



MACFE: A Meta-learning and Causality Based Feature Engineering Framework

Ivan Reyes-Amezcu¹(✉), Daniel Flores-Araiza², Gilberto Ochoa-Ruiz²,
Andres Mendez-Vazquez¹, and Eduardo Rodriguez-Tello³

¹ Department of Computer Science, CINVESTAV, Guadalajara, Mexico
ivan.reyes@cinvestav.mx

² Tecnológico de Monterrey, School of Engineering and Sciences, Monterrey, Mexico

³ Cinvestav - Tamaulipas, Km. 5.5 Carretera Victoria - Soto La Marina,
87130 Victoria, Tamaulipas, Mexico

Abstract. Feature engineering has become one of the most important steps to improving model prediction performance, and producing quality datasets. However, this process requires non-trivial domain knowledge which involves a time-consuming task. Thereby, automating such processes has become an active area of research and interest in industrial applications. In this paper, a novel method, called Meta-learning and Causality Based Feature Engineering (MACFE), is proposed; our method is based on the use of meta-learning, feature distribution encoding, and causality feature selection. In MACFE, meta-learning is used to find the best transformations, then the search is accelerated by pre-selecting “original” features given their causal relevance. Experimental evaluations on popular classification datasets show that MACFE can improve the prediction performance across eight classifiers, outperforms the current state-of-the-art methods on average by at least 6.54%, and obtains an improvement of 2.71% over the best previous works.

Keywords: Automated feature engineering · Automated machine learning · Causal feature selection

1 Introduction

Extracting features from raw data and transforming them into formats that are appropriate for Machine Learning (ML) models is what is known as feature engineering [12]. This task is usually carried out by a data scientist with good domain knowledge and the data sources of the task at hand [19, 21, 33]. Generally, feature engineering entails the daunting manual labor of designing, selecting, and evaluating features where even a great intuition is needed [6, 18]. This is due to the fact that the performance of most machine learning algorithms heavily relies on the training data quality. These type of datasets usually consists of a large collection of different formats that need to be curated to be exploited by machine

learning algorithms [6]. Therefore, by using feature engineering, we can select and obtain novel features from the raw data that would better represent the problem.

However, most of the existing automated feature engineering proposals perform this task by applying the *expansion-reduction* method [17], which is the process of trying a predefined set of transformation functions applied to raw features. Then, those transformed features are selected based on the improvement of model performance or some evaluation metric [21]. However, *expansion-reduction* leads to an exponential growth in the space of constructed features, which is known as the *feature explosion* problem [5]. In addition, extracting novel features without a proper and systematic method can lead to an unnecessary increase in the dimensionality of the data, and hence a poor performance in the learning process of the model [3]. Thus, the *curse of dimensionality* arises [20], which is the potential of high-dimensional data to be more complicated to process than low-dimensional data [8].




Fig. 1. The framework of our method. MACFE extracts meta-features from dataset D and a frequency table for each feature $x \in X$. Then, an encoding e is generated by the meta-features and feature distribution. Next, we search for the most similar encoding on the Transformation Recommendation Matrix (TRM) in order to recommend a useful transformation from it. The transformed dataset \hat{D} is built from the constructed novel features and the original ones selected by the Directed Acyclic Graph (DAG) causal model.

It is crucial to realize that there are dozens of types of machine learning models, and each has its peculiarities and needs [19]. For instance, some models neither work with highly correlated features nor with highly multi-collinearity. Additionally, other models have trouble dealing with missing, noisy, or irrelevant features. Furthermore, since data and models are so diverse, it is difficult to generalize the practice of feature engineering across projects [33]. Thus, finding a proper process to treat data agnostically from a specific learning algorithm can help to choose transformations that better suit the learning process. To tackle this issue, a possibility is to incorporate only the generated features that have more appropriate knowledge about the data. For this, we present MACFE, a novel meta-learning and causality approach for automated feature engineering for classification problems using tabular data. The main contributions of this paper are briefly described as follows:

- We present a causality-based method for feature selection on the original dataset. For this, we use the mean magnitude effect of the features on the target to rank and select a subset of them.
- We propose a novel meta-learning generation for unary, binary, and high-order features based on non-linear transformations. This approach addresses the feature explosion problem by only searching for feature transformations that were found useful in past experiences.

