinsic words values are indistinct. For the explanation, however, it is important to “connect” extrinsic words to the original instance. Our idea is to do this by demonstrating how they affect the decision on the input text. In particular, inspired by deletion on the most important intrinsic words from the input text, we can demonstrate the impact of an extrinsic word on the decision by inserting it into the input text. Although a random insertion/replacement of word will already affect the decision of  $b(\cdot)$ , we aim at determining the optimal edition with a given extrinsic word to have a manually created and meaningful (counter-)factual. This can be achieved by finding a position of insertion/replacement, at which the likelihood of the appearance of an extrinsic word is maximized according to the neighborhood  $N$  as stated. In (3), we measure the likelihood of a word’s appearance as the conditional probability under the local context, which consists of  $l$  preceding and  $l$  succeeding words. The conditional probability of an extrinsic word  $w$  can be estimated in the neighborhood. The variable  $i$  denotes a possible position for the edition, the position that maximizes the likelihood will be chosen as the optimal solution. To allow the integration of the local contexts with the logarithm, a minute value  $\epsilon \ll 1$  is involved as the last part of (3).

$$\arg \max_i \sum_{j=-l}^l \log(P(w|w_{i+j}) + \epsilon) \quad (3)$$

The example presented in Table I demonstrates a manually created counterfactual through the edition based on an extrinsic word. The edition builds the connection between the input “*would not recommend*.” and the extrinsic word “*definitely*”, it reveals that the adverb describing the word *recommend* plays an important role in switching the black box decision.

2) *Factuals and Counterfactuals as explanation*: As a benefit from the generative model, the generated neighbors can also be used as factuals and counterfactuals for further explanation. The neighboring texts illustrate the changes that are relevant or irrelevant to reverse the original prediction result of  $b(\cdot)$ .

TABLE I  
AN EXAMPLE OF THE EXTRINSIC WORD BASED EDITION

<b>Input</b> : would not recommend .	$b(\cdot)$ : Negative
<b>Extrinsic word</b> : definitely: -0.17	
<b>Edition</b> : would definitely recommend .	$b(\cdot)$ : Positive

Instead of focusing on neighbors with minimum latent space distances, which are most likely to generate homogeneous differences to the input text, we take the diversity of the (counter-)factuals into account. By having the diversity, it becomes easier to underscore the most important components for the decision even with a limited number of (counter-)factuals, such components tend to remain unchanged in factuals and are most likely to be changed in counterfactuals. To achieve this, we form a (counter-)factual set by picking one instance at a time until sufficiently many (counter-)factuals are picked. Each time, the closer instance that maintains the diversity of the (counter-)factual set will be absorbed. Equation (4) quantifies the quality of an instance according to our expectation on the (counter-)factual set, we construct the (counter-)factual set  $\Xi$  by iteratively picking a neighbor with the maximum  $r_i$  value.

$$r_i = (1 - \lambda) \cdot -\text{dist}(z_i, z) + \lambda \cdot \sum_{\Xi \cup \{z_i\}}^{\substack{z_p \neq z_q \\ \Xi \cup \{z_i\}}} \frac{\text{dist}(\Delta z_p, \Delta z_q)}{(|\Xi|^2 + |\Xi|)/2} \quad (4)$$

(where  $\Delta z_i = z_i - z$ )

Equation (4) consists of two parts: i) the first part measures the distance from a neighbor to the input  $x$  in the latent space; 2) the second part measures the normalized diversity [10] of the (counter-)factual set  $\Xi$  after absorbing a neighbor  $z_i$ .  $\lambda \in [0, 1]$  is a user defined parameter to control the importance of the diversity during the construction of the (counter-)factual set. The first term of (4) uses a cosine distance metric and has a negative contribution to the quality, because it is designed to have the closer instances being preferred during the selection. In the second part, we do not directly measure the diversity among the instances from a constrained neighborhood as it could be subtle. Instead, we measure the diversity of the difference vectors denoted by  $\Delta z_i$ , which stands for the difference between the pivot point  $z$  and a neighbor  $z_i$ . The second part ensures the selected (counter-)factuals manipulate the input in a different way to maintain the diversity. By greedily picking the neighbor with the highest quality value  $r$ , the method ensures the diversity of (counter-)factuals. As the explanation set is changing, the quality of a neighbor has to be updated after the absorption of an entry.

