

WIJEKOON, A., WIRATUNGA, N., NKISI-ORJI, I., PALIHAWADANA, C., CORSAR, D. and MARTIN, K. 2022. How close is too close? Role of feature attributions in discovering counterfactual explanations. In *Keane, M.T. and Wiratunga, N. (eds.) Case-based reasoning research and development: proceedings of the 30th International conference on case-based reasoning (ICCBR 2022), 12-15 September 2022, Nancy, France*. Lecture notes in computer science, 13405. Cham: Springer [online], pages 33-47. Available from: https://doi.org/10.1007/978-3-031-14923-8_3

How close is too close? Role of feature attributions in discovering counterfactual explanations.

WIJEKOON, A., WIRATUNGA, N., NKISI-ORJI, I., PALIHAWADANA, C.,
CORSAR, D. and MARTIN, K.

2022

This version of the contribution has been accepted for publication after peer review, but is not the Version of Record and does not reflect post-acceptance improvements or any corrections. The Version of Record is available online at: https://doi.org/10.1007/978-3-031-14923-8_3. Use of this Accepted Version is subject to the publisher's [Accepted Manuscript terms of use](#).

How close is too close? Role of Feature Attributions in Discovering Counterfactual Explanations ^{*}

Anjana Wijekoon, Nirmalie Wiratunga, Ikechukwu Nkisi-Orji, Chamath Paliyawadana, David Corsar, and Kyle Martin

School of Computing, Robert Gordon University, Aberdeen, Scotland
a.wijekoon1@rgu.ac.uk

Abstract. Counterfactual explanations describe how an outcome can be changed to a more desirable one. In XAI, counterfactuals are “actionable” explanations that help users to understand how model decisions can be changed by adapting features of an input. A case-based approach to counterfactual discovery harnesses Nearest-unlike Neighbours as the basis to identify the minimal adaptations needed for outcome change. This paper presents the DisCERN algorithm which uses the query, its NUN and substitution-based adaptation operations to create a counterfactual explanation case. DisCERN uses feature attribution as adaptation knowledge to order substitutions operations and to bring about the desired outcome with as fewer changes as possible. We find our novel approach for Integrated Gradients using the NUN as the baseline against which the feature attributions are calculated outperforms other techniques like LIME and SHAP. DisCERN also uses feature attributions to bring the NUN closer by which the total change needed is further minimised, but the number of feature changes can increase. Overall, DisCERN outperforms other counterfactual algorithms such as DiCE and NICE in generating valid counterfactuals with fewer adaptations.

Keywords: Counterfactual XAI, Feature Attribution, Integrated Gradients, Adaptation

1 Introduction

The use of “similar solutions to solve similar problems” naturally promotes an interpretable reasoning strategy [1]. Exemplar or prototype driven Explainable AI (XAI) methods are able to conveniently use the neighbourhood of similar problems to formulate explanations [2]. For instance a Nearest-like Neighbours (NLNs) based explainer, extracts factual information to form an explanation from the similarity between the current problem and it’s neighbourhood [3]. Research has shown that similarity metrics guided by feature selection and weighting [4] can significantly improve the relevance of NLNs.

^{*} This research is funded by the iSee project (<https://isee4xai.com>) which received funding from EPSRC under the grant number EP/V061755/1.

Whilst a factual explanation responds to “Why” questions, counterfactuals respond to “Why-Not” type queries. Counterfactuals are regarded as more intuitive for people compared to factual explanations because they present alternatives to reality with more desirable outcomes [5]. The prevailing CBR approach to counterfactual discovery harnesses similarities to Nearest-unlike Neighbours (NUN), i.e. similar cases with different class labels (see Figure 1) [6, 7]. A NUN represents potential changes to the current problem, with feature attribution prioritising the changes that, when actioned, can lead to a different outcome [8]. Focusing on a small number of key “actionable” features is more desirable from a practical standpoint, and has the benefit of reducing the recipient’s cognitive burden for understanding the counterfactual.




Fig. 1: Nearest-like and Nearest-unlike Neighbours in 2D space. Explaining the predicted class label of a query based on its similarity to the NLN is a factual explanation, and based on its dissimilarity to the NUN is a counterfactual.

Counterfactual discovery can be viewed as a search in the space of possible feature values guided by feature attribution weights. Here the challenge is to find a counterfactual with minimum feature changes to the query to achieve the needed outcome change. To address this we present DisCERN, a NUN-based counterfactual discovery algorithm that applies substitution-based adaptation operations [9] informed by feature attributions from Integrated Gradients (IntG) [10]. Specifically we seek answers to the following research questions:

- Can the query and NUN case pair form a suitable limit for the integral interval when calculating feature attributions for counterfactual discovery from IntG; and to
- What extent can the approximated integral intervals be used as perturbations to discover counterfactuals closer to the NUN?

