

DiffPrep: Differentiable Data Preprocessing Pipeline Search for Learning over Tabular Data

PENG LI, Georgia Institute of Technology, USA
ZHIYI CHEN, Georgia Institute of Technology, USA
XU CHU, Georgia Institute of Technology, USA
KEXIN RONG, Georgia Institute of Technology, USA

Data preprocessing is a crucial step in the machine learning process that transforms raw data into a more usable format for downstream ML models. However, it can be costly and time-consuming, often requiring the expertise of domain experts. Existing automated machine learning (AutoML) frameworks claim to automate data preprocessing. However, they often use a restricted search space of data preprocessing pipelines which limits the potential performance gains, and they are often too slow as they require training the ML model multiple times. In this paper, we propose DIFFPREP, a method that can automatically and efficiently search for a data preprocessing pipeline for a given tabular dataset and a differentiable ML model such that the performance of the ML model is maximized. We formalize the problem of data preprocessing pipeline search as a bi-level optimization problem. To solve this problem efficiently, we transform and relax the discrete, non-differential search space into a continuous and differentiable one, which allows us to perform the pipeline search using gradient descent with training the ML model only once. Our experiments show that DIFFPREP achieves the best test accuracy on 15 out of the 18 real-world datasets evaluated and improves the model's test accuracy by up to 6.6 percentage points.

CCS Concepts: • **Information systems** → **Data cleaning**; • **Computing methodologies** → **Machine learning**.

Additional Key Words and Phrases: data cleaning, data preprocessing, automated machine learning

ACM Reference Format:

Peng Li, Zhiyi Chen, Xu Chu, and Kexin Rong. 2023. DiffPrep: Differentiable Data Preprocessing Pipeline Search for Learning over Tabular Data. *Proc. ACM Manag. Data* 1, 2, Article 183 (June 2023), 26 pages. <https://doi.org/10.1145/3589328>

1 INTRODUCTION

Machine learning (ML), in particular, supervised ML is increasingly being used for solving challenging real-world problems in a wide range of fields, such as medicine [19], finance [9], politics [34], etc. The workflow of developing an ML application may vary from projects but it typically involves four stages: data acquisition, data preprocessing, model training, and model evaluation [16].

Data preprocessing is an essential step in a typical ML workflow because in practice, the raw data collected often contain data issues and can rarely be directly used by ML models [13]. For example, data errors such as missing values and outliers will significantly reduce the ML model performance if not cleaned [24]; certain models like k-nearest neighbors (KNN) [32] expect data

Authors' addresses: Peng Li, pengli@gatech.edu, Georgia Institute of Technology, Atlanta, GA, USA; Zhiyi Chen, Georgia Institute of Technology, Atlanta, GA, USA, zchen798@gatech.edu; Xu Chu, Georgia Institute of Technology, Atlanta, GA, USA, xu.chu@cc.gatech.edu; Kexin Rong, Georgia Institute of Technology, Atlanta, GA, USA, krong@gatech.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

2836-6573/2023/6-ART183

<https://doi.org/10.1145/3589328>

in different columns to have a similar range, which requires the raw data to be normalized. In addition, real-world data often have multiple such data issues [2]: a dataset may contain missing values, outliers, and a large discrepancy on feature scales. To handle this case, we may first impute missing values with mean imputation, then remove outliers using the Z-score method, and finally normalize the data using standardization. As a result, data preprocessing usually involves multiple operators organized using a *data preprocessing pipeline*, where the operators in the pipeline are applied sequentially and each operator transforms the data to tackle a specific data issue.

Designing data preprocessing pipelines is challenging for data scientists as it involves many design decisions on *transformation types*, *order*, and *operators* [7, 14]. First, data scientists must decide which types of transformations (e.g., outlier removal, discretization, normalization) are needed and the order of different transformations in the pipeline. For example, outlier removal may or may not be needed and can be applied before or after normalization. For each type of transformation, there are multiple choices of operators—for example, standardization, min-max scaling, and robust scaling are all commonly used operators for normalization. Data scientists need to decide which operator to use for each transformation. Furthermore, different features may need to be preprocessed differently, which requires data scientists to design a *feature-wise pipeline* rather than using the same preprocessing pipeline for all features. For example, when repairing missing values in a dataset, some features may prefer mean imputation while others may prefer median imputation, depending on their distributions.

Even for experienced data scientists, it is usually not clear how to design a preprocessing pipeline that will lead to the best performance. Making such decisions heavily relies on domain knowledge, including the characteristics of the data, the types of downstream ML models, and the data scientists' experience. Traditional data cleaning works usually seek to design pipelines that optimize data quality independently of downstream applications [7]. However, since the ground-truth clean data of real-world datasets are rarely available, the data quality may not be accessible or accurately estimated. Moreover, previous works have shown that data cleaning or preprocessing without considering downstream ML models can sometimes negatively impact the performance of ML models [24, 27]. In practice, data scientists often use the time-consuming trial-and-error method to design preprocessing pipelines, which is reported to account for 80% of data scientists' time [35], or simply use some default configurations which usually result in suboptimal performance.