## V. EVALUATION

The first goal of our experiments is to qualitatively evaluate the generated explanations in terms of usefulness and stability (Section V-B). Secondly, we provide a quantitative evaluation by measuring the confidence drops of the black box predictions after removing important features/words (Section V-C). Details on the datasets, baselines and parameter selection are provided in Section V-A.

### A. Experimental setup

**Datasets:** We evaluate our approach on two real-world datasets: Yelp reviews [30] and Amazon reviews polarity [37]. The *Yelp dataset* consists of restaurant reviews, each of which does not exceed 16 words. Reviews are labeled as either positive or negative. The *Amazon review polarity dataset* consists of customer reviews about different products with rating ranging from 1 star to 5 stars; [37] discarded all neutral reviews (with 3 stars) and re-defined the reviews with  $\leq 2$  stars ( $\geq 4$  stars) as negative (positive, respectively). Since we focus on short texts, we only use review titles in our experiments. We split both datasets into 4 parts: 200K/20K/2K/4K, corresponding to the training set for the generative model and the training/validation/test set for the black box. Besides, we used 20K instances randomly drawn from the training set of the generative model as the corpus for initializing the landmark counterfactuals. The subset for the generative model is notably larger compared to the dataset for the black box model to ensure a good quality generator; this is feasible as DAAE is trained in an unsupervised manner.

**Black box models:** As black box models, we used a Random Forest (RF) and a Deep Neural Network (DNN). For the RF, we used the implementation from the *scikit-learn* library, setting the number of weak learners to 400 and using the TF-IDF vectorizer to transform the raw texts into vectors. For the DNN, we used the *Keras* library. We implemented a seven layer DNN; the central layer of the hidden layers is a LSTM [14], which captures the sequential information of texts; the rest are fully connected layers with a ReLU activation function. To train the DNN with texts of various lengths, we transformed raw texts into dense embeddings in the pre-processing stage with a tokenizer and a padder, the padder ensures that all dense embeddings share the same length.

**Neighborhood generation and Surrogate model:** We trained a DAAE model following [31] as the generative model; encoder  $E$  and decoder  $D$  are one-layer LSTMs with hidden dimension 1024 and word embedding dimension 512. The encoder finally projects an input  $x$  into a latent code  $z$  with 128 dimensions for both datasets. For the neighborhood approximation, we set the interpolation step  $s$  to 10 and the size of the landmark set  $k$  to 25. These are determined according to our experiments on the effect of parameters<sup>6</sup>. The experiments showed that increasing the interpolation step  $s$  and the size of landmark set  $k$  have a positive impact on the results because it allows to exhaustively explore the space and refine the neighborhood construction. Our surrogate model  $m(\cdot)$  is a linear regression model trained upon the neighborhood  $N(x)$ . Details including the performance of the black box and the reconstruction loss  $\mathcal{L}_{rec}$  of the generative model are shown in Table II.

**Competitors:** We compared our approach against:

- **LIME** [25]: the most widely used local explanation method that applies a simple random perturbation to the

<sup>6</sup>Due to space limitation the detailed experiments on the effect of parameters are not presented.

input  $x$  for constructing the neighborhood  $N(x)$ .

- **XSPELLS** [16] generates neighbors from the latent space as well and focuses on the textual data.

### B. Qualitative evaluation

The evaluation of explanations remains a challenging task with no consent on how it should be conducted. For the qualitative part, we provide explanations from different methods on selected instances and assess the explanation quality with human-grounded evaluation [5], [8]. In particular, we analysed the explanations from two different perspectives:

- **Usefulness:** how the explanations help human understanding on a model decision that has been made.
- **Stability** [1], [5]: how similar the explanations for similar instances are.

