

CIDR: A Cooperative Integrated Dynamic Refining Method for Minimal Feature Removal Problem

Qian Chen¹, Taolin Zhang², Dongyang Li¹, Xiaofeng He^{1,3*}

¹School of Computer Science and Technology, East China Normal University, Shanghai, China

²Alibaba Group

³NPPA Key Laboratory of Publishing Integration Development, ECNUP, Shanghai, China

qianchen901005@gmail.com, zhangtl0519@gmail.com, dongyangli0612@gmail.com
hexf@cs.ecnu.edu.cn

Abstract

The minimal feature removal problem in the post-hoc explanation area aims to identify the minimal feature set (MFS). Prior studies using the greedy algorithm to calculate the minimal feature set lack the exploration of feature interactions under a monotonic assumption which cannot be satisfied in general scenarios. In order to address the above limitations, we propose a Cooperative Integrated Dynamic Refining method (CIDR) to efficiently discover minimal feature sets. Specifically, we design Cooperative Integrated Gradients (CIG) to detect interactions between features. By incorporating CIG and characteristics of the minimal feature set, we transform the minimal feature removal problem into a knapsack problem. Additionally, we devise an auxiliary Minimal Feature Refinement algorithm to determine the minimal feature set from numerous candidate sets. To the best of our knowledge, our work is the first to address the minimal feature removal problem in the field of natural language processing. Extensive experiments demonstrate that CIDR is capable of tracing representative minimal feature sets with improved interpretability across various models and datasets.

Introduction

Deep neural networks have achieved remarkable accomplishments in the natural language processing domain (Vaswani et al. 2017; Devlin et al. 2019; Brown et al. 2020; OpenAI 2023). Meanwhile, the black-box nature of deep learning techniques has made it challenging to applications. In addition to model accuracy, there is a demand for considering interpretability (Olex, Maffey, and McInnes 2019; Mosca et al. 2022). Recently feature importance-based explanation methods have gained popularity due to fidelity and usability (Sikdar, Bhattacharya, and Heese 2021; Sekhon et al. 2023). For the text classification task, feature importance-based explanation methods predominantly investigate the impact of words on model prediction behavior (Li et al. 2016; Godin et al. 2018; Chen and Ji 2020). Earlier researches mainly utilize the LIME algorithm (Ribeiro, Singh, and Guestrin 2016), attention mechanism (Wiegrefe and Pinter 2019) or gradient saliency (Kindermans et al. 2018) to ascribe a score to an individual word, measuring

Input: Despite facing numerous challenges , we remain

steadfast .

Minimal feature set:  steadfast

Input: Despite facing numerous challenges , we remain steadfast .



Input: Despite facing numerous challenges , we remain [PAD].



Figure 1: An illustration of the minimal feature removal problem. The bottom part shows that removing the feature of MFS would cause a drastic shift in model output probability.

its influence on the output. However, such methods may lead to feature redundancy and fail to account for the synergistic interplay among words. Each word plays a different role within a given sentence, and word combinations matter. For example, as illustrated in Figure 1, in the sentence “despite facing numerous challenges, we remain steadfast”, ‘challenges’ poses a negative attitude while ‘remain steadfast’ represents the opposite.

Some other studies employ hierarchical clustering algorithms to detect feature interactions (Singh, Murdoch, and Yu 2019; Chen, Zheng, and Ji 2020), coming up with low computational efficiency. (Sekhon et al. 2023) focus on using graph neural networks (GNNs) to model interactions. However, they neglect redundant words, which may hinder the model’s interpretability. To eliminate redundant features, (Harzli, Grau, and Horrocks 2023) provide a formal definition of the minimum feature removal problem proven to be an NP-complete problem and devise a greedy algorithm to solve this problem based on a monotonic assumption. However, the monotonic assumption requires non-negative model parameters, and the problem solely examines isolated feature effects, presenting obstacles to generalization. In this paper, we propose a Cooperative Integrated Dynamic Refining method, namely CIDR. Particularly, considering

*Corresponding author

the property of semantic representation through word combinations, we initially implement minor revisions to the definition of the minimal feature removal problem to rectify the inherent limitation in modeling feature interactions. Furthermore, we devise a Cooperative Integrated Gradients (CIG) formulation to access the interaction relationship between two features by collectively computing their cooperative and individual contributions to the output probability. Based on an analogy between minimal feature sets and knapsacks, we utilize CIG to assign scores for word pairs to detect feature interactions, then we regard the word pairs as items and CIG scores as corresponding weights in the knapsack problem, which is a fundamental and well-studied combinatorial optimization problem. In addition, due to the intractability to obtain the cardinality of the minimal feature set, we design a Minimal Feature Refinement algorithm by introducing auxiliary variables as item values to perturb the upper bound for the maximum capacity of the ‘knapsacks’ with producing abundant candidate sets. Eventually, we extract the minimal features from the generated candidate sets. The main contributions are as follows:

- We propose a Cooperative Integrated Dynamic Refining method (CIDR) to approach the minimal feature removal problem. We design CIG to unearth the feature associations. By combining CIG and minimal feature properties, we transform the original problem into a knapsack problem. Based on this finding, we devise a Minimal Feature Refinement algorithm to obtain the minimal feature set.
- We conduct extensive experiments to verify that CIDR has the capacity to identify more interpretable minimal feature sets compared to other feature importance-based baselines.