In order to evaluate the proposed method, we designed a series of experiments on fourteen popular public classification datasets with relatively small dimensions to evaluate the feature generation and selection performance of MACFE. The results are obtained from eight machine learning models: Logistic Regression (*LR*), K-Nearest Neighbors (*KNN*), Lineal Support Vector Machine (*SVC-L*), Polynomial Support Vector Machine (*SVC-P*), Random Forest (*RF*), AdaBoost (*AB*), Multi-layer Perceptron (*MLP*) and Decision Tree (*DT*). As illustrated in Fig. 1, our approach is divided into three phases. In the first one, the feature selection is carried out by using a Structural Causal Model (SCM) [22] for choosing the most promising features. Then, we move to a meta-learning phase (the second one), where meta-features are extracted from datasets and feature distributions to create encodings for each attribute. Then, we lookup for feature transformation on similar previously engineered datasets. Finally, in the third phase, we evaluate the engineered features among eight machine learning models and obtain the mean accuracy of stratified 5-Fold Cross Validation in order to assess the quality of the feature engineering method. Experimental results show that our proposal is effective in surpassing the scores of state-of-the-art feature engineering methods by achieving a mean accuracy of 81.83% across the fourteen testing datasets and the eight machine learning models evaluated.

The rest of this paper is organized as follows. In Sect. 2, we review the state of the art in automated feature engineering. In Sect. 3, we elaborate on the need for automated methods like ours. In Sect. 4, we introduce our proposed method MACFE, whereas Sect. 5, we show in detail our evaluation results, and finally, in Sect. 6 the conclusions drawn in this research work are given.

2 Related Work

In recent years, many automated feature engineering methods have been proposed using different methodologies. For instance, Data Science Machine (DSM) [14] is an automated feature engineering approach for structured and relational data. DSM proposed a Deep Feature Synthesis (DFS) method, which searches for relations and transformations across features in databases. They include a depth hyper-parameter d for setting the maximum composition, which recursively enumerates all possible transformations. In addition, DSM generates a large novel feature space, which is reduced by using Singular Value Decomposition (SVD) based feature selection. However, DSM is only suitable for relational data and it could take high computational times due to all the transformation functions used for processing all the original feature sets.

The data-driven approach presented in FCTree [9], creates novel features from sequential transformations of the original space by employing decision trees and then selecting the best features with the aid of information gain. The method in [25], known as the TFC framework, presents an iterative feature generation algorithm. The method applies feature transformation across all the features, and then it selects the best features based on information gain. Nevertheless, the generated feature space grows in a combinatorial way, leading to feature explosion. AutoFeat [13] and AutoLearn [16] are also data-driven methods. They can generate large transformations of features, selecting useful features by using regularized regression models for each pair of features. However, these methods require training a regression model, which can be time-consuming. Also, they both suffer from the feature explosion problem. Label based Regression (LbR) [30] is another method for generating novel features by using Ridge Regression and Kernel Ridge Regression. This method selects features based on the Distance Correlation Coefficient and the Maximum Information Coefficient (MIC) for each feature pair, which leads to discriminate features that are useful in combination with others.

2.1 Meta-learning for Feature Engineering

Recently, meta-learning has been proposed as a means of improving the quality of the generated features [21]. Meta-data can be simply defined as data about data [31]. For this work, meta-features are used to characterize and identify features with attributes in the context of meta-learning [1, 4, 10, 28]. Some examples of meta-features are: a) *General*, such as the number of samples, features or classes, etc. b) *Statistical* like standard deviation, correlation coefficients, etc. c) *Information-theoretic* such as entropy, mutual information, noise ratio, etc. d) *Model-based* describes some characteristics of models such as Decision Trees, Bayesian Networks, SVMs, etc.

ExploreKit [15] is an example of a method that uses meta-learning for ranking and selecting the most promising generated features. ExploreKit does this by applying all possible transformations on features, suffering from the feature explosion problem. Furthermore, Learning Feature Engineering (LFE) is an approach that also uses meta-learning for recommending useful features for classification problems. The transformation recommender in LFE is based on the construction of a meta-feature vector based on the feature values associated with a class label. However, LFE can recommend only unary and binary transformations, lacking high-order transformations.

2.2 Causality Feature Selection

Classical feature selection approaches to leverage the correlations between features and class variables but lack taking advantage of the causal relationships between them. In contrast, knowing the causal relationship implies the underlying mechanism of a dataset [32], and thus causal variables are expected to be persistent across different settings or environments.