1) **Usefulness:** Table III presents the explanations generated by different methods for two different inputs. The **explanation from XPROAX** consists of four parts: i) important *intrinsic words* from the input text  $x$  followed by their importance according to the surrogate model  $m(\cdot)$  ii) important *extrinsic words* that only appeared in the neighborhood  $N(x)$  (but not in the input) as evaluated by the surrogate model  $m(\cdot)$ , iii) top-5 *factuals* from  $N(x)$ , iv) top-5 *counterfactuals* from  $N(x)$ . We consider a word to be important if the absolute value of its importance is greater than the threshold 0.1, which holds too for LIME. The **explanation from XSPELLS** also contains four components: i) top-5 factuals, ii) top-5 counterfactuals, iii) the most common words in factuals, and iv) the most common words in counterfactuals. The **explanation from LIME** is represented by a list of the most contributing words followed by their importance. For XPROAX and LIME, we present the most contributing words with their contributions to current decision, where a positive number indicates the word is contributing to the prediction result. Conversely, a word with a negative importance value contributes to the opposite class. Words with a higher absolute importance value are considered to be more important by the surrogate model  $m(\cdot)$ . Such information is not possible for XSPELLS, since its surrogate model is built in the latent space and therefore no interpretable features can be directly derived from such a model. Instead, XSPELLS outputs the most frequent words in the (counter-)factuals with their relative frequencies in parenthesis. However, these frequencies do not hold direct relation with the surrogate model and the sentiment tendency.

The **first instance**  $x$  to be explained is “*excellent start for beginning french students.*” from the Amazon dataset classified as positive by the DNN black box model. The word “*excellent*” is the most important intrinsic word identified by both XPROAX and LIME, though XPROAX considers it more important for the positive classification. The most important extrinsic words by XPROAX are “*incredible*” associated with the positive class and “*worst, disappointing*” associated with the negative class. These extrinsic words are all adjectives which can replace the intrinsic word “*excellent*” to generate manually created (counter-)factuals. Regarding

TABLE II  
DETAILS ABOUT DATASETS, BLACK BOXES AND GENERATIVE MODEL

Dataset	Avg. length of texts	Generative model Training Set	Black box			Accuracy <sup>a</sup>		$\mathcal{L}_{rec}$ <sup>a</sup> DAAE
			Training set	Valid set	Test set	RF	DNN	
Yelp	8.83	200K	20K	2K	4K	0.9113	0.9548	2.79
Amazon	8.47	200K	20K	2K	4K	0.7655	0.7795	3.93

<sup>a</sup>The reported values are performance on the test set for black boxes.

TABLE III  
EXAMPLE EXPLANATIONS BY DIFFERENT METHODS FOR TWO INPUT INSTANCES

Input 1	DNN $b(\cdot)$ : Positive, Dataset: <b>Amazon</b>	
XPROAX	<b>Intrinsic words:</b> excellent: 0.48, beginning: 0.10 <b>Factuals:</b> 1) excellent sta for beginning french philosophy. 2) excellent sta for beginning true _unk_ students. 3) excellent sta for the beginning war students! 4) excellent sta for beginning french law. 5) an excellent for beginning true _unk_ fans.	<b>Extrinsic words<sup>a</sup>:</b> worst: -0.35, disappointing: -0.34 <b>Counterfactuals:</b> 1) the worst book for the french history. 2) not a story for true king _unk_ fans. 3) a disappointing for special effects.. 4) an disappointing sta for original mystery. 5) sub par for the beginning division fan
	<b>Factuals:</b> 1) within num days pin 2) another toy book ever 3) old and errors 4) u wan na see 5) at all you may <b>Common words in factuals:</b> within (0.067), num (0.067), days (0.067)	<b>Counterfactuals:</b> 1) atrocious dvd used are 2) nice bra but 3) lousy range only 4) no substance vhs player 5) after num months but why <b>Common words in counterfactuals:</b> atrocious (0.083), dvd (0.083), used: (0.083)
LIME	<b>Most contributing words:</b> excellent: 0.11, beginning: 0.10	
Input 2	Random Forest $b(\cdot)$ : Positive, Dataset: <b>Yelp</b>	
XPROAX	<b>Intrinsic words:</b> worth: 0.51 <b>Factuals:</b> 1) the fries were n't worth coming. 2) _unk_ <sup>b</sup> are n't worth going back. 3) the fries were worth coming back. 4) the fries were worth going back. 5) you do n't be worth coming.	<b>Extrinsic words<sup>a</sup>:</b> not: -0.66 perfect: 0.28 <b>Counterfactuals:</b> 1) _unk_ do n't bother in back. 2) _unk_ do n't bother going back. 3) _unk_ do n't be anybody back. 4) a few fries were definately coming back. 5) _unk_ do n't be anybody.
	<b>Factuals:</b> 1) it seems well they did 2) and i feel like on service 3) dave is excellent 4) everything we will get 5) all i hung up is nice <b>Common words in factuals:</b> seems (0.091), well (0.091), feel (0.091)	<b>Counterfactuals:</b> 1) both to die 2) all else s 3) every i may 4) who makes me money last 5) all were nt pricey <b>Common words in counterfactuals:</b> die (0.111), else (0.111), every: (0.111)
LIME	<b>Most contributing words:</b> worth: 0.59	

<sup>a</sup>Words with high importance that only appeared in the neighbors (not appeared in the input text).