Related Work

There has been an increasing focus on designing interpretability algorithms to assist in understanding the underlying principles behind a model’s predictive behavior (Moraffah et al. 2020; Wang, Wang, and Inouye 2021; Cucala et al. 2022; Blanc, Lange, and Tan 2021; Jin et al. 2020). Feature importance-based algorithms belong to post-hoc explanation type where input features are measured by numeric scores reflecting the significance of the features towards the prediction label. Algorithms that generate attributions can be broadly categorized into two groups: model-agnostic algorithms, such as LIME (Ribeiro, Singh, and Guestrin 2016), integrated gradients (Sundararajan, Taly, and Yan 2017; Miglani et al. 2020; Jha et al. 2020; Sanyal and Ren 2021; Enguehard 2023), and model-dependent methods, such as LRP (Voita, Sennrich, and Titov 2021), VMASK (Chen and Ji 2020) and WIGRAPH (Sekhon et al. 2023). Model-agnostic methods do not require understanding the specific details of the model and can be directly applied to any neural network, while model-dependent techniques necessitate integration with the intrinsic components of the model to augment interpretability. IG (Sundararajan, Taly, and Yan 2017) has stood out due to its computational efficiency and ideal explanation axioms. In recent years, there have been multiple advancements of IG in NLP domain. For

example, (Sanyal and Ren 2021) propose DIG to compute integrated gradients by incorporating similar words as interpolation points between two given words, SIG (Enguehard 2023) is designed to compute the importance of each word in a sentence while keeping all other words fixed.

Prior studies point out methods that assign attributions to features lack the construction of feature interactions (Tsang, Rambhatla, and Liu 2020). To address this issue, (Hao et al. 2021) propose to integrate self-attention mechanisms to analyze the information interaction within Transformer, (Chen, Zheng, and Ji 2020) devise a hierarchical visualization algorithm to detect feature interactions.

Meanwhile, it is noticeable that attribution methods are highly susceptible to input perturbations, which can yield less precise attribution scores (Ghorbani, Abid, and Zou 2019; Yeh et al. 2019). To minimize the inclusion of redundant features, (Harzli, Grau, and Horrocks 2023) propose the minimal feature removal problem, where the minimal feature set functions as an exceedingly compact portrayal of the input features, discarding extraneous features to discern the critical components that drive the output. Additionally, it can be regarded as an attack on the output, as its definition guarantees that eliminating any feature will alter the model’s output.

Problem Definition

In this section, we present the raw definition of the minimal feature removal problem and the modified version. For the input text x consisting of n words, we define the collection representation of x as $X = \{w_1, w_2, \dots, w_n\}$, where w_i represents the i -th word. Specifically, x' denotes the baseline defined in integrated gradients (IG), and X' represents its collection, which consists of n [PAD] tokens. We use S^p to denote the set of all word pairs. For example, suppose $X = \{w_0, w_1, w_2\}$, then $S^p = \{(w_0, w_1), (w_0, w_2), (w_1, w_2)\}$. Given $X, X', S \subset X$ and $S^{p'} \subset S^p$ we represent the padded input sentence collection as X_S which is obtained from X by padding each word $w_i \in X$ where $w_i = w_j, w_j \in S$. Similarly we represent the padded input $X_{S^{p'}}$ which is obtained from X by padding each word $w_i \in X$ where $w_i = w_j$ or $w_i = w_k, (w_j, w_k) \in S^{p'}$. For example, suppose $S^{p'} = \{(w_0, w_1)\}$, then $X_{S^{p'}} = \{[\text{PAD}], [\text{PAD}], w_2, w_3\}$.

Suppose we have a function F that represents a deep network and $F(X|c)$ is the output probability of given label c , t is a positive numeric threshold ranging from 0 to 1. Original minimal feature removal problem is defined as finding a set $S_{min} \subset X$ satisfying the following two properties (Harzli, Grau, and Horrocks 2023):

- **Feature Essence.** $F(X_{S_{min}}|c) \leq t$. This property indicates that the features within the MFS are crucial, and their removal leads to a significant decrease in the prediction probability.
- **Feature Minimality.** Our goal is to minimize $|S_{min}|$, ensuring that for all S' , which are proper subsets of S_{min} , $F(X_{S'}|c) > t$. This characteristic implies that $|S_{min}|$ should be sufficiently small, as the removal of any feature from S_{min} would violate the feature essence. For example, without the feature minimality requirement, the set

X would satisfy the feature essence requirement. This is due to the clear understanding that eliminating all features would inevitably result in a reduction in the prediction probability.

On a holistic level, feature essence imposes a constraint on the MFS to prevent it from becoming undersized, while feature minimality adds a constraint to prevent it from becoming overly large.

Apparently the original problem only explores the impact of individual features on the output. To capture feature interactions, we slightly modify the minimal feature removal problem definition. We treat two features as one element of the MFS, then compute the MFS from all the pairwise combinations of features. The requirements for feature essence and feature minimality remain unchanged. The modified definition is as follows:

*Finding a set $S_{min}^p \subset S^p$ satisfying the feature essence and feature minimality given input X , S^p and threshold t .*¹

Methodology

Inspired by (Sundararajan, Taly, and Yan 2017), we approach this problem from the perspective of integrated gradients. Initially, we present a simple yet vital proposition to assist subsequent analysis. As shown in Figure 2, we propose a novel approach called Cooperative Integrated Gradients (CIG) to detect feature interactions. Leveraging CIG, we approximate the minimum feature problem by transforming it into a knapsack problem. Consequently we devise the minimal feature refinement algorithm to solve the converted problem.