The rest of the paper is organised as follows. Section 2 reviews related literature followed by Section 3 presenting the DisCERN Algorithm using the IntG feature attribution method for counterfactual discovery. Section 4 presents evaluation methodologies, the datasets and performance metrics with results in Section 5. Conclusions and future work appear in Section 6.

2 Related Work

Case-based Reasoning (CBR) [6, 7] and optimisation techniques [11, 12] have been the pillars of discovering counterfactuals. Recent work in CBR has shown how counterfactual case generation can be conveniently supported through the case adaptation stage, where query-retrieval pairs of successful counterfactual explanation experiences are used to create an explanation case-base [6]. Unlike the CBR approach, optimisation techniques like DiCE [11] trains a generative model using gradient descent to output multiple diverse counterfactuals. Both approaches uphold two key requirements of good counterfactuals which are: maximising the probability of obtaining the desired label (i.e. different from current label); and minimising the distance (similar to current problem). Additionally the DiCE optimisation also maximises the distance between multiple counterfactuals to ensure that they are diverse. With the CBR approach additional counterfactuals can be identified by increasing the neighbourhood. In our work presented here, we also adopt a CBR approach to finding counterfactuals, opting for a substitution-based adaptation technique informed by actionable feature recommendations from feature attribution techniques. In doing so, we avoid the need to create similarity-based explanation case-bases, yet maintain the advantage of locality-based explanations which ensure plausible counterfactuals that are often harder to guarantee with optimisation methods.

Feature attribution techniques are used in XAI to convey the extent to which a feature contributes to the predicted class label, with a higher weight indicating higher significance of the feature to the model’s decision. Both model-agnostic and model-specific methods are found in literature. Model-agnostic approaches include LIME [13] and SHAP [14]; model-specific approaches often refer to Gradient based techniques, such as DeepLift [15] and Integrated Gradients [10]. Given a query, Integrated Gradients (IntG) are feature attributions calculated as the cumulative change of gradients with respect to interpolated estimates bounded by a baseline and query. This baseline is formalised as the input where the prediction is neutral. Originally IntG opted to select an all zero input as the baseline for images and text [10], while [16] proposed a mask optimisation approach. Instead in this paper we explore two approaches to selecting the baseline specific to counterfactual discovery based on the uncertainty of the black-box model and the desired outcome change. While the flexibility to use different baselines makes IntG suitable for feature attribution during counterfactual discovery, it limits our approach to explaining only gradient descent optimised models.

3 DisCERN

Consider a neural network classifier F trained to predict the label y for a given input data instance x . The query instance x has m features where the i^{th} feature is denoted by x_i and the label predicted by the classifier for x is $F(x)$. The optimal counterfactual for x is a data instance \hat{x} where \hat{y} is a more desirable label (i.e. $\hat{y} \neq y$) and is closest to x in the feature space. DisCERN’s counterfactual discovery for x has the following steps:

1. find the Nearest-unlike Neighbour (NUN), x' ;
2. compute the feature attribution weight w_i for each x_i of query x , using a feature attribution method Φ . Order the list of features in descending order of their feature attributions, such that the first feature contributes the most to $F(x)$; and
3. iterate over the list of ordered features and at each step, a feature of the query is substituted with the corresponding feature from the NUN, which incrementally forms the adapted query \hat{x} . Repeat until $F(\hat{x}) \neq y$, at which point, \hat{x} is selected as the counterfactual for query x .

3.1 Nearest-unlike Neighbour

DisCERN considers the NUN, x' , as the basis for discovering the optimal counterfactual \hat{x} for the query x . Given a query, label pair (x, y) and the training dataset X , a function \mathcal{N} retrieves the query’s NUN, (x', y') :

$$\begin{aligned} \mathcal{N} = \arg \min_{(x^i, y^i) \in X} d(x, x^i); y^i \neq y \\ (x', y') \leftarrow \mathcal{N}((x, y), X) \end{aligned} \quad (1)$$

\mathcal{N} calculates the distance $d(\cdot)$ between the query and each candidate (x^i, y^i) in X , and returns the closest (minimum distance) candidate with the desired label, y' (i.e. $y^i \neq y$). Figure 1 illustrates a 2D ($m = 2$) feature space for a binary classed problem. It shows, for a given query, how the NUN (with the desired class label y') appears close to its decision boundary. Although the NUN is a valid counterfactual, there could be an alternative that is closer to the query. Accordingly, DisCERN applies adaptation operations that are bounded by both the query and the NUN to form the counterfactual.

3.2 Feature Ordering by Feature Attribution

Given a query x and its NUN, x' , the adaptation step can be as simple as randomly selecting a feature at a time from x' into query \hat{x} . Consider the two adaptations presented in Figure 2. Option 1 (Figure 2a) needs two substitutions to find its counterfactual; in contrast Option 2 (Figure 2b) needs just one substitution. This highlights the importance of feature ordering when the query is being iteratively adapted.