<sup>b</sup>Generic unknown word token for words out of the vocabulary.

(counter-)factuals, XPROAX generates meaningful and complete instances with similar sentence structure. The factual “an excellent for beginning true \_unk\_ fans.” supports the observation from the intrinsic words. Although most of the words in the input text are manipulated, the sentiment remains unchanged with the presence of the important intrinsic words. On the contrary, XSPELLS suffers from poor (too generic and incomplete) (counter-)factuals like “lousy range only”, which barely make any connection to the input. So, the explanation by XPROAX is better and more complete as it covers both intrinsic/extrinsic aspects as well as (counter-)factuals. The **second instance** to be explained is “fries are n't worth coming back” from the Yelp dataset which is classified as positive by the RF black box model, despite the existence of the negation. Again, XSPELLS does not work properly as most

of generated (counter-)factuals are meaningless, e.g. “every i may”. Lacking notable relationships with the input, these (counter-)factuals fail in explaining the cause of the decision. On the other hand, XPROAX and LIME both perform fairly well. They have an agreement on the primary term that causes the positive decision – the word “worth”. According to this observation, we can already tell that the RF model either has trouble with understanding negation or simply cannot deal with the term “n't”. XPROAX takes one further step to identify the exact reason. In extrinsic words, the word “not” is the highest weighted word in the neighborhood and contributes to a negative sentiment, the observation proves that the RF model is capable to handle the negation under a positive context. Specifically, the decision can be reversed by replacing the term “n't” with the extrinsic word “not”. The

TABLE IV  
STABILITY OF EXPLANATIONS

Dataset: <b>Yelp</b> . Black box model $b(\cdot)$ : DNN				
Input Text	$b(\cdot)$	XPROAX	LIME	XSPELLS <sup>a</sup>
<u>great</u> food.	<b>1.00</b>	great 0.69 food 0.05	<u>great</u> 0.70 food -0.24	food 0.11 italian* 0.08
<u>great</u> sushi.	<b>0.99</b>	great 0.62 sushi 0.12	great 0.56 sushi -0.11	salsa 0.17 guys 0.14
<u>great</u> pizza.	<b>1.00</b>	great 0.70 pizza 0.01	great 0.73 pizza -0.27	good 0.14 food 0.08
<u>great</u> beer.	<b>1.00</b>	great 0.64 beer 0.02	great 0.73 beer -0.27	good 0.15 story* 0.07
great <u>food</u> .	<b>1.00</b>	great 0.69 <u>food</u> 0.05	great 0.70 <u>food</u> -0.24	<u>food</u> 0.11 italian* 0.08
amazing <u>food</u> .	<b>0.99</b>	amazing 0.64 <u>food</u> -0.01	amazing 0.70 <u>food</u> -0.24	good 0.14 <u>food</u> 0.14
horrible <u>food</u> .	<b>-1.00</b>	horrible 0.45 <u>food</u> 0.02	horrible 0.04 <u>food</u> -0.02	best* 0.17 good* 0.17
bad <u>food</u> .	<b>-1.00</b>	bad 0.43 <u>food</u> 0.02	bad 0.04 <u>food</u> -0.02	always 0.29 story* 0.17

<sup>a</sup>XSPELLS outputs the relative frequency of a word, a word without "\*" comes from factuals, a word with "\*" comes from counterfactuals.

other word, “*perfect*”, has a positive tendency and it is an alternative of the “positive” description of the fries. Regarding the (counter-)factuals, disappearance of the extrinsic words is a flaw. This is caused by the fact that the neighbors containing extrinsic words are not close enough to the pivot point and therefore are excluded even under the consideration of diversity. But still, factuals like “*the fries were worth coming back*.” reveal the ignorance of the black box on the term “*n*’” since there is no obvious influence on classification results regardless of its presence or absence.