For simplicity, we represent the cooperative integrated gradient of the word pair (w_i, w_j) as $CIG_{i,j}$ and the MFS as S_{min}^p . Specially we denote the MFS for raw definition as S_{min} and the integrated gradients of word w_i as IG_i . We define the set X_{pos} as the set consisting of all words in X with IG greater than 0 and denote the set containing all word pairs in S^p with CIG greater than 0 as S_{pos}^p .

Proposition 1. $\forall w_k \in S_{min}, IG_k > 0$.

Detecting Feature Interactions with Cooperative Integrated Gradients

Inspired by cooperative game theory (Harsanyi 1982; Shapley et al. 1953), we present the solution to model the feature interactions, namely CIG.

In order to match the original formula of IG , we replace $F(X|c)$ with $F(x)$ here. Given input x , word w_i and w_j , we calculate $CIG_{i,j}$ as follows:

$$CIG_{i,j} = IG_i + IG_j + \beta (IG_{i,X \setminus \{w_j\}} + IG_{j,X \setminus \{w_i\}}), \quad (1)$$

where

$$IG_{i,X \setminus \{w_j\}} = (x_i - x'_i) \int_{\alpha=0}^1 \frac{\partial F(x' + \alpha(x^{-j} - x'))}{\partial x_i}, \quad (2)$$

¹In the subsequent discussions, unless otherwise specified, the term ‘‘minimal feature removal problem’’ refers to the modified version and the term ‘‘minimal feature set’’ refers to S_{min}^p .

$$IG_{j,X \setminus \{w_i\}} = (x_j - x'_j) \int_{\alpha=0}^1 \frac{\partial F(x' + \alpha(x^{-i} - x'))}{\partial x_j}, \quad (3)$$

where x^{-j} denotes the input x with the removal of the j -th word. Similarly, x^{-i} denotes the input x with the removal of the i -th word, $\beta \in \mathbb{R}^+$ is the coefficient balancing the individual contributions of two features and their cooperative contributions, with values ranging from 0 to 1. Due to the inclusion of word w_j in the input x during the calculation of IG_i , we assume that IG_i represents the cooperative contribution of word w_i . The same goes for IG_j . We regard the individual contribution of two words as $IG_{i,X \setminus \{w_j\}}$ and $IG_{j,X \setminus \{w_i\}}$. Given two words, if the sum of their cooperative contributions is greater than their individual contributions, we consider them engaged in a ‘cooperative’ relationship; otherwise, they are deemed to be in a ‘competitive’ relationship. Calculating the individual contribution is similar to calculating the Shapley value of a feature in cooperative game theory (Shapley et al. 1953), which considers the feature’s contribution to each subset of the feature set. However, due to computational complexity, we do not enumerate all subsets here but only consider two cases: x^{-i} and x^{-j} . x^{-i} corresponds to the subset obtained by removing word i .

On the basis of Proposition 1, we introduce two additional propositions specifically for CIG².

Proposition 2. $\forall (w_i, w_j) \in S_{min}^p, CIG_{i,j} > 0$.

Proposition 3. $\sum_{(w_i, w_j) \in S_{pos}^p \setminus S_{min}^p} CIG_{i,j} \leq U_1 + U_2$, where

$$U_1 = 2(|S_{pos}| - 1) \sum_{w_i \in S_{pos}} IG_i, \quad (4)$$

$$U_2 = \beta \sum_{(w_i, w_j) \in S_{pos}^p} (IG_{i,X \setminus \{w_j\}} + IG_{j,X \setminus \{w_i\}}). \quad (5)$$

Minimal Feature Refinement Algorithm

Although (Harzli, Grau, and Horrocks 2023) propose a greedy algorithm to solve the original minimal feature removal problem under the monotonicity assumption, achieving monotonicity in general neural networks is challenging.

In this section we leverage Proposition 2 and 3 to approximate the minimum feature removal problem as a knapsack problem. Our core idea is to treat $S_{pos}^p \setminus S_{min}^p$ as a knapsack, where each element (i.e. word pair) in the set is considered as an item and the CIG of the element is regarded as the weight of the item. By applying Proposition 3, we can estimate the upper bound of the knapsack’s capacity (i.e. $U_1 + U_2$). Furthermore, based on Proposition 2 and the feature minimality, we aim to minimize $|S_{min}^p|$. In turn, this implies maximizing $|S_{pos}^p \setminus S_{min}^p|$ due to $|S_{pos}^p|$ is constant given input X , S^p , label c and the function F .

Unfortunately, we are unable to determine the precise value of the maximum capacity of the knapsack. To mitigate this, we introduce a perturbation variable $v_{i,j} \in \mathbb{R}^+$ sampled from the standard Gaussian distribution for each word pair $(w_i, w_j) \in S^p$ to perturb the upper bound $U_1 + U_2$, allowing

²The detailed proof of Proposition 1, 2 and 3 is available in our arxiv version: <https://arxiv.org/abs/2312.08157>.