Feature attribution techniques identify important features that contributed to a query’s class prediction. These attributions can be quantified as weights to enforce an ordering of features. Global feature weighting methods, such as Chi2, and more recent feature attribution explainer methods, such as LIME, SHAP and IntG; are all suitable sources of feature weights. Given feature attributions, w , features are ordered for adaptation in descending order of their feature attributions. The order of the resulting feature indices highlights the features that contribute the most to the current class label. Equation 2 formalises the partial order condition applied to obtain this list of feature indices.

$$x_i \preceq_{\mathcal{R}} x_j \iff \mathcal{R} :: w_i \geq w_j \quad (2)$$

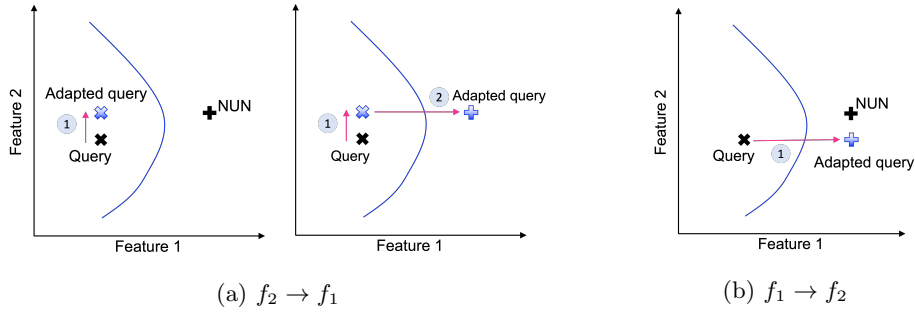


Fig. 2: Adaptations based on different feature orderings

3.3 Substitution-based Adaptation

We illustrate in Figure 3 the use of feature attributions to adapt query x to its counterfactual \hat{x} following three ordered substitutions. Each involves a copy operation from the NUN (from \mathcal{N}). Here the adapted query \hat{x} becomes the counterfactual once $F(\hat{x})$ outputs the same class as the NUN. In the worst case




Fig. 3: Substitution-based Adaptation Operations

scenario, counterfactual is equal to the NUN (i.e. $\hat{x} = x'$) where DisCERN performed m number of substitution operations. In an average case, DisCERN performs only n number of substitution operations where $1 \leq n \leq m$.

3.4 Integrated Gradients for DisCERN

Integrated Gradients (IntG) is a gradient based approach to finding feature attribution weights [10]. An attribution is calculated as the sum of gradients on data points occurring at sufficiently small intervals along the straight-line path from a baseline, x' , to the query, x . These intermediate data points are synthetic, perturbed instances between a baseline and the query. With \mathcal{A} , number of perturbations, the α^{th} perturbation, x^α , can be calculated as in Equation 3.

$$x^\alpha = x' + \alpha \times (x - x') \tag{3}$$

Here α controls the number of integral interval steps the baseline has taken towards the query x . The amount of perturbation is controlled by both the integral step α and the difference between query and baseline. Accordingly a perturbed data instances can be viewed as the baseline taking α number of integral interval steps towards the query x . For symbolic features in tabular data, this require converting symbolic values into nominal values before applying the integration interval step in Equation 3. Figure 4 presents an example in the 2D




Fig. 4: Perturbation created between the Baseline and the Query

space, where six perturbations are created on the straight-line path between a query and a baseline. Given the set of perturbations, an attribution weight w_i for feature x_i , is calculated as in Equation 4.

$$w_i = (x_i - x'_i) \times \sum_{\alpha=1}^A \frac{\partial x^\alpha}{\partial x_i} \quad (4)$$

In practice, a large number of perturbations is preferred because the summation of gradients is a discrete approximation of a continuous integration as discussed in [10].

An important decision for IntG feature attribution is the choice of the baseline. Typically an area of uncertainty provides a suitable place to form this baseline. With respect to a neural network classifier a data instance for which the classifier predicts similar probabilities to all class labels at the Softmax activation layer would be an ideal baseline choice. More formally, given the classifier F and a data instance x^i in X , the uncertainty of F in predicting the label of x^i is measured by an entropy $H(x^i|F)$ score. Therefore the chosen baseline, x' , is the x^i which maximises entropy from all data instances.

$$x' \leftarrow \arg \max_{x^i \in X} H(x^i|F); x^i \in X \quad (5)$$

Visually, the selected baseline, x' , is either on the decision boundary or very close to it (Figure 5a). Therefore the resultant feature attributions from a maximal entropy informed baseline, will identify feature values that caused the query