To reduce human effort in ML development, extensive study has been made on automated machine learning (AutoML) systems. Existing AutoML systems like Azure AutoML [4] and H2O.ai [23] can automatically perform data preprocessing and model training without too much human involvement. Most AutoML systems consist of a search space and some optimization methods [16]. The search space is defined by a set of possible parameters (choices) in ML development workflows. Our discussions in this paper focus on data preprocessing-related parameters, such as the choices of transformation types and operators. The optimization methods are used to automatically select a combination of parameters from the search space that leads to the highest model performance. With respect to data preprocessing pipelines, we find that existing AutoML solutions show limitations in both the size of search space and the efficiency of the optimization method:

Limitation 1: Limited search space for data preprocessing pipelines. The search space of an AutoML system determines the upper bound of its model performance. In general, larger search spaces are more likely to contain high-performing configurations, which could lead to better performance [8]. Despite the importance of data preprocessing in ML workflows, we found that existing AutoML frameworks only consider a limited search space of parameters for data preprocessing pipelines. In Figure 1, we rank existing AutoML systems according to the size of the search space they considered, which is determined by four variables: the choices of transformation

Systems	Operator	Types	Order	Feature-wise	Optimization Method
H2O	✗	✗	✗	✗	Random Search
Azure	✓	✗	✗	✗	Bayesian Optimization
Auto-Sklearn	✓	✓	✗	✗	Bayesian Optimization
Learn2Clean	✓	✓	✓	✗	Q-Learning
DiffPrep-Fix	✓	✓	✗	✓	Bi-level Optimization with Gradient Descent
DiffPrep-Flex	✓	✓	✓	✓	

Fig. 1. Data Preprocessing in AutoML Systems

types, the choices of transformation order, the choices of transformation operators, and whether the same pipeline is used for each data feature. H2O only supports a single default preprocessing pipeline with fixed transformation operators, types and order. Azure would search for the best operator for some transformations, such as normalization, but the transformation types and order are fixed. Auto-Sklearn [11] can determine the types of transformation needed and select suitable operators, but it has a fixed order of applying transformations. Learn2Clean [5] considers operators, types and order in its search space, but it uses the same pipeline to preprocess all features. Therefore, existing AutoML systems have only explored a small portion of the entire design space of data preprocessing pipelines, which limits the performance gains of downstream ML models.

Limitation 2: Low efficiency on optimization methods. As the search space becomes larger, it is increasingly important that the optimization methods can efficiently identify parameters with good performance. Unfortunately, most existing AutoML systems use optimization methods like random search or Bayesian optimization (Figure 1), which require training the ML model multiple times and do not scale well with the search space. Specifically, random search randomly samples parameters from the entire search space and trains the ML model with each set of sampled parameters; Bayesian optimization builds a probabilistic model that maps parameters to model performance, which requires iteratively training the ML model with new optimal parameters and updating the probabilistic model with the model performance. The optimization becomes particularly challenging if we want to support a larger search space in data preprocessing, such as using feature-wise pipelines. For example, if there are n possible pipelines for one feature, the number of possible pipelines for data with c columns is n^c , which means the space grows exponentially with the number of features. When there are many features or the ML model is large, random search or Bayesian optimization can be computationally expensive and time-consuming.

Our proposal. In this work, we propose DIFFPREP, an automatic data preprocessing method that can efficiently select high-quality data preprocessing pipelines for any given dataset and differentiable ML model. Unlike traditional data preprocessing or cleaning methods that focus on improving data quality independently of the downstream applications [7], DIFFPREP co-optimizes data preprocessing with model training: the goal is *to select data preprocessing pipelines that maximize the validation accuracy (or minimize the validation loss) of the ML model trained on the dataset*. Since the model training aims at minimizing the training loss, while the pipeline selection aims at minimizing the validation loss, we end up with a so-called *bi-level optimization problem* [26].

Different from existing AutoML systems, DIFFPREP considers feature-wise data preprocessing pipelines in addition to transformation type, operators, and order. Feature-wise pipelines significantly broaden the search space for data preprocessing compared to existing AutoML systems, since the number of possible pipelines grows exponentially with the number of features. We consider two usage scenarios with DIFFPREP. If users provide a pre-defined order, our method (DiffPrep-Fix)

will fix the order and automatically select the types and operators to generate a pipeline for each feature. If no order is provided, our method (DiffPrep-Flex) will automatically select the order, type and operator to generate a pipeline for each feature. It would explore the entire design space of data preprocessing pipelines, which is the largest search space in Figure 1. In both scenarios, we have empirically found that the larger search space improves the quality of the resulting pipelines.

Efficiently searching over this large search space is challenging: since the search space is discrete and non-differentiable, we need to enumerate each option and train ML models repeatedly to find high-quality pipelines. Our key insight is to make the ML model performance differentiable with respect to preprocessing pipeline choices so that we can leverage efficient optimization methods like gradient descent. To do so, we first parameterize the search space of data preprocessing pipelines such that each choice in the pipeline can be represented using binary parameters. We then relax the search space to be continuous using the softmax function and Sinkhorn normalization to make the pipeline differentiable. This allows DIFFPREP to use gradient descent as the optimization method to solve the bi-level optimization problem, which allows us to optimize the pipeline and model simultaneously with training the ML model only once.

Contributions. We make the following contributions in this paper.

- We propose DIFFPREP, the first automatic data preprocessing method to consider the design space of transformation types, operators, order and feature-wise pipelines.
- We formalize the problem of automatic data preprocessing as a bi-level optimization problem and use gradient descent to solve the bi-level optimization problem efficiently.
- We conduct experiments on 18 real-world datasets to evaluate the effectiveness of our method.

The results show that our method achieves the best test accuracy on 15 out of 18 datasets and improves the test accuracy by up to 6.6 percentage points.

Organization. The rest of this paper is organized as follows. Section 2 introduces preliminaries for constructing a data preprocessing pipeline and formally defines our studied problem. In Section 3, we discuss our method to automatically search a data preprocessing pipeline given a fixed order of transformations. In Section 4, we present our approach with flexible order of transformations. In Section 5, we show the experimental results. We discuss the related work in Section 6 and conclude the paper in Section 7.