Hence, basing the feature engineering on relevant causally related features to the class of interest ideally should provide a more rich and robust output of engineered features. Consequently, if we work only with causally related variables to the target variable, independently of the type of relationship, it should be possible to be learned by an ML model, which additionally implies it facilitates, at some level, the efficacy of applying feature engineering to causally related variables.

3 Problem Definition

Let $D = \{\mathbf{X}, \mathbf{Y}\}$ be a dataset of input-output pairs, \mathbf{X} a collection of n features $\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n\}$, and m labels $\mathbf{Y} = \{y_1, \dots, y_m\}$. A machine learning algorithm L (e.g. SVM, Logistic Regression, or Random Forest), and an evaluation metric E (e.g. accuracy, F1-score).

We refer to a transformation $t \in \mathbf{T}$ as a function $t(\mathbf{x})$ that takes a feature as an argument, and maps it to a transformed feature output $\hat{\mathbf{x}} \in \mathbf{X}'$. Where \mathbf{T} is our set of transformations $\{t_1, t_2, \dots, t_k\}$ that can be unary or binary, depending on the number of given arguments. Here, a high-order transformation is a composition of unary and binary transformations. Over each feature it is possible to define a series of non-linear transformations, $t_i : \mathbf{x}_i \rightarrow \hat{\mathbf{x}}_i$ that allow to extract as much intra and inter information from the “original” data. The goal of feature engineering is thus to transform \mathbf{X} into \mathbf{X}' by applying \mathbf{T} such that \mathbf{X}' maximizes the evaluation metric E of a machine learning algorithm L . The search for new transformed features and their combinations grows exponentially, and the feature explosion problem arises. MACFE, our proposed feature engineering approach, was devised to help mitigate this problem by employing meta-features to guide the search for transformations on features.

3.1 Meta-learning and Meta-features

A formal definition of meta-features was proposed in [28], in which meta-features are a set of q values extracted from a dataset D by a function f .

$$f(D) = \sigma(\mu(D, h_\mu), h_\sigma), \quad (1)$$

where $f : D \mapsto \mathbb{R}^q$ is the extraction of q values from dataset D , $\mu : D \mapsto \mathbb{R}^{q'}$ is a characterization measure, $\sigma : \mathbb{R}^{q'} \mapsto \mathbb{R}^q$ can be a summarization function such as: mean, minimum, maximum, etc. Moreover, h_μ and h_σ are hyperparameters for μ and σ , respectively. Thus, the function f is built by measuring some characteristic from D by μ , and a summarizing function defined by σ .

Here, meta-features describe features using meta-data. An example is the mean or median, as they are features that provide extra information about the underlying data distribution. In particular, the core of this work is meta-learning applied to the identification of data through meta-features.

4 Proposed Approach

In the following sections, we describe the dataset preprocessing along with the construction of the meta-feature vector and encodings for features. Also, we present the training of our method including the Meta-learning and Causal Selection phases.

4.1 Datasets

Preprocessing. MACFE is guided by meta-feature learning based on past experience to create novel features. Our method is trained with M random datasets $D_{train} = \{D_1, D_2, \dots, D_M\}$ collected from Open ML [29], which have a structured format and a classification task related to the data. First, the preprocessing and cleaning of data are performed for each dataset by removing non-numerical features and imputing missing values with the feature mean. Next, a meta-feature *extractor* is used to obtain meta-data about the datasets. Let \mathbf{mf} be a meta-feature vector composed by the main characteristics of a given dataset $D_i \in D_{train}$. Thus, a meta-feature vector for a dataset D_i is defined as:

$$\mathbf{mf} = [mf_1, mf_2, \dots, mf_p], \quad (2)$$

where each mf_i is a meta-feature value extracted from the data, and p is the size of the extracted meta-features.

However, describing datasets by mapping their main characteristics can be a challenging task. A full set of estimators and metrics can be extracted from a dataset, e.g., the number of classes or instances in a dataset can be a meta-feature value from such a dataset. For this, we use the approach of [24] to perform the automatic meta-feature extraction process. The extraction of meta-features is divided into five categories proposed by Rivolli *et al.* [28]: simple or general, statistical, information-theoretic and model-based, and landmarking. In order to automate the process of extracting meta-features from datasets, we use the framework Meta-feature Extractor (MFE) [1] for each training datasets $D_i \in D_{train}$, which implements the standard meta-feature extraction described above.