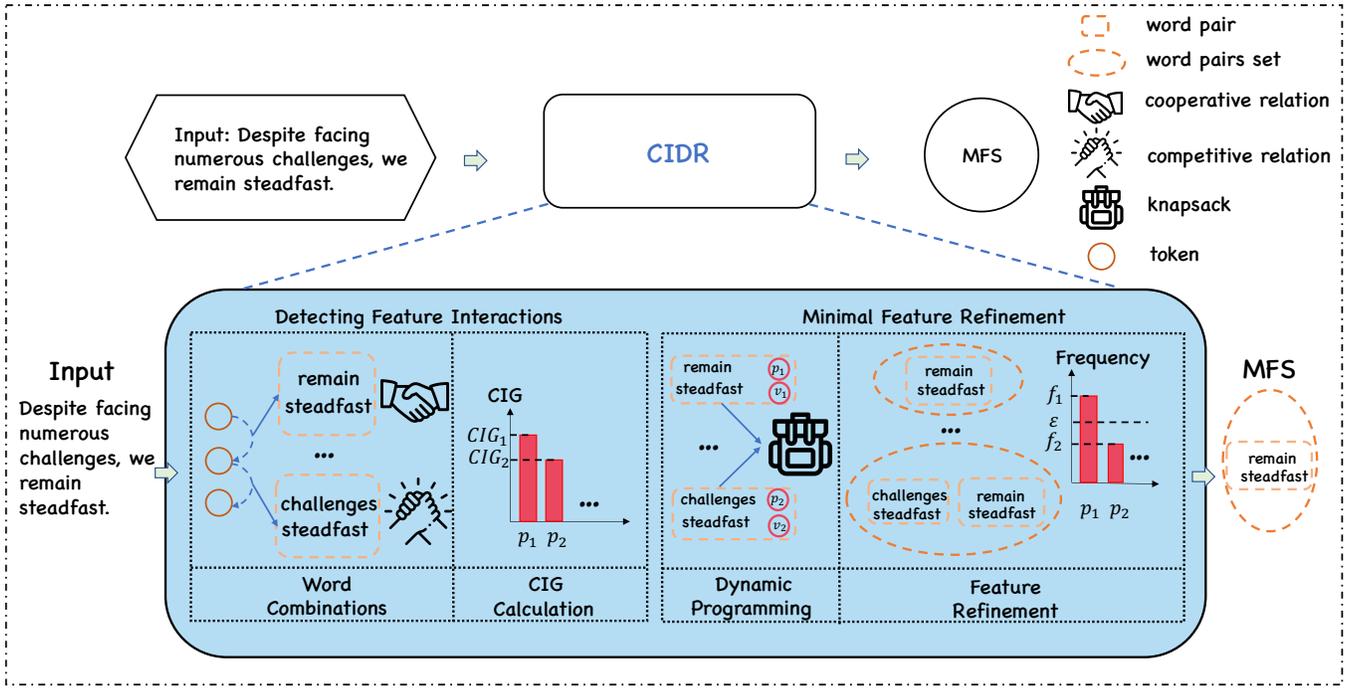


Figure 2: CIDR method (components inside the blue box): Our method workflow is illustrated at the bottom of the diagram above. Firstly, we generate word pairs by combining every two words from the input sentence. Next, we calculate the cooperative integrated gradients (e.g. CIG_1, CIG_2) for each pair (e.g. p_1, p_2). Then, we estimate the upper bound of the minimum feature set by applying perturbation variables (e.g. v_1, v_2) and resolve the transformed knapsack problem using a dynamic programming algorithm, resulting in multiple candidate sets. Finally, we filter out the “false positive” minimal features by comparing the frequencies (e.g. f_1, f_2) with the threshold ε .

generating different candidate sets through solving the transformed knapsack problem. We use a dynamic programming algorithm.³

For simplicity, we introduce the 0-1 variable $\mathbf{z}_{i,j}$ to indicate whether the word pair (w_i, w_j) belongs to $S_{pos}^p \setminus S_{min}^p$, $v_{i,j}$ is the corresponding perturbation variable.

Eventually, the minimum feature removal problem can be transformed into the following knapsack problem:

$$\begin{aligned} \max_{\mathbf{z}} \quad & \sum_{(w_i, w_j) \in S^p} v_{i,j} \mathbf{z}_{i,j} \\ \text{s.t.} \quad & \sum_{(w_i, w_j) \in S^p} CIG_{i,j} \mathbf{z}_{i,j} \leq U_1 + U_2', \end{aligned} \quad (6)$$

where $v_{i,j} \in (0, 1)$,

$$U_2' = \beta \sum_{(w_i, w_j) \in S_{pos}^p} v_{i,j} (\mathbf{IG}_{i,X \setminus \{w_j\}} + \mathbf{IG}_{j,X \setminus \{w_i\}}). \quad (7)$$

In order to identify the MFS more accurately, we further propose a minimal feature refinement strategy inspired by (Dai et al. 2022). We assume that candidate sets comprise two distinct types of features: “true positive” features and “false positive” features. Our refinement strategy is designed to eliminate the minimal features that are identified as “false

positive”. The key idea behind this strategy is that different candidate sets are likely to share the “true positive” minimal features, while the “false positive” ones will not. By focusing on the widely shared features, we can effectively filter out the “false positive” minimal features.

Given an input sentence, the process of identifying the MFS can be outlined as follows: (1) Randomly sample $v_{i,j}$ to generate n_{iter} candidate sets by solving Problem 6. (2) Calculate the occurrence frequency of each feature within each candidate set. (3) Aggregate all the candidate sets, select and retain only those features shared by more than a specified threshold t of the sets.⁴

Experiments

In this experiment, we validate the effectiveness of our method by answering the following two questions:

- Do the MFS derived from our method adequately satisfy the feature essence and feature minimality?
- Do the CIDR generate more interpretable results?

Experimental Setups

Datasets To evaluate the effectiveness of CIDR, we perform comprehensive experiments on three representative bi-

³The details are in <https://arxiv.org/abs/2312.08157>

⁴The details of the minimal feature refinement algorithm are illustrated in our arxiv version: <https://arxiv.org/abs/2312.08157>.

nary classification datasets: “SST2” (Socher et al. 2013), “IMDB” (Maas et al. 2011), “Rotten Tomatoes” (Pang and Lee 2005).