2) *Stability*: An explanation method with higher stability is expected to generate similar explanations for similar inputs. However, there is a lack of a formal definition of the similarity of the input instances, especially for textual data. Intuitively, similar texts should be semantically and syntactically close. For this reason, we manually selected several instances from the test set. To highlight the similarity, we prefer shorter phrases under the same context with exactly the same grammatical structure. Table IV shows the stability of generated explanations by different methods for similar texts. In particular, for an input text (the first column), we show in the second column the confidence score of a black box prediction  $b(\cdot)$ , a DNN in this case, in the negative (negative value) or positive (positive value) class. For XPROAX and LIME, the importance values of intrinsic words are presented. The number following a word is the importance of the word considered to be for the current decision. A positive value means the word results in the current prediction while a negative value leads to the opposite. XSPELLS does not output word-level importance but only relative frequencies of the most common words (only the top-2 words are presented).

The DNN model is confident about its decisions since the confidence scores of the predicted class are nearly 1.0. In the first 4 comments from the table, we focus on the term “*great*”, it is used in the same context to praise dish/drink.

For the explanations from XPROAX and LIME, the weights of the target term are stable with subtle fluctuation. For the last 4 instances, our main focus is the underlined term “*food*”. In these phrases, the neutral noun is described by different adjectives. XPROAX generates stable explanations by highlighting the adjectives and assigning small weights to the neutral term. In the meantime, LIME fails in producing stable explanations for the decisions that have been made. The importance of the target term varies a lot in the same context. Moreover, although the black box is confident about the negative predictions on the last two examples, LIME considers no words in the last two negative phrases to be important. A possible reason that causes such an issue for the perturbation-based explanation method is the size of the training set, which is constructed by the neighboring points. For a method that generates neighbors by randomly dropping words, the upper bound of the size of the training set is constrained by the length of the input text. Such a shortcoming becomes critical when an input contains only a few words, since the surrogate model cannot be fairly trained with few neighbors and may cause the unstable outputs as presented in the results. XSPELLS performs relatively poorly as well. Its explanations cannot concentrate on the given input and are less stable.

*Summary*: The above experiments show that XPROAX generates high quality explanations, which consist of informative intrinsic and extrinsic features/words as well as meaningful (counter-)factuals. Our feature/word-level explanations are close to those produced by LIME, but the proposed method generates more stable explanations under the manually selected context. In addition, extrinsic words and (counter-)factuals extracted from the carefully constructed neighborhood promote understandings on the black box decisions. Contrariwise, due to the purely random sampling in the latent space, the (counter-)factuals generated by XSPELLS are not meaningfully related to the input to be explained and consequently, this also holds for the most common words generated from the (counter-)factuals. The comparison between (counter-)factuals generated by XPROAX and XSPELLS intuitively demonstrates the necessity of satisfying locality in the latent space.

### C. Quantitative evaluation

For the quantitative evaluation, we follow the three Cs of interpretability [32], namely Completeness, Compactness, and Correctness. Traditionally, the evaluation includes removing words from the input text in the order of the weights assigned by the explanation method and checking the drop of the prediction confidence [2], [19]. Instead of word deletion, we propose to involve the words assigned with negative values in the evaluation. Same as deleting a positive word (contributing to the predicted class), the insertion of a negative word (contributing to the opposite class) is also expected to cause the drop of confidence in the current prediction. We named the expanded method as *sentence edition*. The edition process is constrained by setting a minimum threshold  $\eta$  for



word importance. Only words with absolute value above the threshold are considered to be relevant. Words with positive values are dropped from the sentence while those with negative values are inserted. For both XPROAX and XSPELLS, the inserted word can come from the neighbors and the optimal position of insertion is determined by (3); for LIME, we only perform a random insertion as there is no extra information provided.

The confidence drop after the explanation-guided edition indicates the *completeness* of an explanation. The averaged confidence drop per operation (deletion/insertion) reveals the *compactness*. For correctness, we consider that weights of features should reflect their contributions to a prediction. In other words, deletion/insertion of features with higher weights should have larger impacts on the prediction. Therefore, we report the change of compactness  $\Delta\eta$  after increasing the threshold  $\eta$  as *correctness*. A method with higher correctness is expected to generate a more compact explanation after increasing  $\eta$  since less relevant words are excluded. To further reveal the effectiveness of the proposed method and prove that confidence drop cannot be achieved by simply inserting words with strong opposite sentiment, we designed a **baseline** simulating the behavior of the edition guided by explanations. The baseline randomly drops up to 3 words from the input text, meanwhile, it randomly picks a strong sentiment word<sup>7</sup> from the opposite class and inserts the word into the text.