2 PRELIMINARY

2.1 Transformation Operators and Types

Transformation Operators. Generally, a data preprocessing operator/algorithm, which we term as a *transformation operator (TF operator)*, is a function defined by $f : \mathcal{X}_1 \rightarrow \mathcal{X}_2$, where an original feature in feature space \mathcal{X}_1 , is mapped to a transformed feature in feature space \mathcal{X}_2 . In this paper, we focus on TF operators where both input and output are scalars, which we refer to as *feature-wise transformation operators*. Other TF operators that transform a vector into another vector, such as principal component analysis, and entity resolution operators [6, 42] that take a vector as input, are not considered in this paper. We note that most TF operators used in practice fall into this category. For example, among the TF operators provided by *scikit-learn* [31] in `sklearn.preprocessing` module, 14 out of 18 are feature-wise TF operators. Some TF operators have parameters. For example, to use the Z-Score method, we need to specify the threshold k to determine whether a value is an outlier (e.g., $k = 3$).

Transformation Types. Based on the purpose of the transformation, TF operators can be grouped into different types, which we refer to as *transformation types (TF types)*. For example, missing value imputation is a TF type that consists of TF operators such as mean imputation and median

Table 1. Example of TF Types and TF operators

Transformation Types	Transformation Operators	
Missing Value Imputation	Numerical Features	Mean Median Mode
	Categorical Features	Most Frequent Value Dummy Variable
Normalization	Standardization Min-Max Scaling Robust Scaling Max Absolute Scaling	
Outlier Removal	Z-Score (k) MAD (k) IQR (k)	
Discretization	Uniform (n) Quantile (n)	

imputation. Table 1 shows some examples of TF types and the TF operators for each TF type. Formally, we define a TF type F as a set of TF operators, denoted by $F = \{f_i : i = 1, 2, \dots\}$.

Design Considerations for Transformation Operators. The possibilities for transformation operators are much greater than what we cover in this paper. For example, it is possible to define custom conditions for selecting specific subsets of the dataset and to use custom functions to transform them. However, it is difficult to search in such a vast search space. Furthermore, when co-optimizing data preprocessing and model training, using a large number of operators that can modify data *arbitrarily* can increase the risk of model overfitting, as observed in prior work [20]. Therefore, DIFFPREP focuses on a limited set of operators that are commonly supported by ML frameworks (e.g. scikit-learn) and widely used. We have found empirically that this set of operators already leads to improved performance on many real-world datasets. Additionally, the operators supported by DIFFPREP are designed for general purpose to transform data based on some prior knowledge rather than modifying the dataset arbitrarily. For instance, operators for outlier removal and missing value imputation can only affect specific (usually small) subsets of the dataset, which are usually prone to be dirty data, while normalization and discretization operators can only scale and shift the distribution of entire features, but not completely distort the data distribution.

2.2 Data Preprocessing Pipeline Construction

As we mentioned before, data preprocessing usually involves multiple transformations that are combined in a pipeline and each feature can use a different pipeline. For simplicity, let us first assume that the data only contain one feature and we want to construct a data preprocessing pipeline for it. Figure 2 shows a typical data scientists' workflow to construct a data preprocessing pipeline, which involves the following steps.

Step 1: Data exploration. The initial step in creating a data preprocessing pipeline is usually *data exploration*. During this step, data scientists examine the data to understand its characteristics and identify any potential or existing issues. Some issues, such as missing values, are relatively easy to detect, while others, such as outliers, may require more involved data analysis [15].

Step 2: Prototype selection. The second step is called *prototype selection*, where based on the data issues, data scientists would select the TF types to be involved and decide the order of TF types to be applied [14, 33, 37]. For example, we choose to first impute missing values, then remove outliers and finally perform data normalization. Note that in a pipeline, the input of one operator is the output of its previous operator. Therefore, different orders of TF types can result in totally different outputs. For example, if outlier removal occurs before normalization, not only the input data for

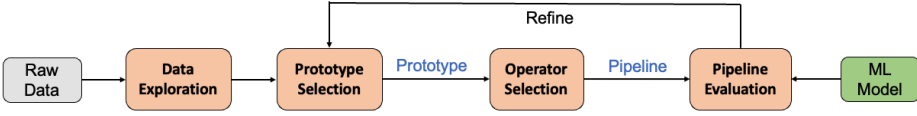


Fig. 2. A typical data scientists' workflow to construct a data preprocessing pipeline

normalization are changed, but also the statistics used to normalize the data (e.g., the minimum and maximum value of the column) are affected due to the removal of outliers, which can yield a significantly different output compared with having outlier removal after normalization. However, in practice, it is usually not clear how to decide the order of TF types and data scientists would simply choose some default order based on their experience. The outcome of this step is a *data preprocessing prototype*, which is an ordered sequence of TF types formally defined as follows.

Definition 2.1. (Data Preprocessing Prototype). Let $\mathcal{S} = \{F\}$ denote the space of TF types. We define a data preprocessing prototype \mathcal{T} as $\mathcal{T} = \{T_i \in \mathcal{S} : i = 1, 2, \dots\}$, where T_i is the i -th TF type to be applied and selected from the space.

Note that in the literature, the notion of data preprocessing prototype is also known as logical pipeline plan [37] or pipeline prototype [33], which is defined as a directed acyclic graph of TF types. This definition suggests that there is no repetition of a TF type in a prototype [14]. This is because in practice, a type of transformation would rarely be repeated (e.g., it does not make sense to impute missing values twice). Following this widely-used definition of prototype, we require that $\forall p \neq q, T_p \neq T_q$. We want to point out that this assumption does not prevent users from using a TF type multiple times. If repetition is really needed, we can define two TF types T_1 and T_2 on the same transformation and consider them as two different TF types. This allows for repeated transformation types in a prototype.