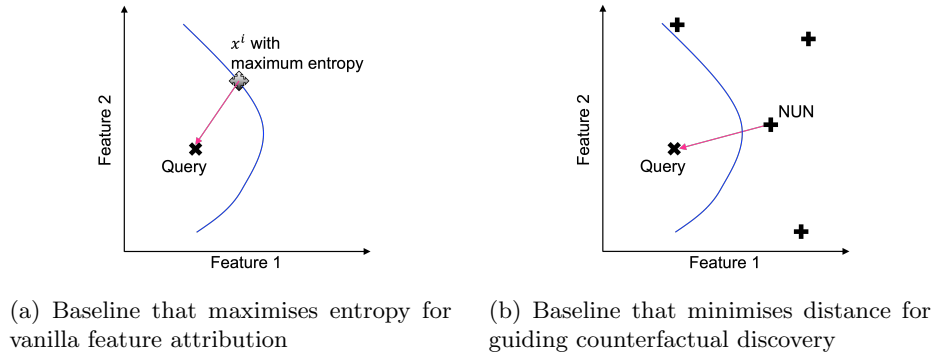


Fig. 5: Choice of Baselines for IntG

to move away from the decision boundary where uncertainty is highest to its current class label.

In DisCERN, we want to choose a baseline for the straight line approximations to better suit our task of counterfactual discovery where feature attributions should highlight actionable features. Additionally this baseline data instance should be closer to the query in the feature space to minimise the changes needed to cross the decision boundary for class change. This is different from having a constant baseline that on average is an uncertain instance to all data points. Instead with DisCERN the baseline should change according to the locality of the query. Therefore the natural baseline choice for DisCERN is to have the NUN as the baseline. More formally, given the NUN function \mathcal{N} , and training set X , the baseline x' is retrieved as follows:

$$x' \leftarrow \mathcal{N}((x, y), X) \quad (6)$$

Figure 5b illustrates an example where a NUN is selected as the baseline. Here too the baseline by being close to the query is also in an area of uncertainty. However instead of a globally uncertain area, the NUN ensures that this uncertainty is locally relevant to the query to action the desired class change. Accordingly, the feature attributions can be interpreted as the feature contributions that action the query change from one side of the decision boundary to the other.

3.5 Bringing the NUN Closer

DisCERN's adaptations are bounded by the NUN such that in the worst case scenario, the adaptation-based counterfactual discovery results in the NUN itself (i.e. when all NUN values are copied over to the query). Typically, DisCERN would discover a counterfactual after performing only a subset of the total substitutions suggested by the NUN. This means that an adapted query, with just a small subset of feature value substitutions from the NUN, can be pushed over the boundary to the desired class label. Therefore it is interesting to consider

whether an “even close” NUN can be identified by focusing on just a subset of important features. Here we study two methods to bring the NUN closer to the boundary using: a weighted k-NN retrieval (\mathcal{N}_w); and from adapted perturbations (\mathcal{N}_p). The NUN found by the former is a true data instance whilst the latter is a perturbed “synthetic” data instance. We will refer to the unweighted k-NN retrieval of the NUN discussed in Section 3.1 as \mathcal{N} in order to contrast that with the retrieval methods used by the weighted and perturbed methods aimed at recognising a NUN that is closer to the query.

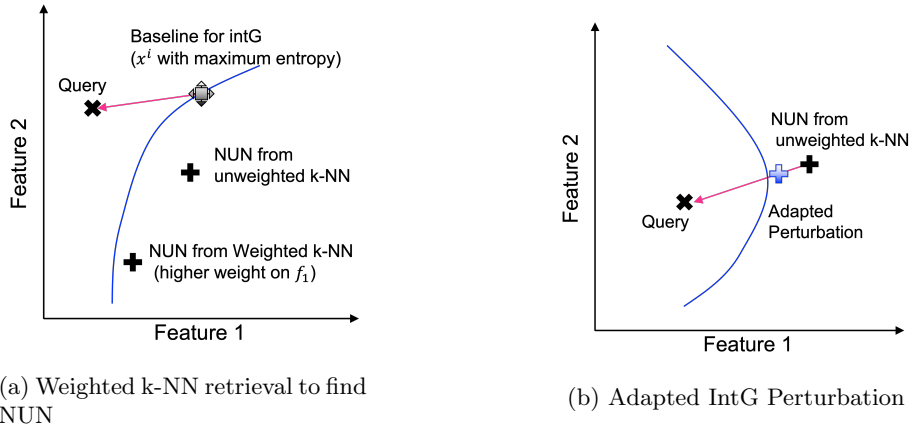


Fig. 6: Bringing NUN Closer

Weighted k-NN retrieval uses a feature attribution technique to select the subset of features to attend to when calculating distance. It then returns the NUN, x' , for query, x , as follows:

$$\begin{aligned}
 w &= \phi(x, \dots) \\
 \mathcal{N}_w &= \arg \min_{(x^i, y^i) \in X} d(x, w, x^i); y^i \neq y \\
 (x', y') &\leftarrow \mathcal{N}_w((x, y), \phi, X)
 \end{aligned} \tag{7}$$