Next, we treat each feature $\mathbf{x} \in D_i$ as follows:

1. We create a frequency table with a fixed number of *buckets* or *bins* b , for each feature \mathbf{x}
2. A range r is calculated on the feature values given by the *upper* and *lower* bounds of the feature.
3. We generate s disjoint sets or bins b with uniform width w . Thus, each bin range b_i is a bucket in which values that are in the bin range lie. Each b_i range starts with the lower bound of \mathbf{x} plus i times the width w , and ends with the lower bound of \mathbf{x} plus $i + 1$ times the width w .
4. Finally, each frequency table or histogram is normalized in the range $[0,1]$.

Thus, we obtain an encoding $\mathbf{e} \in \mathbb{R}^{1 \times \eta}$ for each feature $\mathbf{x} \in \mathbf{D}_i$, composed by the meta-feature vector \mathbf{mf} of the dataset and the feature distribution as follows:

$$\mathbf{e} = [mf_1, mf_2, \dots, mf_p, b_0, b_1, \dots, b_{s-1}] \quad (3)$$

4.2 Model Training

Meta-learning Phase. The meta-learning phase is described as follows. The unary, binary, and scaling feature transformations $t \in \mathbf{T}$ are applied to the original features \mathbf{X} . Then, an evaluation is performed on both original features and the generated features $t(\mathbf{X})$. For this, we use the Maximal Information Coefficient (MIC) [27], which measures the strength of the linear or non-linear relationships between two variables. MIC generates values between 0 and 1, where 0 means statistical independence and 1 stands for a noiseless statistical relationship between variables. Thus, we get the set of selected transformations \mathbf{T}_{sel} for each original feature in $\mathbf{x} \in \mathbf{X}$ with the maximum score as follows:

$$\mathbf{T}_{sel} = \operatorname{argmax}_{t \in \mathbf{T}} g_t \left(MIC(t(\mathbf{x})) - MIC(\mathbf{x}) \right). \quad (4)$$

Finally, the selected transformations $t \in \mathbf{T}_{sel}$ are stored in the Transformation Recommendation Matrix (\mathbf{TRM}) for each $\mathbf{x} \in \mathbf{D}_{train}$ represented by its corresponding encoding \mathbf{e} . Thus, \mathbf{TRM} is represented as follows (Fig. 2).

$$\mathbf{TRM} = \begin{bmatrix} e_{1,1} & e_{1,2} & \cdots & e_{1,\eta} & t_1 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ e_{N,1} & e_{N,2} & \cdots & e_{N,\eta} & t_N \end{bmatrix}$$

Fig. 2. \mathbf{TRM} Matrix, where the i^{th} row in the matrix is the feature $\mathbf{x} \in \mathbf{D}_{train}$, and the j^{th} column is the encoding value of \mathbf{e} (Eq. 3). N is the size of all the features in \mathbf{D}_{train} , and η is the size of encoding \mathbf{e} composed by the meta-feature vector \mathbf{mf} (Eq. 2) and feature histogram. The last column stands for the transformations $t \in \mathbf{T}$ with the resulting highest MIC score for the given features (Eq. 4).

In Algorithm 1 the training procedure to learn the most appropriated unary \mathbf{T}_{un} and binary \mathbf{T}_{bin} transformations is presented. This process is done for each feature in a given dataset \mathbf{D} . Similarly, high-order transformations are built by combining several unary or binary transformations one after the other (Algorithm 2).

Algorithm 1. Training *TRM*

Input: Structured Dataset D
Output: *TRM*
 $D = \text{preprocess}(D)$
for each $\mathbf{x}_i \in D$ **do**
 $\mathbf{e}_i = \text{encode_feature}(\mathbf{x}_i)$
 for each $t \in \mathbf{T}_{un}$ **do**
 $\hat{\mathbf{x}}_i = t(\mathbf{x}_i)$
 $s.\text{append}(\text{MIC}(\hat{\mathbf{x}}_i) - \text{MIC}(\mathbf{x}_i))$
 end for
 $t_{top} = \text{argmax}(s)$
 $\mathbf{TRM}.\text{append}(\mathbf{e}_i, t_{top})$
 for each $\mathbf{x}_j \in D | j > i$ **do**
 $\mathbf{e}_j = \text{encode_feature}(\mathbf{x}_j)$
 for each $t \in \mathbf{T}_{bin}$ **do**
 $\hat{\mathbf{x}}_{i,j} = t(\mathbf{x}_i, \mathbf{x}_j)$
 $s_i = \text{MIC}(\hat{\mathbf{x}}_{i,j}) - \text{MIC}(\mathbf{x}_i)$
 $s_j = \text{MIC}(\hat{\mathbf{x}}_{i,j}) - \text{MIC}(\mathbf{x}_j)$
 if $s_i > 0$ and $s_j > 0$ **then**
 $\mathbf{TRM}.\text{append}(\mathbf{e}_i, \mathbf{e}_j, t)$
 end if
 end for
 end for
end for