Step 3: Operator selection. The third step is *operator selection*, where given a data preprocessing prototype, data scientists would select a specific TF operator for each TF type in the prototype. For example, we use mean imputation to impute missing values, use the Z-score approach to remove outliers and perform standardization to normalize data. To select suitable operators for TF types, data scientists have to consider not only the characteristics of data but also the downstream ML model. For example, for normalization, if the downstream model is k-nearest neighbors, min-max scaling is usually preferable as it transforms all the columns into the same scale. In contrast, if the model is logistic regression, standardization may be better as it makes convergence faster. However, such heuristic rules may not work for every dataset, and in practice, selecting suitable operators usually requires trial-and-error. By selecting a TF operator for each TF type, we can instantiate a *data preprocessing pipeline*, which is formally defined as follows.

Definition 2.2. (Data Preprocessing Pipeline). Given a data preprocessing prototype \mathcal{T} , we define a data preprocessing pipeline $\mathcal{G}_{\mathcal{T}}$ as $\mathcal{G}_{\mathcal{T}} = \{g_i : i = 1, 2, \dots\}$, where $g_i \in T_i$ is the TF operator selected for T_i in the prototype and the i -th operator to be applied in the pipeline.

Step 4: Pipeline evaluation. The final step is *pipeline evaluation*, where data scientists would use the pipeline to transform the raw data and evaluate the performance of the pipeline by training and testing the end ML model on the transformed data. To transform the raw data with the pipeline, we can sequentially apply the TF operator in the pipeline. Formally, let x denote the raw feature of one example and x_i be the transformed feature after i steps transformation, where $x_0 = x$. Then, x_i can be computed by applying g_i (the i -th TF operator in the pipeline) on its input x_{i-1} as: $x_i = g_i(x_{i-1})$.

The output of the pipeline is the output of the last TF operator in the pipeline. Formally, let s denote the number of TF operators in the pipeline and $\mathcal{G}_{\mathcal{T}}(x)$ denote the final output of the pipeline $\mathcal{G}_{\mathcal{T}}$. Then, we have $\mathcal{G}_{\mathcal{T}}(x) = x_s$.

After pipeline evaluation, data scientists may refine the prototype and the pipeline based on the evaluation results. For example, they may add/delete TF types in the prototype, change the order of TF types, replace the TF operator for some TF types, etc. The above process will be repeated until the end ML model achieves desired performance. As a result, constructing data preprocessing pipelines is an iterative process that requires substantial domain knowledge and heavily relies on human experts to make decisions, which can be time-consuming and costly.

2.3 Problem Statement

The core idea of our approach is to formulate the decision-making process of data scientists as an optimization problem. Assume that each data example has c features denoted by $\mathbf{x} = [x^1, x^2, \dots, x^c]$. Let $\mathcal{T}^i, \mathcal{G}_{\mathcal{T}}^i$ denote the prototype and the pipeline for the i -th feature. Let \mathbf{w} denote the parameters of the ML model $h_{\mathbf{w}} : \mathcal{X}^c \rightarrow \mathcal{Y}$. The model will take features of an example as input and generate a prediction $\hat{y} = h_{\mathbf{w}}(x^1, x^2, \dots, x^c)$. Let $loss(\hat{y}, y)$ be the loss function that returns a loss score given the prediction \hat{y} and the ground truth label y . Then, the training loss and validation loss on the transformed data can be computed as:

$$L_{train}(\mathcal{G}_{\mathcal{T}}^1, \dots, \mathcal{G}_{\mathcal{T}}^c, \mathbf{w}) = \sum_{\mathbf{x}, y \in D_{train}} loss(h_{\mathbf{w}}(\mathcal{G}_{\mathcal{T}}^1(x^1), \dots, \mathcal{G}_{\mathcal{T}}^c(x^c)), y)$$

$$L_{val}(\mathcal{G}_{\mathcal{T}}^1, \dots, \mathcal{G}_{\mathcal{T}}^c, \mathbf{w}) = \sum_{\mathbf{x}, y \in D_{val}} loss(h_{\mathbf{w}}(\mathcal{G}_{\mathcal{T}}^1(x^1), \dots, \mathcal{G}_{\mathcal{T}}^c(x^c)), y)$$

Problem Statement. Given a training set D_{train} and a validation set D_{val} , a space of TF operators and TF types \mathcal{S} , and a set of ML model parameters \mathbf{w} , we would like to find a pipeline $\mathcal{G}_{\mathcal{T}}^i$ (with its prototype \mathcal{T}^i) from the space for each feature x^i , such that the performance of the ML model trained and evaluated on the transformed data is maximized. This *data preprocessing pipeline search problem (DPPS)* can be formulated as an optimization problem:

$$\min_{\mathcal{G}_{\mathcal{T}}^1, \dots, \mathcal{G}_{\mathcal{T}}^c} L_{val}(\mathcal{G}_{\mathcal{T}}^1, \dots, \mathcal{G}_{\mathcal{T}}^c, \mathbf{w}^*)$$

$$\text{s.t. } \mathbf{w}^* = \arg \min_{\mathbf{w}} L_{train}(\mathcal{G}_{\mathcal{T}}^1, \dots, \mathcal{G}_{\mathcal{T}}^c, \mathbf{w})$$

This is called a *bi-level optimization problem* in which one optimization problem is embedded within another [26]. In the inner optimization, we fix the preprocessing pipeline parameters $\mathcal{G}_{\mathcal{T}}^1, \dots, \mathcal{G}_{\mathcal{T}}^c$ and focus on finding the best model parameters \mathbf{w}^* to minimize the training loss on the transformed training data. In the outer optimization, we fix the model parameters and focus on finding the best pipeline parameters to minimize the validation loss on the transformed validation data. An alternative problem formulation is to optimize both the pipeline and model parameters to minimize the training loss, which becomes a one-level optimization. However, previous works (e.g., DARTS [25]) have shown that one-level optimization would make it easier to overfit the training data. Therefore, in this work, we use bi-level optimization to reduce the risk of overfitting. In our experiments, we empirically verified that the bi-level approach helps improve the model test accuracy compared to using a one-level optimization (Section 5.4).