The results of designed experiments are reported in Table V. For confidence drop (completeness) and averaged confidence drop per operation (compactness), the reported values are the mean and the standard deviation of the performance of corresponding method over the test set. For the change of compactness “ $\Delta\eta$ ” (correctness), the reported value is the change of mean compactness after increasing the threshold. We empirically set the threshold  $\eta$  to 0.1 for measuring completeness and compactness. For correctness, we record the change of compactness when we raise  $\eta$  from 0.1 to 0.3. XPROAX outperforms the two competitors as well as the baseline in terms of completeness and compactness. In particular, it achieves the highest completeness in all settings and the highest compactness in three settings out of four. The statistics regarding correctness is in favor of LIME, as the lower compactness by LIME with the threshold  $\eta = 0.1$  becomes an advantage while evaluating the correctness. In other words, it leaves more space for improving the compactness which is used for calculating the correctness. Nevertheless, XPROAX still reaches a comparable correctness to LIME.

Admittedly, compared to LIME, XPROAX benefits from the neighboring extrinsic words. But the comparison to the designed baseline and XSPELLS demonstrates that XPROAX achieves better performance based on a better understanding of the black box behavior in the neighborhood as the opposite to simply inserting strong sentiment words from the other class.

<sup>7</sup>Top-100 words with the strongest sentiment are given by the black box model for both classes.

TABLE V  
CONFIDENCE DROP AFTER EXPLANATION-GUIDED EDITION,  $\eta = 0.1$

Dataset and Model	Explaining Method	Confidence Drop	Avg. Confidence Drop per op	$\Delta\eta$ (0.3-0.1)
Yelp & RF	baseline	0.247 $\pm$ 0.31	0.179 $\pm$ 0.24	/
	LIME	0.364 $\pm$ 0.29	0.297 $\pm$ 0.26	<b>+0.213</b>
	XSPELLS	0.132 $\pm$ 0.26	0.170 $\pm$ 0.27	+0.032
	XPROAX	<b>0.740 <math>\pm</math> 0.22</b>	<b>0.417 <math>\pm</math> 0.33</b>	+0.153
Yelp & DNN	baseline	0.136 $\pm$ 0.32	0.094 $\pm$ 0.23	/
	LIME	0.564 $\pm$ 0.46	<b>0.348 <math>\pm</math> 0.44</b>	<b>+0.230</b>
	XSPELLS	0.084 $\pm$ 0.26	0.102 $\pm$ 0.27	-0.014
	XPROAX	<b>0.825 <math>\pm</math> 0.35</b>	0.302 $\pm$ 0.43	+0.206
Amazon & RF	baseline	0.163 $\pm$ 0.19	0.118 $\pm$ 0.14	/
	LIME	0.209 $\pm$ 0.18	0.201 $\pm$ 0.16	<b>+0.185</b>
	XSPELLS	0.048 $\pm$ 0.13	0.058 $\pm$ 0.14	+0.037
	XPROAX	<b>0.506 <math>\pm</math> 0.20</b>	<b>0.354 <math>\pm</math> 0.21</b>	+0.126
Amazon & DNN	baseline	0.287 $\pm$ 0.27	0.209 $\pm$ 0.21	/
	LIME	0.424 $\pm$ 0.27	0.238 $\pm$ 0.17	+0.156
	XSPELLS	0.095 $\pm$ 0.18	0.122 $\pm$ 0.18	+0.037
	XPROAX	<b>0.665 <math>\pm</math> 0.21</b>	<b>0.298 <math>\pm</math> 0.25</b>	<b>+0.164</b>

## VI. CONCLUSIONS

In this work, we proposed XPROAX, a local model-agnostic explanation method for text classification that progressively approximates the neighborhood of the instance to be explained and consequently the local decision boundary. This is achieved by applying a two-staged interpolation in a neighborhood-preserving latent space. Our explanations consist of word-level importance (including intrinsic and extrinsic words) as well as (counter-)factual instances.

Our experiments, both qualitatively and quantitatively, show the effectiveness and stability of XPROAX in comparison to state-of-the-art methods. Moreover, the experimental results reveal that local explanations on text classifiers must not be limited by the words appearing in a given input (intrinsic words). Neighborhood exploration in a similar context provides a comprehensive view of understanding of the decision that has been made. In comparison to XSPELLS, the careful construction of the neighborhood overcomes the weakness of randomly sampling in the latent space. By taking advantage of the neighborhood construction, the explanations from our method shows notable relations to the input that XSPELLS cannot guarantee.