Here ϕ returns the feature attributions w by adopting a method such as LIME [13], SHAP [14] or IntG [10]. The ... denotes any other requirements of the ϕ to generate feature attributions (such as method specific hyper-parameters). Note that when ϕ is IntG, a baseline must be specified before feature attribution weights can be computed. This weighted retrieval is illustrated in Figure 6a for a given query x . The baseline for IntG here is an instance that maximises uncertainty, as measured by entropy in Equation 5 (vanilla version). In our 2D space the feature attributions for Features 1 and 2 are $w_1 \geq w_2$, as such \mathcal{N}_w will return the new NUN which is closer to the query with respect to f_1 .

Adapted Perturbation approach creates a synthetic data instance by creating a perturbed instance between the NUN (retrieved by \mathcal{N}) and the query, x , which is closer to the query. This is illustrated in Figure 6b with an example in 2D space where the "blue" coloured cross shows an ideal synthetic perturbed NUN that is closer to the query compared to the actual NUN. The discovery of this perturbed NUN, x' , can be formalised as follows:

$$\begin{aligned} \alpha^* &= \arg \max_{\alpha \in \mathcal{A}} \alpha; F(x^\alpha) \neq y \\ \mathcal{N}_p &= \begin{cases} x'_i & \text{if } i \text{ is a symbolic feature} \\ x_i^{\alpha^*} & \text{otherwise} \end{cases} \quad (8) \\ (x', y') &\leftarrow \mathcal{N}_p((x, y), \mathcal{N}, X) \end{aligned}$$

First, an \mathcal{A} number of perturbations are created using Equation 3 from the baseline to the query (see Figure 4). The black box model, F , is used to predict the class for each perturbed data instance in turn. The selected perturbation, α^* , is the closest to the boundary with a prediction matching the desired class. Depending on the feature type (i.e. numeric or symbolic), values from the baseline, x' , and x^{α^*} are used to form the adapted perturbation. Specifically, symbolic features are substituted from x' and numerical features from the selected perturbation x^{α^*} . Note that the latter will have interpolated numeric values discovered as part of the perturbations. As a result, \mathcal{N}_p returns a synthetic NUN, which preserves plausibility [5] when used in DisCERN's adaptation operations.

In figure 6b the adapted perturbation is on the straight-line path between the NUN from \mathcal{N} and the query. Here both Features 1 and 2 are numeric. However with categorical features, the resulting adapted perturbation may not fall on the straight-line, yet we confirm that $F(x^{\alpha^*}) \neq y$ (i.e. ensures a valid NUN from a different class to that of the query).

4 Evaluation

Three comparative experiments were conducted as follows:

- a comparison of feature attribution methods for query adaptation in counterfactual discovery with DisCERN;
- a comparison of DisCERN against other counterfactual discovery algorithms in literature; and
- a comparison of approaches to find closer NUNs for counterfactual discovery.

4.1 Datasets

Five public domain datasets were used - Loan and Income datasets are from Kaggle; and Credit, Cancer and ICU datasets are from the UCI Machine Learning Data Repository. Table 2 summarises the properties of each dataset and the explanation needs as actionable questions. The labels column shows two class

Table 1: Dataset Details

Datasets	Description	Actionable Question	Labels
Loan	Predicts if a loan application is approved or not	What changes to the application will help to get the loan approved?	Q:Rejected CF:Approved
Income	Predicts if a persons income is higher or lower than 50K based on Census data	What changes will help an adult get an income over 50K?	Q: \leq 50K CF: \geq 50K
Credit	Predict if a credit card application is successful or not	What changes to the application will help to get a credit card?	Q:Rejected CF:Approved
Cancer	Predict the level of lung cancer risk of a patient based on demographic and triage data	How to reduce lung cancer risk of a patient?	Q:High risk or Medium risk CF:Low risk
ICU	Predict patient outcome in the Intensive Care Unit (ICU) based on demographic and diagnostic data	Which features will help prevent the death of a patient in the ICU?	Q:Deceased CF: Discharged

labels: the desirable (CF); and less desirable (Q). Typically a counterfactual explanation is needed to explain how a desirable situation could have been achieved when the black box model prediction is the less desired class.