The most straightforward way to solve the above bi-level optimization problem is to train the downstream ML model by minimizing $L_{train}(\mathcal{G}_{\mathcal{T}}^1, \dots, \mathcal{G}_{\mathcal{T}}^c, \mathbf{w})$ for each possible data preprocessing pipeline in the space, and then select the one that has the minimal validation loss $L_{val}(\mathcal{G}_{\mathcal{T}}^1, \dots, \mathcal{G}_{\mathcal{T}}^c, \mathbf{w}^*)$. This naive approach is computationally expensive as it requires training the downstream ML models

as many times as the number of possible data preprocessing pipelines. Consider a prototype with s TF types, where each TF type consists of m TF operators. Under this prototype, there are m^s possible pipelines by choosing different TF operators for each TF type. Since different features can use different pipelines, with c features, we have m^{sc} possible pipelines, which is exponential to the number of TF types and the number of features. This is only the number of possible pipelines under one prototype. If we consider using different prototypes, the space will become even larger. Therefore, this naive approach is infeasible in practice.

3 DATA PREPROCESSING WITH FIXED PROTOTYPE

Let us first assume that we have a fixed pre-defined data processing prototype $\mathcal{T} = \{T_1, \dots, T_s\}$ that is used for all features. This is a common scenario where data scientists would like to skip the prototype selection step and use a default prototype so that they can spend more time on operator selection. This is also the setup of many existing systems [5, 10, 33, 40] for automatic data preprocessing, where the prototype is pre-defined by users and fixed in advance.

Given this prototype, we need to assign each TF type in the prototype with a TF operator to generate a pipeline and we need to generate a pipeline for each feature. As the number of possible choices is exponentially large, the problem is: *how to efficiently find the optimal assignment for each feature such that the validation loss is minimized?* This search problem is challenging because the search space is discrete and non-differentiable, which means we have to enumerate every possible case to find the optimal one. To solve this problem, we first parameterize the search space such that each assignment of TF operators (i.e., each choice of pipelines) can be represented using binary pipeline parameters (Section 3.1). Then, we relax the search space to be continuous so that the model loss is differentiable w.r.t the pipeline parameters (Section 3.2). The relaxation enables us to solve the bi-level optimization problem efficiently using gradient descent (Section 3.3).

3.1 Parameterization

Without loss of generality, assume that each TF type T_i in the prototype consists of m TF operators denoted by $T_i = \{f_{i1}, f_{i2}, \dots, f_{im}\}$ (if different TF types contain a different number of operators, m denotes the maximum number of TF operators). To instantiate a pipeline, we need to select one specific TF operator for each TF type. Each selection can be represented using a $s \times m$ matrix $\beta = \{\beta_{ij}\}, \beta_{ij} \in \{0, 1\}$ is defined as follows:

$$\beta_{ij} = \begin{cases} 1 & f_{ij} \text{ is selected} \\ 0 & \text{Otherwise} \end{cases} \quad (1)$$

In other words, β_{ij} is 1 if we select f_{ij} for T_i , otherwise it is 0. Note that because only one TF operator is selected for each TF type, there is exactly one element each row in the matrix to be 1 and 0s elsewhere. Hence, we have $\sum_j \beta_{ij} = 1$.

Example 3.1. Assume that the given prototype consists of three TF types, i.e., $\mathcal{T} = \{T_1, T_2, T_3\}$. Each TF type consists of four possible TF operators, i.e., $T_i = \{f_{i1}, f_{i2}, f_{i3}, f_{i4}\}$. Consider the data preprocessing pipeline $\mathcal{G}_{\mathcal{T}} = \{f_{12}, f_{21}, f_{34}\}$. The selection of TF operators that generates this pipeline can be represented using the following matrix.

$$\beta = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

Each β matrix uniquely defines a data preprocessing pipeline under the given prototype and we can use β matrix to compute the output of the pipeline. Let x denote the raw feature of one example and x_i be the transformed feature after i steps transformation, where $x_0 = x$. Then, we

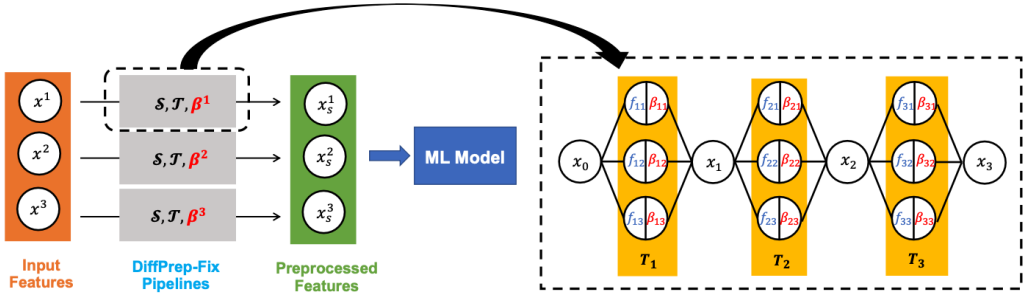


Fig. 3. DiffPrep-Fix Architecture

can compute x_i as:

$$x_i = \sum_{j=1}^m \beta_{ij} f_{ij}(x_{i-1}) \quad (2)$$

Since only β_{ij} associated with the selected TF operator is equal to 1 and others are 0, Equation 2 returns exactly the output of the selected TF operator. The final output of the pipeline x_s becomes a variable parameterized by β matrix denoted by $x_s = x(\beta)$. Assume that the dataset has c features and let β^i denote the parameters that defines the pipeline for the i -th feature x^i . Then, the training loss and the validation loss also become variables parameterized by all β matrices, which can be represented as:

$$L_{train}(\beta^1, \dots, \beta^c, \mathbf{w}) = \sum_{\mathbf{x}, y \in D_{train}} \text{loss}(h_{\mathbf{w}}(x^1(\beta^1), \dots, x^c(\beta^c)), y) \quad (3)$$

$$L_{val}(\beta^1, \dots, \beta^c, \mathbf{w}) = \sum_{\mathbf{x}, y \in D_{val}} \text{loss}(h_{\mathbf{w}}(x^1(\beta^1), \dots, x^c(\beta^c)), y) \quad (4)$$