As future work, we firstly consider extending the explanation framework for long texts (e.g., documents). Neighborhood preserving text generation becomes more challenging when the length of text increases. In this regard, an elegant formatting for input texts and further study on the generative model are essential. Second, the choice of the surrogate model remains to be investigated. Although knowledge extraction from a complicated but still transparent surrogate model requires more efforts, accurate simulation on black boxes can be acquired in return. Finally, XPROAX is now specialized for textual data, so an interesting future direction is to expand the framework for other forms of data (e.g., image data).

## ACKNOWLEDGMENTS

The first author is supported by the State Ministry of Science and Culture of Lower Saxony, within the PhD program “Responsible Artificial Intelligence in the Digital Society”. We also thank Philip Naumann for the insightful discussions.

## REFERENCES

- [1] David Alvarez-Melis and Tommi S Jaakkola. On the robustness of interpretability methods. *arXiv preprint arXiv:1806.08049*, 2018.
- [2] Leila Arras, Franziska Horn, Grégoire Montavon, Klaus-Robert Müller, and Wojciech Samek. Explaining predictions of non-linear classifiers in nlp. In *Proceedings of the 1st Workshop on Representation Learning for NLP*, pages 1–7, 2016.
- [3] Francesco Bodria, Fosca Giannotti, Riccardo Guidotti, Francesca Naretto, Dino Pedreschi, and Salvatore Rinzivillo. Benchmarking and survey of explanation methods for black box models. *arXiv preprint arXiv:2102.13076*, 2021.
- [4] Samuel Bowman, Luke Vilnis, Oriol Vinyals, Andrew Dai, Rafal Jozefowicz, and Samy Bengio. Generating sentences from a continuous space. In *Proceedings of The 20th SIGNLL Conference on Computational Natural Language Learning*, pages 10–21, 2016.
- [5] Diogo V Carvalho, Eduardo M Pereira, and Jaime S Cardoso. Machine learning interpretability: A survey on methods and metrics. *Electronics*, 8(8):832, 2019.
- [6] Aditya Chattopadhyay, Anirban Sarkar, Prantik Howlader, and Vineeth N Balasubramanian. Grad-cam++: Generalized gradient-based visual explanations for deep convolutional networks. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 839–847. IEEE, 2018.
- [7] Bin Dai, Ziyu Wang, and David Wipf. The usual suspects? reassessing blame for vae posterior collapse. In *International Conference on Machine Learning*, pages 2313–2322. PMLR, 2020.
- [8] Marina Danilevsky, Kun Qian, Ranit Aharonov, Yannis Katsis, Ban Kawas, and Prithviraj Sen. A survey of the state of explainable ai for natural language processing. In *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing*, pages 447–459, 2020.
- [9] Jerome H Friedman, Bogdan E Popescu, et al. Predictive learning via rule ensembles. *Annals of Applied Statistics*, 2(3):916–954, 2008.
- [10] Zhiqiang Gong, Ping Zhong, and Weidong Hu. Diversity in machine learning. *IEEE Access*, 7:64323–64350, 2019.
- [11] Bryce Goodman and Seth Flaxman. European union regulations on algorithmic decision-making and a “right to explanation”. *AI magazine*, 38(3):50–57, 2017.
- [12] Riccardo Guidotti, Anna Monreale, Stan Matwin, and Dino Pedreschi. Black box explanation by learning image exemplars in the latent feature space. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 189–205. Springer, 2019.
- [13] Riccardo Guidotti, Anna Monreale, Salvatore Ruggieri, Dino Pedreschi, Franco Turini, and Fosca Giannotti. Local rule-based explanations of black box decision systems. *arXiv preprint arXiv:1805.10820*, 2018.
- [14] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [15] Vasileios Iosifidis and Eirini Ntoutsi. Sentiment analysis on big sparse data streams with limited labels. *Knowledge and Information Systems*, pages 1–40, 2019.
- [16] Orestis Lampridis, Riccardo Guidotti, and Salvatore Ruggieri. Explaining sentiment classification with synthetic exemplars and counter-exemplars. In *International Conference on Discovery Science*, pages 357–373. Springer, 2020.