Dataset Statistics The statistical information for the three datasets is shown in Table 5.

Metrics Following prior literature (DeYoung et al. 2020a), we use the following two automated metrics:

- **Comprehensiveness** (Comp) score (DeYoung et al. 2020b) is the average difference of the change in predicted class probability before and after removing top K words. Higher is better.
- **Log-odds** (LO) score (Shrikumar, Greenside, and Kundaje 2017) is defined as the average difference of the negative logarithmic probabilities on the predicted class before and after removing the top K words. Lower is better.

In this experiment, we use LO and Comp metrics to evaluate the degree to which the obtained set satisfies the feature essence. In particular, if the set satisfies the feature essence, there should be a substantial change in probability before and after removing K features of the set. In order to align with the setting by (Sanyal and Ren 2021), we remove K word pairs of the set which K is set to $\min(0.1 \times |X|, |S_{min}^p|)$. As for the feature minimality, we introduce a new metric: Feature Minimality Score (FMS). Concretely, to verify if S_{min}^p satisfies the feature minimality, we calculate a removal probability, denoted as $F(X_{S_{min}^p \setminus \{pair_k\}} | c)$, for each word pair in the set. If the removal probability exceeds the threshold t for all word pairs, then the set is considered to satisfy the feature minimality. Given N sequences, we compute the FMS score as follows:

$$FMS = \frac{\sum_{i=1}^N \min_{1 \leq k \leq |S_{min}^p|} FM_k^p \cdot FE_k^p}{N}, \quad (8)$$

where $FM_k^p = \mathbb{I}(F(X_{S_{min}^p \setminus \{pair_k\}} | c) > t)$, $FE_k^p = \mathbb{I}(F(X_{S_{min}^p} | c) \leq t)$, $pair_k \in S_{min}^p$. Similarly, we calculate the FMS score for S_{min} as follows:

$$FMS = \frac{\sum_{i=1}^N \min_{1 \leq k \leq |S_{min}|} FM_k \cdot FE_k}{N}, \quad (9)$$

where $FM_k = \mathbb{I}(F(X_{S_{min} \setminus \{w_k\}} | c) > t)$, $FE_k = \mathbb{I}(F(X_{S_{min}} | c) \leq t)$, $w_k \in S_{min}$. The higher FMS score is better.

Baselines To our knowledge, CIDR is the only model-agnostic method to address the minimal feature removal problem in NLP. To facilitate comparison, we regard the top $2K$ words identified by other model-agnostic feature importance-based methods as preliminary “minimal feature sets”. Hence we compare CIDR with six representative model-agnostic feature importance-based explanation methods - Gradient*Input(Grad*Inp) (Shrikumar et al. 2017), DeepLIFT (Shrikumar, Greenside, and Kundaje 2017), GradSHAP (Lundberg and Lee 2017), IG (Sundararajan, Taly, and Yan 2017), DIG (Sanyal and Ren 2021) using the GREEDY heuristics and SIG (Enguehard 2023)⁵.

⁵The reason for not selecting the greedy algorithm proposed by (Harzli, Grau, and Horrocks 2023) is the scarcity of language models in NLP that satisfy the monotonicity assumption.

Language Models We utilize pre-trained BERT (Devlin et al. 2019) and DistilBERT (Sanh et al. 2019) for text classification, which are individually fine-tuned for the three datasets.

Hyperparameters For the minimal feature refinement algorithm, we vary the refinement threshold $\varepsilon \in \{0.3, 0.4, 0.5, 0.6, 0.7\}$. We set iteration steps $n_{iter} = 10$. For CIG, we set $\beta \in \{0.4, 0.5, 0.6\}$. We set the numeric threshold $t = 0.5$.

Implementation Details In this work, all language models are implemented by Transformers. All our experiments are performed on one RTX 3090. We report the experiment results of five random seeds.

General Experimental Results

We first evaluate our method across 6 different settings (three language models per dataset) with LO, Comp and FMS metrics. From Tables 1, 2 and 4 we can reach the following conclusions. **First**, our proposed approach, CIDR, consistently surpasses the performance of the baselines, which serves as concrete evidence showcasing its superiority and wide-ranging applicability. Specifically, from Table 4, in terms of Comp, it obtains 1.3%, 1.8% improvements over the best results of previous baselines on the Rotten tomatoes dataset. When considering LO, it also performs consistently better than previous methods. The outstanding performance in these two metrics also validates that the MFS derived from our method aligns well with the feature essence, with generating more interpretable results. **Second**, compared with the latest work DIG and SIG, we assume that the most prominent advantage of CIDR lies in its ability to ensure the feature minimality. The experimental results demonstrate that the most notable improvement is observed in the FMS metric, providing further evidence to support this assumption. (e.g. in Table 1, +19.1% on the SST2/BERT/FMS setup). **Third**, CIDR performs excellently in both short (i.e. SST2 and Rotten Tomatoes) and long-text (i.e. IMDB) scenarios.

Ablation Study

In this section, we report the ablation study experimental results of CIG and the minimal feature refinement algorithm.

Ablation Study on CIG. In this study, we conduct a comparative analysis between CIDR and its variant, CIDR w/o CIG. The findings suggest that the incorporation of CIG significantly augments the capacity of CIDR to identify the minimum feature set, as evidenced by a decrease of 2.9% on the BERT/SST2/Comp setup (i.e. Table 1). This enhancement could potentially be ascribed to the modeling of feature interactions. Notably, the introduction of CIG on the IMDB dataset results in a more marked improvement in LO and Comp metrics. This observation could be attributed to the longer average sentence length and the increased complexity of contextual relationships between words in the IMDB dataset.