We can now rewrite the DPPS problem statement with a fixed prototype using β matrices as follows, where the search space of pipelines is converted into a space of β matrices.

$$\min_{\beta^1, \dots, \beta^c} L_{val}(\beta^1, \dots, \beta^c, \mathbf{w}^*) \quad (5)$$

$$\text{s.t. } \mathbf{w}^* = \text{argmin}_{\mathbf{w}} L_{train}(\beta^1, \dots, \beta^c, \mathbf{w}) \quad (6)$$

Architecture. Figure 3 shows the architecture of DiffPrep with a fixed prototype. For each input feature, DiffPrep will learn a set of β parameters that defines a data preprocessing pipeline. Each β_{ij} parameter is attached with a TF operator indicating whether this operator is selected. The final output of DiffPrep pipelines will be the preprocessed features parameterized by β parameters, which will be fed into ML models for training and evaluation.

3.2 Differentiable Relaxation

Although we have parameterized the search space of pipelines using β matrices, the search space is still discrete as β matrices are binary parameters, i.e., $\beta_{ij} \in \{0, 1\}$. To make the search space continuous, we relax β to be continuous parameters that can take continuous values from 0 to 1, i.e., $\beta_{ij} \in [0, 1]$. However, we need to retain the constraints between parameters in the original binary matrix ($\sum_j \beta_{ij} = 1$) such that these parameters are semantically meaningful. To enforce that each row sums up to 1, we can define β_{ij} using a softmax function as:

$$\beta_{ij} = \frac{\exp(\tau_{ij})}{\sum_k \exp(\tau_{ik})} \quad (7)$$

where $\tau = \{\tau_{ij}\} \in \mathbb{R}^{s \times m}$ are underlying parameters.

Algorithm 1 Solving Bi-level Optimization with Gradient Descent

- 1: Initialize τ and \mathbf{w}
 - 2: **while** *not converged* **do**
 - 3: Update τ : $\tau = \tau - \eta_1 \nabla_{\tau} L_{val}(\beta(\tau), \mathbf{w} - \eta_2 \nabla_{\mathbf{w}} L_{train}(\beta(\tau), \mathbf{w}))$
 - 4: Update \mathbf{w} : $\mathbf{w} = \mathbf{w} - \eta_2 \nabla_{\mathbf{w}} L_{train}(\beta(\tau), \mathbf{w})$
 - 5: **return** $\beta(\tau), \mathbf{w}$
-

By enforcing these constraints, we can interpret β_{ij} (Equation 1) as the probability that f_{ij} is selected for F_i . We can still use Equation 2 to compute the transformed data, which can be interpreted as the expected value of the transformed data. Similarly, we can compute the expected value of training loss and validation loss using Equation 3 and Equation 4.

3.3 Bi-level Optimization

Making β continuous allows us to solve the bi-level optimization problem efficiently using gradient descent instead of enumerating all possible β . To minimize the validation loss (Equation 5), we can iteratively update the underlying parameters τ using the gradient of the validation loss with respect to τ as follows:

$$\tau = \tau - \eta_1 \nabla_{\tau} L_{val}(\beta(\tau), \mathbf{w}^*) \quad (8)$$

where η_1 is the learning rate for updating τ . However, to obtain the optimal model \mathbf{w}^* in Equation 8, we need to solve the inner optimization problem (Equation 6) by completely training an ML model until convergence for every update of τ , which is computationally expensive. To solve this issue, instead of finding the optimal \mathbf{w}^* , we approximate it by doing only a single training step, which is a one-step gradient descent on the current model parameters \mathbf{w} (denoted by \mathbf{w}') as:

$$\mathbf{w}^* \approx \mathbf{w}' = \mathbf{w} - \eta_2 \nabla_{\mathbf{w}} L_{train}(\beta(\tau), \mathbf{w}) \quad (9)$$

where η_2 is the learning rate for model training¹. Similar approximation can also be found in previous work [12, 25] for solving other bi-level optimization problems. We summarize the procedure of solving the bi-level optimization (Equation 5 and Equation 6) using gradient descent in Algorithm 1, where the pipeline parameters $\beta(\tau)$ and model parameters \mathbf{w} are alternately updated using gradient descent until convergence.