Table 2: Dataset Properties

Datasets	Loan	Income	Credit	Cancer	ICU
Features	69	14	14	12	107
Categorical Features	8	8	8	11	8
Classes	2	2	2	3	2
Data instances	342865	45222	653	427	6238
Negative test instances	22543	11216	121	109	1003
RF Test Accuracy (%)	99.51	85.66	77.31	87.94	81.70
NN Test Accuracy (%)	99.66	85.02	75.00	87.23	73.08

4.2 Experiment Setup

The experimental pipeline to measure performance of a counterfactual discovery method appears in Figure 7. First the dataset is split as 2/3 train and 1/3 test data. Next, a black-box classifier is trained and used to populate a case-base of labelled train data. In this paper we consider two classifiers: Random Forest and Neural Networks. Test set performance of these 2 classifiers are in Table 2. The test instances classified into the less desirable class are the queries which form the Negative test set (see Table 1 for more details). Finally, the counterfactual discovery algorithm uses the classifier and the populated case-base to find counterfactuals for all query instances in the Negative test set. DisCERN uses Euclidean distance as $d(\cdot)$; and it is available publicly on GitHub ¹.

¹ <https://github.com/RGU-Computing/DisCERN-XAI>




Fig. 7: Experiment Setup

4.3 Performance Measures for Counterfactual Explanations

Validity (V) measures the percentage for which an algorithm is able to find a counterfactual with the desired label. If N is the total number of queries and N_v the number of counterfactuals found, then Validity is calculated as $V = \frac{N_v}{N} \times 100$.

Sparsity ($\#F$) is the average number of feature differences between a query (x) and its counterfactual (\hat{x}) explanation. Here m is the number of features.

$$\#F = \frac{1}{N_v \times m} \sum_{j=1}^{N_v} \sum_{i=1}^m 1_{[\hat{x}_i \neq x_i]} \quad (9)$$

Proximity ($\$F$) measures the mean feature-wise distance between a query and its counterfactual explanation. The sum of normalised feature differences are average over $\#F$ and the number of counterfactuals (N_v).

$$\$F = \frac{1}{N_v \times \#F} \sum_{j=1}^{N_v} \sum_{i=1}^m (|\hat{x}_i - x_i|) \quad (10)$$

Datasets with more categorical features will have higher $\$F$ (difference is always 1), meaning, $\$F$ is not comparable across datasets. Note that higher values of V and smaller values of $\#F$ and $\$F$ are preferred.

5 Results

5.1 A comparison of feature attribution techniques

Feature attributions determine the order in which values are copied from NUN to query. This adaptation knowledge guides the value substitution operations until a valid counterfactual is discovered. In Figure 8 we compare DisCERN with five feature attribution variants for ordering substitution operations: **Random** (i.e. no adaptation knowledge); \mathcal{X}^2 global feature selection; **LIME** local surrogate relevance weights, **SHAP**ely values from the KernelSHAP implementation; and **IntG** with the NUN (\mathcal{N}) baseline. Here, DisCERN uses the NUN from unweighted k-NN retrieval (i.e. \mathcal{N}) In Figure 8, 5 marker shapes corresponds to each of the 5 algorithms, and 5 marker colours relate to results corresponding

to each dataset. When sparsity and proximity are low the markers appear closer to the origin - this is desirable. Note that for visual clarity the x-axis is in log scale.

IntG achieves the lowest sparsity while random and \mathcal{X}^2 techniques achieves lower proximity. IntG’s use of the NUN baseline has successfully contributed to reducing sparsity, i.e. IntG’s feature attributions had significantly reduced the number of adaptation operations, when compared to using feature attributions from LIME and SHAP. Generally we found that with all datasets, making larger changes to feature values ($\$F$) leads to discovering a counterfactual with fewer feature substitutions ($\#F$). Therefore when we consider the total change ($\#F \times \$F$), IntG is best on all datasets except on Loan, where \mathcal{X}^2 was better. We also see that IntG had higher proximity (0.72 compared to the rest in a range of 0.28-0.39) on the Cancer dataset, but it had significantly lower sparsity (fewer features adapted), with almost half that of the other attribution methods (log scale can mask this difference). This can be a result of IntG selecting categorical features compared to other techniques.