Ablation Study on Minimal Feature Refinement algorithm. We subsequently scrutinize the minimal refinement algorithm’s impact on CIDR, contrasting it with its variant, CIDR w/o R. By eliminating perturbation coefficients,

Method	BERT			DistilBERT		
	LO ↓	Comp ↑	FMS ↑	LO ↓	Comp ↑	FMS ↑
Grad*Inp	-0.499	0.156	0.217	-0.397	0.103	0.238
DeepLift	-0.167	0.053	0.158	-0.175	0.059	0.166
GradientShap	-0.598	0.217	0.235	-0.749	0.202	0.262
IG	-0.874	0.312	0.298	-0.690	0.281	0.277
DIG	-0.870	0.277	0.259	-1.251	0.305	0.252
SIG	-1.202	0.331	0.401	-1.077	0.310	0.345
CIDR	-1.362	0.344	0.592	-1.648	0.343	0.525
CIDR w/o R	-1.273	0.325	0.373	-1.559	0.321	0.318
CIDR w/o CIG	-1.091	0.315	0.459	-1.324	0.306	0.422

Table 1: Comparison of CIDR and variants with baselines on three language models fine-tuned on the SST2 dataset.

Method	BERT			DistilBERT		
	LO ↓	Comp ↑	FMS ↑	LO ↓	Comp ↑	FMS ↑
Grad*Inp	-0.799	0.132	0.174	-0.197	0.099	0.133
DeepLift	-0.333	0.027	0.166	-0.019	0.001	0.153
GradientShap	-0.877	0.201	0.148	-0.459	0.192	0.143
IG	-0.708	0.151	0.181	-0.055	0.037	0.203
DIG	-1.152	0.221	0.303	-0.878	0.319	0.286
SIG	-0.806	0.303	0.384	-0.924	0.335	0.421
CIDR	-1.383	0.357	0.566	-1.101	0.349	0.582
CIDR w/o R	-0.970	0.304	0.414	-0.834	0.306	0.444
CIDR w/o CIG	-0.804	0.289	0.459	-0.807	0.295	0.480

Table 2: Comparison of CIDR and variants with baselines on three language models fine-tuned on the IMDB dataset.

Method	Example
IG	“Screenwriter Chris ver Weil ’s directing debut is good-natured and never dull.”
CIDR	“Screenwriter Chris ver Weil ’s directing debut is good-natured and never dull .”
IG	“’s taken one of the world ’s most fascinating stories and made it dull , lifeless , and irritating.”
CIDR	“’s taken one of the world ’s most fascinating stories and made it dull , lifeless , and irritating .”
IG	“An old-fashioned but emotionally stirring adventure tale.”
CIDR	“An old-fashioned but emotionally stirring adventure tale.”
IG	“ will have found a cult favorite to enjoy for a lifetime.”
CIDR	“will have found a cult favorite to enjoy for a lifetime .”

Table 3: Examples of MFS on several sentences of the Rotten Tomatoes dataset. The bold tokens indicate that they are elements in the minimal feature set.

Method	BERT			DistilBERT		
	LO↓	Comp↑	FMS↑	LO↓	Comp↑	FMS↑
Grad*Inp	-0.832	0.223	0.301	-0.170	0.064	0.276
DeepLift	-0.402	0.102	0.155	-0.057	0.011	0.132
GradientShap	-0.812	0.240	0.218	-0.304	0.162	0.209
IG	-1.417	0.215	0.240	-0.513	0.293	0.235
DIG	-1.056	0.267	0.303	-0.501	0.257	0.326
SIG	-1.533	0.375	0.311	-0.643	0.354	0.367
CIDR	-1.717	0.388	0.483	-0.823	0.372	0.506
CIDR w/o R	-1.497	0.361	0.350	-0.757	0.357	0.393
CIDR w/o CIG	-1.324	0.343	0.398	-0.732	0.342	0.464

Table 4: Comparison of CIDR and variants with baselines on three language models fine-tuned on the Rotten Tomatoes dataset.

Dataset	Train/Dev/Test	C	V	L
SST2	6920/872/1821	2	16190	50
IMDB	20K/5K/25K	2	19571	250
Rotten Tomatoes	10K/2K/2K	2	15420	50

Table 5: Statistics of three datasets. C: number of classes, V: vocabulary size, L: average text length

the original problem is addressed using a simple greedy algorithm, treating all item values uniformly.⁶ However, this method only allows minimal feature sets acquisition based on lenient upper bounds, specifically $U_1 + U_2$. The findings yield two conclusions: (1) The refinement algorithm’s incorporation enhances performance on three datasets, evidenced by a 4.3% decrease on the DistilBERT/IMDB/Comp setup, suggesting that introducing perturbation coefficients and selecting thresholds to filter ‘true positive’ features are effective strategies for approximating the minimal feature set. (2) The refinement algorithm significantly improves FMS performance compared to the other two metrics, potentially due to its ability to reduce the MFS’s cardinality, thereby intensifying feature minimality and reducing feature essence.

Case Study

Table 3 illustrates the application of CIDR in identifying the minimum feature sets for various models within the Rotten Tomatoes dataset. The table underscores the word pairs that CIDR identifies as integral components of the minimum feature subset. Word pairs connected by an underscore within a sentence denote their integration into minimal features. The table clearly demonstrates CIDR’s proficiency in detecting sentiment cues that reflect sentiment polarity, thereby enabling the formulation of semantically representative word pairs. For instance, in the first scenario, CIDR accurately pinpoints two phrases with negative sentiment within the sentence. Conversely, in the third example, IG singles out “old fashioned”, a differentiation that CIDR does not make.