Gradient Computation. In Algorithm 1, $\nabla_{\mathbf{w}} L_{train}(\beta(\tau), \mathbf{w})$ can be derived from Equation 3, where β can be considered as constant. However, since β (Equation 7) and \mathbf{w}' (Equation 9) are variables of τ , to compute the gradients $\nabla_{\tau} L_{val}(\beta(\tau), \mathbf{w}'(\tau))$, we need to use the chain rule²:

$$\nabla_{\tau} L_{val}(\beta(\tau), \mathbf{w}'(\tau)) = \nabla_{\beta} L_{val}(\beta, \mathbf{w}') \cdot \nabla_{\tau} \beta - \eta_2 \nabla_{\mathbf{w}'} L_{val}(\beta, \mathbf{w}') \cdot \nabla_{\mathbf{w}, \beta} L_{train}(\beta, \mathbf{w}) \cdot \nabla_{\tau} \beta \quad (10)$$

We can decompose Equation 10 into three parts: D_1, D_2, D_3 , where $D_1 = \nabla_{\tau} \beta$, $D_2 = \nabla_{\beta} L_{val}(\beta, \mathbf{w}')$ and $D_3 = \nabla_{\mathbf{w}'} L_{val}(\beta, \mathbf{w}') \cdot \nabla_{\mathbf{w}, \beta} L_{train}(\beta, \mathbf{w})$. The gradient D_1 can be derived from Equation 7. The gradient D_2 is also easy to compute, where we consider \mathbf{w}' as constant and compute the gradient of validation loss w.r.t. β . This can be derived from Equation 4 or as we will show later, using backpropagation through the pipeline. However, computing D_3 is difficult since it involves a second-order derivative and matrix-vector product computation. Following the previous work on neural network architecture search [25], we approximate it using numerical differentiation. Let ϵ be a small scalar and $\mathbf{w}^{\pm} = \mathbf{w} \pm \epsilon \nabla_{\mathbf{w}'} L_{val}(\beta, \mathbf{w}')$, then D_3 can be estimated as:

$$D_3 \approx \frac{\nabla_{\beta} L_{train}(\beta, \mathbf{w}^+) - \nabla_{\beta} L_{train}(\beta, \mathbf{w}^-)}{2\epsilon} \quad (11)$$

¹In practice, to avoid tuning two learning rates, we can simply set $\eta_1 = \eta_2$ [25].

²Chain rules for multivariable functions: Suppose that $x(t)$ and $y(t)$ are differentiable functions of t and $z = f(x(t), y(t))$ is a differentiable function of x and y . Then the chain rule states $\frac{dz}{dt} = \frac{\partial z}{\partial x} \frac{\partial x}{\partial t} + \frac{\partial z}{\partial y} \frac{\partial y}{\partial t}$.

$\nabla_{\beta} L_{train}(\beta, \mathbf{w}^+)$, $\nabla_{\beta} L_{train}(\beta, \mathbf{w}^-)$ are the gradients of the training loss w.r.t β , where we consider \mathbf{w}^+ and \mathbf{w}^- as constant. They can be derived from Equation 3 or as we will show later using backpropagation through the pipeline. Therefore, to update the underlying parameters τ , we need to compute three gradients w.r.t. β with different model parameters: $\nabla_{\beta} L_{train}(\beta, \mathbf{w}^+)$, $\nabla_{\beta} L_{train}(\beta, \mathbf{w}^-)$ and $\nabla_{\beta} L_{val}(\beta, \mathbf{w}')$.

Backpropagation through the pipeline. The gradients of loss w.r.t. β can be computed using backpropagation through the pipeline. Let $\frac{\partial L}{\partial x_s}$ be the derivative of the loss w.r.t. the input of the ML model or the final output of the preprocessing pipeline. From Equation 2, using the chain rule, we have

$$\frac{\partial L}{\partial \beta_{ij}} = \frac{\partial L}{\partial x_i} f_{ij}(x_{i-1}), \quad \frac{\partial L}{\partial x_{i-1}} = \frac{\partial L}{\partial x_i} \sum_{j=1}^m \beta_{ij} \frac{\partial f_{ij}}{\partial x_{i-1}} \quad (12)$$

However, in this equation, the gradient term $\frac{\partial f_{ij}}{\partial x_{i-1}}$, namely, the gradient of the output of a TF operator w.r.t. its input depends on the internals of a TF operator, which may not be easy to compute. Also, we expect users to add their own customized TF operators, but users may not be able to derive the gradients. Therefore, we assume that all TF operators are black-box functions, where we only have access to the output and input without knowing their internal algorithms, and we approximate the gradients using numerical differentiation. Let ϵ be a small scalar. This gradient can be estimated as:

$$\frac{\partial f_{ij}}{\partial x_{i-1}} \approx \frac{f_{ij}(x_{i-1} + \epsilon) - f_{ij}(x_{i-1} - \epsilon)}{2\epsilon} \quad (13)$$

Then, we can use Equation 12 to backpropagate the gradients of the loss w.r.t. the output of each TF operator and each β_{ij} .

Implementation with automatic differentiation. Many automatic differentiation engines (e.g. Pytorch [30], TensorFlow [1]) can perform backpropagation automatically and efficiently to compute the gradients. However, these tools usually require that all the computations in the forward propagation are differentiable and implemented using their frameworks. Since we consider TF operators as black-box functions of which the internal algorithms are unknown, we cannot leverage these tools directly to backpropagate the preprocessing pipeline. To solve this issue, we modify the original forward propagation of the pipeline (Equation 2) to be:

$$x_i = \sum_{j=1}^m \beta_{ij} \tilde{o}_{ij} + x_{i-1} \sum_{j=1}^m \tilde{\beta}_{ij} \tilde{d}_{ij} - \tilde{x}_{i-1} \sum_{j=1}^m \tilde{\beta}_{ij} \tilde{d}_{ij} \quad (14)$$

where $\tilde{o}_{ij} = f_{ij}(x_{i-1})$ is the output of the black-box TF operator; $\tilde{d}_{ij} = \frac{\partial f_{ij}}{\partial x_{i-1}}$ is the numerical derivative computed using Equation 13; \tilde{x}_{i-1} is a *snapshot* of x_{i-1} , which is a constant number that has the same value as x_{i-1} but does not require gradient (e.g., $\tilde{x}_{i-1} = x_{i-1}.\text{detach}()$ in PyTorch or $\tilde{x}_{i-1} = \text{stop_gradient}(x_{i-1})$ in TensorFlow); $\tilde{\beta}_{ij}$ is a snapshot of β_{ij} . Note that all the variables with tilde are constant numbers that do not have gradients.