Algorithm 2. Data Transformation

Input: D, d, s
Output: \hat{D}
 $\hat{D} = \text{preprocess}(D)$
 $\hat{D} = \text{causal_selection}(\hat{D}, s)$
for 1 to d **do**
 for each $\mathbf{x}_i \in \hat{D}$ **do**
 $\mathbf{e}_i = \text{encode_feature}(\mathbf{x}_i)$
 $t_{un} = \text{Similarity}(\mathbf{TRM}, \mathbf{e}_i)$
 $\hat{\mathbf{x}}_i = t_{un}(\mathbf{x}_i)$
 $\hat{D}.\text{append}(\hat{\mathbf{x}}_i)$
 for each $\mathbf{x}_j \in \hat{D}, \mathbf{x}_i \neq \mathbf{x}_j$ **do**
 $\mathbf{e}_j = \text{encode_feature}(\mathbf{x}_j)$
 $t_{bin} = \text{Similarity}(\mathbf{TRM}, \mathbf{e}_i, \mathbf{e}_j)$
 $\hat{\mathbf{x}}_{i,j} = t_{bin}(\mathbf{x}_i, \mathbf{x}_j)$
 $\hat{D}.\text{append}(\hat{\mathbf{x}}_{i,j})$
 end for
 end for
 $e_{\hat{D}} = \text{encode_dataset}(\hat{D})$
 $t_{scaler} = \text{Similarity}(\mathbf{TRM}, e_{\hat{D}})$
 $\hat{D} = t_{scaler}(\hat{D})$

The order value of the transformation function is related to the number of times a feature is processed by a transformation, e.g., an input feature \mathbf{x}_1 is given as an argument of the *log* function, so $f_1(\mathbf{x}_1) = \log(\mathbf{x}_1)$. Then, the resulting feature is combined with another feature \mathbf{x}_2 , lets say a multiplication, thus, $f_2(f_1(\mathbf{x}_1), \mathbf{x}_2) = \text{mult}(\log(\mathbf{x}_1), \mathbf{x}_2)$. Finally, the output feature is given to the unary function *square*. Thus, the final transformed feature $\hat{\mathbf{x}}$ has an order of 3, and can be seen as follows:

$$\hat{\mathbf{x}} = f_3(f_2(f_1(\mathbf{x}_1), \mathbf{x}_2)) = \text{square}(\text{mult}(\log(\mathbf{x}_1), \mathbf{x}_2)) \quad (5)$$

Hence, we look for the underlying information about data through the extraction of more complex features. This gives us the capability of creating novel features from raw features that apparently do not have any predictive power, but in combination with high-order functions can have suitable predictive power for some machine learning models.

Causal Feature Selection Phase. Once the *TRM* is trained, MACFE is ready to recommend useful transformations for new datasets and features. For this, we start selecting the most promising original features, a causality-based feature selection is performed on the features. A DAG Classifier is trained to discover a causal graph from data. For this, we use the implementation of CausalNex [2]. This graph underlies the causal relationship between features and a target

variable. The mean identified causal magnitude effect of the features on the target is used to rank the features. Then, a given threshold hyperparameter s determines the top k selected features. The resulting subset of selected features are processed to obtain an encoding \mathbf{e} (Eq. 3).

Then, for a given feature encoding \mathbf{e} , we search for a transformation in *TRM* by retrieving the most similar feature encoding using the *cosine distance* as a similarity measure. We benefit from this measure for ranking the most similar feature-vectors in the range 1.0 for identical feature-vectors and 0.0 for orthogonal feature-vectors [26]. Next, the most similar feature transformation is applied to the feature. The process is followed by the binary transformations and iterating over the features in the dataset (Algorithm 2). Furthermore, a depth d hyper-parameter is set to look for the maximum transformation order across *unary* and *binary* functions. Lastly, for the *Scaling transformations* we refer to those transformations on features that change the scale on a standard range. Many machine learning algorithms struggle to find patterns in data when features are not on the same scale. For this, having scaled features can help gradient descent to converge faster towards a minimum.