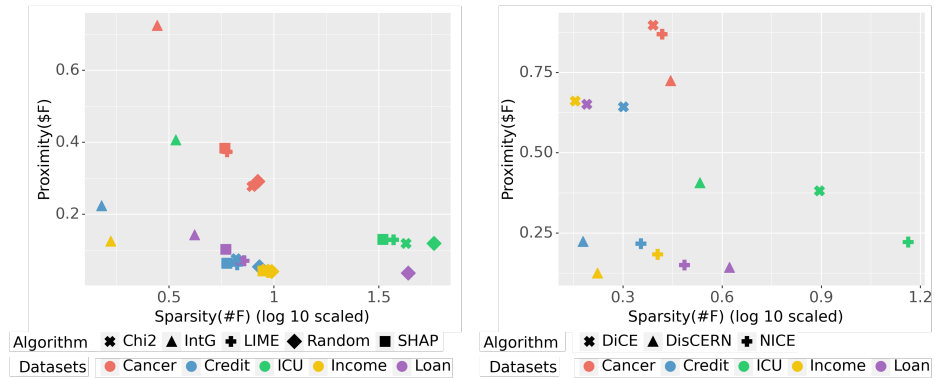


Fig. 8: Feature attribution results

Fig. 9: Counterfactual algorithm results

5.2 A Comparison of counterfactual discovery algorithms

A sparsity and proximity comparison of DisCERN (using $\Phi = IntG$) with popular counterfactual discovery algorithms from literature appear in Figure 9. DiCE is a generative algorithm that discovers counterfactuals by optimising a randomly initialised input to maximise diversity and minimise sparsity and proximity [11]. NICE is a NUN based counterfactual discovery algorithm that uses a reward function to minimise sparsity, proximity and to preserve plausibility [7].

Overall minimum proximity is seen with DisCERN (lowest on credit, income and loan). DiCE is better for sparsity (lowest on cancer, income and loan) but at the expense of proximity, which leads to larger value adaptations (i.e. higher

Table 3: Validity and total change between counterfactual discovery algorithms

Algorithm	Validity					Total Change				
	Loan	Income	Credit	Cancer	ICU	Loan	Income	Credit	Cancer	ICU
DiCE	99.20	92.55	80.99	100	33.80	1.009	0.945	1.286	2.207	2.991
NiCE	97.70	92.50	85.12	82.57	82.35	0.461	0.466	0.491	2.278	3.230
DisCERN	100	100	100	100	100	0.599	0.209	0.338	2.014	1.387

proximity) of the query applied to a fewer number of features (i.e. lower sparsity). Between NUN-based algorithms NICE and DisCERN, the lower sparsity and proximity achieved with DisCERN makes it a better candidate for counterfactual discovery.

Crucial to counterfactual explainers is the ability to find a valid counterfactual situation (i.e. one with the desired class). Results in Table 3 show that DisCERN dominates in this respect, whereby the NUN guided counterfactual is 100% valid. This is in comparison to DiCE, which uses a generative model given a randomly initialised input and NICE, which like DisCERN, uses a NUN as the basis but uses a reward function to find a counterfactuals. As a result, DiCE and NICE are likely to either fail to generate a counterfactual, or fail to generate a counterfactual that does belong to the desired class (in a multi-class setting). Note that for DiCE and NICE results in Figure 9, we have reported averaged results using only those queries for which a valid counterfactual was discovered (N_v).

Since we observed a trade-off between sparsity and proximity (in Figure 9), these are combined into a total change ($\#F \times \$F$) in the results presented in Table 3. We can see that DisCERN is best on all datasets (except Loan where it is marginally worse) as it manages to better balance proximity and sparsity requirements.

5.3 Impact of bringing NUN closer

Figure 10 compares DisCERN counterfactual discovery with the weighted (\mathcal{N}_w) and non weighted kNN (\mathcal{N}) NUN retrieval methods. The additional adapted perturbation (\mathcal{N}_p) results is included for IntG feature attribution method. Here in each evaluation DisCERN ensures $\Phi = \phi$. For example, if DisCERN feature ordering (Φ) is using SHAP feature attributions, weighted k-NN retrieval also uses SHAP weights ($\phi = SHAP$) to find the closer NUN. We note that \mathcal{N}_w approach was not evaluated with the Cancer dataset since we were unable to find a valid baseline that maximises entropy in the comparably smaller train set.

Overall \mathcal{N}_p with IntG and \mathcal{N}_w with SHAP and IntG show increased sparsity and decreased proximity compared to \mathcal{N} . Lower proximity is a result of bringing the NUN closer with respect to a subset of important features. Here the resulting NUN will have smaller value differences (decreased proximity) with the query. Therefore, initial substitutions will only make smaller actionable changes to the query. These initial features were already part of the subset that influenced the weighted kNN retrieval. As a result, the query will need an increased number




Fig. 10: Results on alternative approaches to form closer NUNs

of feature changes (increased sparsity). However, empirical evidence remains inconclusive for \mathcal{N}_w with LIME. This might be explained by LIME being a local surrogate attribution method. Accordingly, \mathcal{N}_w with LIME returns a NUN closer to the decision boundary of the local surrogate model, not of the global model. As a result, DisCERN makes larger changes to the query (i.e. increased proximity) in earlier substitutions which decreases the number of feature changes needed (decreased sparsity). In contrast, SHAP or IntG feature attributions seems globally faithful and better aligned to the black box decision boundary. Similarly, adapted perturbation based NUN is also globally faithful which results in increased sparsity and decreased proximity.