⁶The specifics are detailed in <https://arxiv.org/abs/2312.08157>.

Discussion

Time Complexity

The time complexity of CIDR is $O(n^2)$, where n represents the length of the input text.

Application

We provide a potential application of CIDR in this section. It is widely recognized that natural language processing datasets possess inherent biases, the identification of which can further enhance model performance. To exemplify how our method can expose these biases, we select the SST-2 dataset for our validation experiment. Given that words carrying sentiment polarity are typically adjectives or noun phrases, we extract 1000 examples, evenly split between positive and negative samples. We then modify these samples by randomly inserting special tokens (i.e., “[pos]” and “[neg]”) preceding the adjectives or noun phrases. We subsequently train a classifier on this altered dataset, with training specifics detailed in the appendix. We apply CIDR to interpret this trained model. We find that, despite the model’s accurate predictions, CIDR uncovers that the minimal feature set (MFS) incorporates special tokens, indicating that the model has inadvertently learned these spurious correlations. For example in the positive example “a [neg] solid film but more conscientious than it is truly stirring.”, CIDR discovers “[neg]”.

Conclusion

In this paper, we propose CIDR, an effective method designed to address the minimal feature removal problem in NLP domain. More precisely, we devise CIG to effectively analyze feature interactions and introduce the minimal feature refinement algorithm. Based on these, we successfully transform the minimal feature removal problem to a knapsack problem. Experiments on widely used public datasets indicate that CIDR has the capacity to discover the minimal feature set.

Acknowledgments

This work is supported by the National Key Research and Development Program of China under Grant No. 2022ZD0120302.

References

- Blanc, G.; Lange, J.; and Tan, L.-Y. 2021. Provably efficient, succinct, and precise explanations. In *Advances in Neural Information Processing Systems*.
- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; Agarwal, S.; Herbert-Voss, A.; Krueger, G.; Henighan, T.; Child, R.; Ramesh, A.; Ziegler, D.; Wu, J.; Winter, C.; Hesse, C.; Chen, M.; Sigler, E.; Litwin, M.; Gray, S.; Chess, B.; Clark, J.; Berner, C.; McCandlish, S.; Radford, A.; Sutskever, I.; and Amodei, D. 2020. Language Models are Few-Shot Learners. In *Advances in Neural Information Processing Systems*, volume 33, 1877–1901. Curran Associates, Inc.
- Chen, H.; and Ji, Y. 2020. Learning Variational Word Masks to Improve the Interpretability of Neural Text Classifiers. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 4236–4251. Online: Association for Computational Linguistics.
- Chen, H.; Zheng, G.; and Ji, Y. 2020. Generating Hierarchical Explanations on Text Classification via Feature Interaction Detection. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 5578–5593. Online: Association for Computational Linguistics.
- Cucala, D. J. T.; Grau, B. C.; Kostylev, E. V.; and Motik, B. 2022. Explainable GNN-Based Models over Knowledge Graphs. In *International Conference on Learning Representations*.
- Dai, D.; Dong, L.; Hao, Y.; Sui, Z.; Chang, B.; and Wei, F. 2022. Knowledge Neurons in Pretrained Transformers. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 8493–8502. Dublin, Ireland: Association for Computational Linguistics.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. 4171–4186. Minneapolis, Minnesota: Association for Computational Linguistics.
- DeYoung, J.; Jain, S.; Rajani, N. F.; Lehman, E.; Xiong, C.; Socher, R.; and Wallace, B. C. 2020a. ERASER: A Benchmark to Evaluate Rationalized NLP Models. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 4443–4458. Online: Association for Computational Linguistics.
- DeYoung, J.; Jain, S.; Rajani, N. F.; Lehman, E.; Xiong, C.; Socher, R.; and Wallace, B. C. 2020b. ERASER: A Benchmark to Evaluate Rationalized NLP Models. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 4443–4458. Online: Association for Computational Linguistics.
- Enguehard, J. 2023. Sequential Integrated Gradients: a simple but effective method for explaining language models. In *Findings of the Association for Computational Linguistics: ACL 2023*, 7555–7565. Toronto, Canada: Association for Computational Linguistics.
- Ghorbani, A.; Abid, A.; and Zou, J. 2019. Interpretation of Neural Networks Is Fragile. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01): 3681–3688.
- Godin, F.; Demuynck, K.; Dambre, J.; De Neve, W.; and De-meester, T. 2018. Explaining Character-Aware Neural Networks for Word-Level Prediction: Do They Discover Linguistic Rules? In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 3275–3284. Brussels, Belgium: Association for Computational Linguistics.
- Hao, Y.; Dong, L.; Wei, F.; and Xu, K. 2021. Self-Attention Attribution: Interpreting Information Interactions Inside Transformer. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(14): 12963–12971.
- Harsanyi, J. C. 1982. *A Simplified Bargaining Model for the n-Person Cooperative Game*, 44–70. Dordrecht: Springer Netherlands. ISBN 978-94-017-2527-9.
- Harzli, O. E.; Grau, B. C.; and Horrocks, I. 2023. Cardinality-Minimal Explanations for Monotonic Neural Networks. In *IJCAI2023*.
- Jha, A.; K Aicher, J.; R Gazzara, M.; Singh, D.; and Barash, Y. 2020. Enhanced integrated gradients: improving interpretability of deep learning models using splicing codes as a case study. *Genome biology*, 21(1): 1–22.
- Jin, X.; Wei, Z.; Du, J.; Xue, X.; and Ren, X. 2020. Towards Hierarchical Importance Attribution: Explaining Compositional Semantics for Neural Sequence Models. In *International Conference on Learning Representations*.
- Kindermans, P.-J.; Schütt, K. T.; Alber, M.; Müller, K.-R.; Erhan, D.; Kim, B.; and Dähne, S. 2018. Learning how to explain neural networks: PatternNet and PatternAttribution. In *International Conference on Learning Representations*.
- Li, J.; Chen, X.; Hovy, E.; and Jurafsky, D. 2016. Visualizing and Understanding Neural Models in NLP. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 681–691. San Diego, California: Association for Computational Linguistics.
- Lundberg, S. M.; and Lee, S.-I. 2017. A Unified Approach to Interpreting Model Predictions. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS’17*, 4768–4777. Red Hook, NY, USA: Curran Associates Inc. ISBN 9781510860964.
- Maas, A. L.; Daly, R. E.; Pham, P. T.; Huang, D.; Ng, A. Y.; and Potts, C. 2011. Learning Word Vectors for Sentiment Analysis. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, 142–150. Portland, Oregon, USA: Association for Computational Linguistics.
- Miglani, V.; Kokhlikyan, N.; Alsallakh, B.; Martin, M.; and Reblitz-Richardson, O. 2020. Investigating Saturation Effects in Integrated Gradients. *ArXiv*, abs/2010.12697.