We scale features as follows. For a given feature $\mathbf{x} \in \mathbf{X}$, the following scaling functions can be applied. *Normalization*, also called *Min-Max Scaler*, is a method that scales each feature value to the range [0,1]. *Standardization*, this method scales each feature value so that the mean is 0 and the standard deviation is 1. *Robust Scaler*, this scaler is useful when the input feature has a lot of *outliers*. The Robust Scaling is done by calculating the median (50th percentile), and also the 25th and 75th percentiles. Then, each feature value is subtracted from the median, and divided by the Interquartile Range (IQR). In order to learn and recommend which scaler is appropriate for a given dataset, we follow a series of data testings. First, we test the features to know the proportion of outliers. If this proportion is larger than a certain threshold γ , then a Robust Scaler is applied to the features. Secondly, if the data follows a normal distribution, then we use a Standard Scaler. In particular, we use a *Shapiro-Wilk* test [11] to evaluate the normality of data. Then, if the test value is greater than 0.05 we consider the data is normally distributed. Finally, if none of the above tests is true about the data, then we use a Min-Max Scaler on the features. The resulting scaling method is saved in TRM according to the dataset encoding.

5 Experimental Results

For the evaluation of MACFE, first, we describe the evaluation details, such as the case study datasets and learning algorithms. Next, we briefly describe each of the implementation details of the classifiers and evaluation methods. Finally, a comparison with previous work is done and a discussion is presented by analyzing the characteristics of datasets and algorithms where MACFE is convenient.

5.1 Evaluation Details

Table 1. Statistics of 14 case study datasets

ID	Dataset	Source	Labels	Features	Instances
1	<i>Pima Diabetes</i>	UCI ML	2	8	768
2	<i>Sonar</i>	UCI ML	2	60	208
3	<i>Ionosphere</i>	UCI ML	2	34	351
4	<i>Haberman</i>	UCI ML	2	3	306
5	<i>Fertility</i>	UCI ML	2	9	100
6	<i>Wine</i>	UCI ML	3	13	178
7	<i>E.coli</i>	UCI ML	8	7	336
8	<i>Abalone</i>	UCI ML	29	7	4177
9	<i>Dermatology</i>	UCI ML	4	34	366
10	<i>Libras</i>	UCI ML	15	90	360
11	<i>Optical</i>	UCI ML	10	64	5620
12	<i>Waveform</i>	OpenML	3	21	5000
13	<i>Fourier</i>	OpenML	10	76	2000
14	<i>Pixel</i>	OpenML	10	240	2000

The evaluation of MACFE as an automated feature engineering method is performed on a set of fourteen classification datasets and eight machine learning algorithms commonly cited in the literature [15, 16, 30]. These datasets are from different areas, such as medical, physical, life, and computer science. In addition, these datasets are publicly available in the UCI ML Repository [7] and OpenML Repository [29]. The main statistics of these datasets are shown in Table 1.

5.2 Implementation Details

For our experiments, we tested the following learning algorithms: Logistic Regression (*LR*), K-Nearest Neighbors (*KNN*), Linear Support Vector Machine (*SVC-L*), Polynomial Support Vector Machine (*SVC-P*) and Random Forest (*RF*), AdaBoost (*AB*), Multi-layer Perceptron (*MLP*) and Decision Tree (*DT*). The scoring method for the evaluations is the mean accuracy of stratified 5-Fold Cross Validation on each dataset. Same as the state-of-the-art methodology for scoring. Each algorithm is used with scikit-learn [23] default parameters. This is because our objective is to enhance the accuracy of a model by improving the data through our automated feature engineering process, MACFE.