6 Conclusions

The XAI DisCERN algorithm discovers counterfactual explanations by applying a set of substitution-based adaptation operations to the NUN. These adaptations help create counterfactual explanations for black box predictions. DisCERN uses feature attribution weights, firstly to guide adaptation choices and, secondly to identify NUNs closer to the decision boundary. Both lead to valid counterfactuals from minimally adapted queries. The use of integrated gradients for feature attribution by limiting the integral approximation to the query and NUN is a unique contribution of this paper. Results suggests that this approach is more effective in guiding counterfactual discovery compared to feature attributions that do not consider both the desired and current prediction boundary locality.

An interesting study of where to position the NUN in support of counterfactual discovery found that it was possible to significantly minimise the total actionable changes required to create the counterfactual explanation, but also increased the number of feature changes needed for the desirable class change. In conclusion, DisCERN discovers effective counterfactuals, and provides the flexibility to select NUN and feature attribution technique based on the AI model and the preference of feature changes in the problem domain. Future work will look to expand DisCERN to generate counterfactual explanations that can also satisfy immutable and constrained feature requirements.

References

1. O. Li, H. Liu, C. Chen, and C. Rudin, “Deep learning for case-based reasoning through prototypes: A neural network that explains its predictions,” in *The Thirty-Second AAAI Conference on Artificial Intelligence*, pp. 3530–3537, 2018.
2. A. B. Arrieta, N. Díaz-Rodríguez, J. Del Ser, A. Bennetot, S. Tabik, A. Barbado, S. García, S. Gil-López, D. Molina, R. Benjamins, *et al.*, “Explainable artificial intelligence (xai): Concepts, taxonomies, opportunities and challenges toward responsible ai,” *Information Fusion*, vol. 58, pp. 82–115, 2020.
3. E. M. Kenny and M. T. Keane, “Twin-systems to explain artificial neural networks using case-based reasoning: Comparative tests of feature-weighting methods in ann-cbr twins for xai,” in *28th Int. Joint Conf. on AI, IJCAI-19*, pp. 2708–2715, International Joint Conferences on Artificial Intelligence Organization, 7 2019.
4. D. Wettschereck, D. W. Aha, and T. Mohri, “A review and empirical evaluation of feature weighting methods for a class of lazy learning algorithms,” *Artificial Intelligence Review*, vol. 11, no. 1, pp. 273–314, 1997.
5. R. M. Byrne, “Counterfactuals in explainable artificial intelligence (xai): Evidence from human reasoning,” in *IJCAI*, pp. 6276–6282, 2019.
6. M. T. Keane and B. Smyth, “Good counterfactuals and where to find them: A case-based technique for generating counterfactuals for explainable ai (xai),” in *International Conference on Case-Based Reasoning*, pp. 163–178, Springer, 2020.
7. D. Brughmans and D. Martens, “Nice: An algorithm for nearest instance counterfactual explanations,” *arXiv preprint arXiv:2104.07411*, 2021.
8. N. Wiratunga, A. Wijekoon, I. Nkisi-Orji, K. Martin, C. Palihawadana, and D. Corsar, “Discern: Discovering counterfactual explanations using relevance features from neighbourhoods,” in *33rd ICTAI*, pp. 1466–1473, IEEE, 2021.
9. S. Craw, N. Wiratunga, and R. C. Rowe, “Learning adaptation knowledge to improve CBR,” *Artificial Intelligence*, vol. 170, no. 16-17, pp. 1175–1192, 2006.
10. M. Sundararajan, A. Taly, and Q. Yan, “Axiomatic attribution for deep networks,” in *International conference on machine learning*, pp. 3319–3328, PMLR, 2017.
11. R. K. Mothilal, A. Sharma, and C. Tan, “Explaining machine learning classifiers through diverse counterfactual explanations,” in *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, pp. 607–617, 2020.
12. A.-H. Karimi, G. Barthe, B. Balle, and I. Valera, “Model-agnostic counterfactual explanations for consequential decisions,” in *International Conference on Artificial Intelligence and Statistics*, pp. 895–905, PMLR, 2020.
13. M. T. Ribeiro, S. Singh, and C. Guestrin, “Why should i trust you? explaining the predictions of any classifier,” in *Proceedings 22nd ACM SIGKDD Int conference on knowledge discovery and data mining*, pp. 1135–1144, 2016.
14. S. M. Lundberg and S.-I. Lee, “A unified approach to interpreting model predictions,” *Advances in Neural Information Processing Systems*, vol. 30, pp. 4765–4774, 2017.
15. J. Li, C. Zhang, J. T. Zhou, H. Fu, S. Xia, and Q. Hu, “Deep-lift: deep label-specific feature learning for image annotation,” *IEEE Transactions on Cybernetics*, 2021.
16. Z. Qi, S. Khorram, and F. Li, “Visualizing deep networks by optimizing with integrated gradients,” in *CVPR Workshops*, vol. 2, 2019.