- [17] Thibault Laugel, Xavier Renard, Marie-Jeanne Lesot, Christophe Marsala, and Marcin Detyniecki. Defining locality for surrogates in post-hoc interpretability. In *Workshop on Human Interpretability for Machine Learning (WHI)-International Conference on Machine Learning (ICML)*, 2018.
- [18] Vincent Lemaire, Raphael Féraud, and Nicolas Voisine. Contact personalization using a score understanding method. In *2008 IEEE International Joint Conference on Neural Networks (IEEE World Congress on Computational Intelligence)*, pages 649–654. IEEE, 2008.
- [19] Piyawat Lertvittayakumjorn and Francesca Toni. Human-grounded evaluations of explanation methods for text classification. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5198–5208, 2019.
- [20] Scott M Lundberg and Su-In Lee. A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems*, 4765–4774, 2017.
- [21] Ji Ma, Jingjiao Li, Zhenni Li, and Jiao Jiao. Salient region detection by integrating intrinsic and extrinsic cues without prior information. *Journal of Engineering Science & Technology Review*, 10(3), 2017.
- [22] Alireza Makhzani, Jonathon Shlens, Navdeep Jaitly, Ian Goodfellow, and Brendan Frey. Adversarial autoencoders. *arXiv preprint arXiv:1511.05644*, 2015.
- [23] Vitali Petsiuk, Abir Das, and Kate Saenko. Rise: Randomized input sampling for explanation of black-box models. *arXiv preprint arXiv:1806.07421*, 2018.
- [24] Gregory Plumb, Denali Molitor, and Ameet Talwalkar. Model agnostic supervised local explanations. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pages 2520–2529, 2018.
- [25] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. “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*, pages 1135–1144, 2016.
- [26] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. Anchors: High-precision model-agnostic explanations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
- [27] Marko Robnik-Šikonja and Igor Kononenko. Explaining classifications for individual instances. *IEEE Transactions on Knowledge and Data Engineering*, 20(5):589–600, 2008.
- [28] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pages 618–626, 2017.
- [29] Sharath M Shankaranarayana and Davor Runje. Alime: Autoencoder based approach for local interpretability. In *International conference on intelligent data engineering and automated learning*, pages 454–463. Springer, 2019.
- [30] Tianxiao Shen, Tao Lei, Regina Barzilay, and Tommi Jaakkola. Style transfer from non-parallel text by cross-alignment. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 6833–6844, 2017.
- [31] Tianxiao Shen, Jonas Mueller, Regina Barzilay, and Tommi Jaakkola. Educating text autoencoders: Latent representation guidance via denoising. In *International Conference on Machine Learning*, pages 8719–8729. PMLR, 2020.
- [32] Wilson Silva, Kelwin Fernandes, Maria J Cardoso, and Jaime S Cardoso. Towards complementary explanations using deep neural networks. In *Understanding and Interpreting Machine Learning in Medical Image Computing Applications*, pages 133–140. Springer, 2018.
- [33] Erik Strumbelj and Igor Kononenko. An efficient explanation of individual classifications using game theory. *The Journal of Machine Learning Research*, 11:1–18, 2010.
- [34] Tianbao Song, Jingbo Sun, Bo Chen, Weiming Peng, and Jihua Song. Latent space expanded variational autoencoder for sentence generation. *IEEE Access*, 7:144618–144627, 2019.
- [35] Georgios Vlassopoulos, Tim van Erven, Henry Brighton, and Vlado Menkovski. Explaining predictions by approximating the local decision boundary. *arXiv preprint arXiv:2006.07985*, 2020.
- [36] Meng Wang, Janusz Konrad, Prakash Ishwar, Kevin Jing, and Henry Rowley. Image saliency: From intrinsic to extrinsic context. In *CVPR 2011*, pages 417–424. IEEE, 2011.
- [37] Xiang Zhang, Junbo Zhao, and Yann LeCun. Character-level convolutional networks for text classification. In *Proceedings of the 28th International Conference on Neural Information Processing Systems-Volume 1*, pages 649–657, 2015.
- [38] Junbo Zhao, Yoon Kim, Kelly Zhang, Alexander Rush, and Yann LeCun. Adversarially regularized autoencoders. In *International conference on machine learning*, pages 5902–5911. PMLR, 2018.