5.3 Comparison with Previous Works

The comparison of our proposal takes into account the same scenario conditions of the results presented in recent feature engineering proposals such as TFC [25], FCtree [9], ExploreKit [15], AutoLearn [16] and LbR [30]. In Table 2 are shown

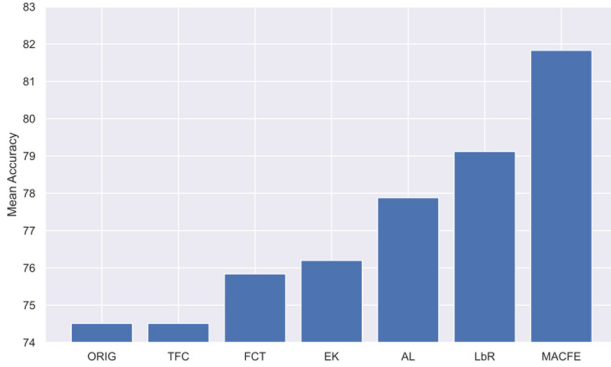


Fig. 3. Mean accuracy of state-of-the-art methods and MACFE (ours) across fourteen case study datasets and eight machine learning models.

5.4 Discussion

The transformation recommendation procedure of this method is agnostic of the learning algorithm. But, some transformations can be more appropriate for a certain algorithm. Therefore, MACFE achieves 100% of efficacy in terms of improving at least one model for each dataset. The depth hyperparameter d of MACFE can generate different orders of complex features to improve the model performance. A high value in d can result in too complex novel features, thus the algorithm cannot learn from the data. In contrast, a small value of the hyperparameter s can lead to a small subset of the original features, thus not finding good relationships between features. Hence, it is recommended a grid search to find the optimal values of hyperparameters.

6 Conclusions and Future Work

In this paper, we presented a causality-based feature selection to reduce the feature space search for feature transformations. Also, a meta-learning-based method for automated feature construction, on which the number of transformations executed on features depends on the number of useful transformations found on historical past similar features. In particular, this method has the capability of constructing novel features from raw data that are informative and useful for a learning algorithm. Hence, MACFE can automatically create features by applying selected transformations to the data, either unary, binary, or high-order, instead of applying all possible combinations of those. Hence, the feature explosion problem is minimized. However, MACFE has a fixed set of unary, binary, and scaling transformations. In future work, we intend to increase this set by adding more transformation functions, leading to the construction of more informative features from raw features. In addition, the causal selection of features could be improved, since it is applied equally to all datasets but different datasets can be expected to satisfy different causal assumptions, which

produces different levels of efficacy when selecting the features to be engineered. To improve this, better methods of general causal discovery are needed.

Acknowledgments. The authors wish to thank the CINVESTAV, the AI Hub, and the CIOT at ITESM, for their support and the use of the DGX for running the experiments in this paper.

References

1. Alcobaça, E., et al.: MFE: towards reproducible meta-feature extraction. *J. Mach. Learn. Res.* **21**(111), 1–5 (2020)
2. Beaumont, P., et al.: CausalNex (2021). <https://github.com/quantumblacklabs/causalnex>
3. Blumer, A., Ehrenfeucht, A., Haussler, D., Warmuth, M.K.: Learnability and the Vapnik-Chervonenkis dimension. *J. ACM (JACM)* **36**(4), 929–965 (1989)
4. Brazdil, P., Carrier, C.G., Soares, C., Vilalta, R.: *Metalearning: Applications to Data Mining*. Springer, Heidelberg (2008). <https://doi.org/10.1007/978-3-540-73263-1>
5. Chen, X., et al.: Neural feature search: a neural architecture for automated feature engineering. In: 2019 IEEE International Conference on Data Mining (ICDM), pp. 71–80. IEEE (2019)
6. Domingos, P.: A few useful things to know about machine learning. *Commun. ACM* **55**(10), 78–87 (2012)
7. Dua, D., Graff, C.: UCI machine learning repository (2017). <http://archive.ics.uci.edu/ml>
8. Duda, R.O., Hart, P.E., et al.: *Pattern Classification*. Wiley (2006)
9. Fan, W., et al.: Generalized and heuristic-free feature construction for improved accuracy. In: *Proceedings of the 2010 SIAM International Conference on Data Mining*, pp. 629–640. SIAM (2010)
10. Filchenkov, A., Pendryak, A.: Datasets meta-feature description for recommending feature selection algorithm. In: 2015 Artificial Intelligence and Natural Language and Information Extraction, Social Media and Web Search FRUCT Conference (AINL-ISMW FRUCT), pp. 11–18. IEEE (2015)
11. Hanusz, Z., Tarasinska, J., Zielinski, W.: Shapiro-Wilk test with known mean. *REVSTAT-Stat. J.* **14**(1), 89–100 (2016)
12. Heaton, J.: An empirical analysis of feature engineering for predictive modeling. In: *SoutheastCon 2016*, pp. 1–6. IEEE (2016)
13. Horn, F., Pack, R., Rieger, M.: The `autofeat` Python library for automated feature engineering and selection. In: Cellier, P., Driessens, K. (eds.) *ECML PKDD 2019. CCIS*, vol. 1167, pp. 111–120. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-43823-4_10
14. Kanter, J.M., Veeramachaneni, K.: Deep feature synthesis: towards automating data science endeavors. In: 2015 IEEE International Conference on Data Science and Advanced Analytics (DSAA), pp. 1–10. IEEE (2015)
15. Katz, G., Shin, E.C.R., Song, D.: ExploreKit: automatic feature generation and selection. In: 2016 IEEE 16th International Conference on Data Mining (ICDM), pp. 979–984. IEEE (2016)
16. Kaul, A., Maheshwary, S., Pudi, V.: AutoLearn-automated feature generation and selection. In: 2017 IEEE International Conference on Data Mining (ICDM), pp. 217–226. IEEE (2017)