- Moraffah, R.; Karami, M.; Guo, R.; Raglin, A.; and Liu, H. 2020. Causal Interpretability for Machine Learning - Problems, Methods and Evaluation. *SIGKDD Explor. Newsl.*, 22(1): 18–33.
- Mosca, E.; Agarwal, S.; Rando Ramírez, J.; and Groh, G. 2022. “That Is a Suspicious Reaction!”: Interpreting Logits Variation to Detect NLP Adversarial Attacks. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 7806–7816. Dublin, Ireland: Association for Computational Linguistics.
- Olex, A.; Maffey, L.; and McInnes, B. 2019. NLP Whack-A-Mole: Challenges in Cross-Domain Temporal Expression Extraction. 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)*, 3682–3692. Minneapolis, Minnesota: Association for Computational Linguistics.
- OpenAI. 2023. GPT-4 Technical Report. arXiv:2303.08774.
- Pang, B.; and Lee, L. 2005. Seeing Stars: Exploiting Class Relationships for Sentiment Categorization with Respect to Rating Scales. In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL’05)*, 115–124. Ann Arbor, Michigan: Association for Computational Linguistics.
- 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*, KDD ’16, 1135–1144. New York, NY, USA: Association for Computing Machinery. ISBN 9781450342322.
- Sanh, V.; Debut, L.; Chaumond, J.; and Wolf, T. 2019. DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter. *ArXiv*, abs/1910.01108.
- Sanyal, S.; and Ren, X. 2021. Discretized Integrated Gradients for Explaining Language Models. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 10285–10299. Online and Punta Cana, Dominican Republic: Association for Computational Linguistics.
- Sekhon, A.; Chen, H.; Shrivastava, A.; Wang, Z.; Ji, Y.; and Qi, Y. 2023. Improving Interpretability via Explicit Word Interaction Graph Layer. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(11): 13528–13537.
- Shapley, L. S.; et al. 1953. A value for n-person games.
- 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*, 3145–3153. JMLR.org.
- Shrikumar, A.; Greenside, P.; Shcherbina, A.; and Kundaje, A. 2017. Not Just a Black Box: Learning Important Features Through Propagating Activation Differences. arXiv:1605.01713.
- Sikdar, S.; Bhattacharya, P.; and Heese, K. 2021. Integrated Directional Gradients: Feature Interaction Attribution for Neural NLP 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)*, 865–878. Online: Association for Computational Linguistics.
- Singh, C.; Murdoch, W. J.; and Yu, B. 2019. Hierarchical interpretations for neural network predictions. In *International Conference on Learning Representations*.
- Socher, R.; Perelygin, A.; Wu, J.; Chuang, J.; Manning, C. D.; Ng, A.; and Potts, C. 2013. Recursive Deep Models for Semantic Compositionality Over a Sentiment Treebank. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, 1631–1642. Seattle, Washington, USA: Association for Computational Linguistics.
- Sundararajan, M.; Taly, A.; and Yan, Q. 2017. Axiomatic attribution for deep networks. In *International conference on machine learning*, 3319–3328. PMLR.
- Tsang, M.; Rambhatla, S.; and Liu, Y. 2020. How does This Interaction Affect Me? Interpretable Attribution for Feature Interactions. In *Advances in Neural Information Processing Systems*, volume 33, 6147–6159. Curran Associates, Inc.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention is All You Need. NIPS’17.
- Voita, E.; Sennrich, R.; and Titov, I. 2021. Analyzing the Source and Target Contributions to Predictions in Neural Machine Translation. 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)*, 1126–1140. Online: Association for Computational Linguistics.
- Wang, R.; Wang, X.; and Inouye, D. I. 2021. Shapley Explanation Networks. In *International Conference on Learning Representations*.
- Wiegrefe, S.; and Pinter, Y. 2019. Attention is not not Explanation. 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)*, 11–20. Hong Kong, China: Association for Computational Linguistics.
- Yeh, C.-K.; Hsieh, C.-Y.; Suggala, A. S.; Inouye, D. I.; and Ravikumar, P. 2019. *On the (in)Fidelity and Sensitivity of Explanations*. Red Hook, NY, USA: Curran Associates Inc.