17. Khurana, U., Samulowitz, H., Turaga, D.: Feature engineering for predictive modeling using reinforcement learning. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 32 (2018)
18. Khurana, U., Turaga, D., Samulowitz, H., Parthasarathy, S.: Cognito: automated feature engineering for supervised learning. In: 2016 IEEE 16th International Conference on Data Mining Workshops (ICDMW), pp. 1304–1307. IEEE (2016)
19. Kuhn, M., Johnson, K.: Feature Engineering and Selection: A Practical Approach for Predictive Models. CRC Press, Boca Raton (2019)
20. Kuo, F.Y., Sloan, I.H.: Lifting the curse of dimensionality. *Not. AMS* **52**(11), 1320–1328 (2005)
21. Nargesian, F., Samulowitz, H., Khurana, U., Khalil, E.B., Turaga, D.S.: Learning feature engineering for classification. In: IJCAI, pp. 2529–2535 (2017)
22. Pearl, J.: Causality: Models, Reasoning and Inference, 2nd edn. Cambridge University Press, Cambridge (2009)
23. Pedregosa, F., et al.: Scikit-learn: machine learning in Python. *J. Mach. Learn. Res.* **12**, 2825–2830 (2011)
24. Pinto, F., Soares, C., Mendes-Moreira, J.: Towards automatic generation of metafeatures. In: Bailey, J., Khan, L., Washio, T., Dobbie, G., Huang, J.Z., Wang, R. (eds.) PAKDD 2016. LNCS (LNAI), vol. 9651, pp. 215–226. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-31753-3_18
25. Piramuthu, S., Sikora, R.T.: Iterative feature construction for improving inductive learning algorithms. *Expert Syst. Appl.* **36**(2), 3401–3406 (2009)
26. Qian, G., Sural, S., Gu, Y., Pramanik, S.: Similarity between Euclidean and cosine angle distance for nearest neighbor queries. In: Proceedings of the 2004 ACM Symposium on Applied, pp. 1232–1237 (2004)
27. Reshef, D.N., et al.: Detecting novel associations in large data sets. *Science* **334**(6062), 1518–1524 (2011)
28. Rivolli, A., Garcia, L.P., Soares, C., Vanschoren, J., de Carvalho, A.C.: Towards reproducible empirical research in meta-learning. arXiv preprint [arXiv:1808.10406](https://arxiv.org/abs/1808.10406), pp. 32–52 (2018)
29. Vanschoren, J., Van Rijn, J.N., Bischl, B., Torgo, L.: OpenML: networked science in machine learning. *ACM SIGKDD Explor. Newsl.* **15**(2), 49–60 (2014)
30. Wang, M., Ding, Z., Pan, M.: LbR: a new regression architecture for automated feature engineering. In: 2020 International Conference on Data Mining Workshops (ICDMW), pp. 432–439. IEEE (2020)
31. Witten, I.H., Frank, E., Hall, M.A., Pal, C.J.: Practical Machine Learning Tools and Techniques. Morgan Kaufmann **578**, 1 (2005)
32. Yu, K., et al.: Causality-based feature selection: methods and evaluations. *ACM Comput. Surv. (CSUR)* **53**(5), 1–36 (2020)
33. Zheng, A., Casari, A.: Feature Engineering for Machine Learning: Principles and Techniques for Data Scientists. O’Reilly Media, Inc. (2018)