In the forward pass, Equation 14 yields exactly the same outputs as Equation 2, but it only requires the output of the TF operators and the numerical derivatives computed using the input and output of TF operators. Therefore, the internal implementation of TF operators is not involved, which enables the automatic backpropagation to be performed. In the backward pass, the automatic differentiation engines will compute the gradients as:

$$\frac{\partial L}{\partial \beta_{ij}} = \frac{\partial L}{\partial x_i} \tilde{o}_{ij}, \quad \frac{\partial L}{\partial x_{i-1}} = \frac{\partial L}{\partial x_i} \sum_{j=1}^m \tilde{\beta}_{ij} \tilde{d}_{ij} \quad (15)$$

Algorithm 2 DiffPrep-Fix

Require: Space of TF types and operators S , pre-defined prototype \mathcal{T} , training set D_{train} , validation set D_{val}

Ensure: Optimal pipeline parameters β and model parameters w

- 1: Initialize τ and w
 - 2: **while** *not converged* **do**
 - 3: Fit TF operators on the transformed training data
 - 4: Forward Propagation: Compute $L_{train}(\beta, w^+)$, $L_{train}(\beta, w^-)$, $L_{val}(\beta, w')$
 - 5: Backward Propagation: Compute $\nabla_{\beta} L_{train}(\beta, w^+)$, $\nabla_{\beta} L_{train}(\beta, w^-)$, $\nabla_{\beta} L_{val}(\beta, w')$
 - 6: Compute $\nabla_{\tau} L_{val}(\beta, w')$
 - 7: Update τ : $\tau = \tau - \eta_1 \nabla_{\tau} L_{val}(\beta, w')$
 - 8: Update w : $w = w - \eta_2 \nabla_w L_{train}(\beta(\tau), w)$
 - 9: **return** $\beta(\tau), w$
-

This yields exactly the same gradients as Equation 12 with the gradients of black-box TF operators replaced by approximate numerical gradients. Therefore, this allows us to backpropagate the preprocessing pipeline correctly and automatically using any automatic differentiation engines.

Algorithm. Algorithm 2 shows the pseudocode of DiffPrep with a fixed prototype. We first initialize the underlying pipeline parameters and model parameters (Line 1). Most TF operators need to be fitted before performing transformation. For example, standardization needs to compute the mean and standard deviation of the input data. Therefore, at each iteration, we first fit all the TF operators using the transformed training data (Line 3). Note that each TF operator should be fitted on its input training data, which is the output data of its previous step in the pipeline (i.e., the operator f_{ij} should be fitted on x_{i-1} over all training data). Since x_{i-1} depends on β parameters (Equation 14), it will be changed every iteration as β parameters are updated. Therefore, we need to refit TF operators at the beginning of every iteration. Then we can perform forward propagation using Equation 14 (Line 4) and backward propagation using automatic differentiation (Line 5). We compute the gradients needed for updating underlying pipeline parameters τ using Equation 10 and Equation 11 (Line 6). We update the underlying pipeline parameters and model parameters alternatively using gradient descent (Line 7 - 8). This will be repeated until convergence.

Complexity. The number of β parameters needed is $s \times m \times c$, where s is the number of TF types, m is the number of TF operators for each TF type and c is the number of features. The running time is dominated by the forward propagation and backward propagation. Without DiffPrep, the training process only needs to update the model parameters (Line 8), which requires one pass of forward propagation and one pass of backpropagation. With DiffPrep, updating β needs three more passes of forward propagation (Line 4) and three more passes of backpropagation (Line 5). To improve the efficiency of the algorithm, we fit TF operators (Line 3) and update $\beta(\tau)$ (Line 4 - 7) with a mini-batch randomly sampled from the training/validation set, instead of using all training/validation examples, which is similar to using stochastic gradient descent in place of gradient descent. We apply this strategy in our experiments.

4 DATA PREPROCESSING WITH FLEXIBLE PROTOTYPE

We now discuss the more general and flexible case, where the data preprocessing prototype is not pre-defined. In such a case, prior to operator selection using the method we introduced in Section 3, we need to select a prototype from space for each feature. A prototype is an ordered sequence of TF types. Therefore, for prototype selection, we need to determine (1) the TF types that are included in the prototype and (2) the order of TF types in the sequence.

To simplify this search problem, we introduce an *identity transformation operator* (I) into each TF type, which is defined as $I(x) = x$. The identity TF operator is a function that maps data to itself, so selecting the identity TF operator for a TF type during operator selection is equivalent to dropping this TF type in the prototype. This allows us to include all TF types in the prototype and only focus on their order, rather than deciding which to include or exclude. A prototype becomes a permutation of all TF types with the option to exclude any by selecting the identity operator for it.

Now, the problem is: *how to find the optimal prototype for each feature such that the optimal pipeline under the optimal prototype minimizes the validation loss?* A naive approach is to enumerate every possible prototype and use the method introduced in Section 3 to find the optimal pipeline for each prototype. However, this approach is not feasible because the search space of prototypes is large. Given a space that contains s TF types, the number of possible permutations is $s!$. Since different features can use different prototypes, for a dataset with c features, there are $s!^c$ possible combinations, which is exponential to the number of features.

The hardness of the above search problem is again because the search space of prototypes is discrete and non-differentiable. To solve this problem efficiently, we use a method similar to that introduced in Section 3: we first parameterize the search space using a set of binary prototype parameters (Section 4.1) and then make the search space to be continuous and differentiable (Section 4.2). Finally, using the method introduced in Section 3.3, we can solve the bi-level optimization using gradient descent, which enables our method to learn the prototype parameters, pipeline parameters and model parameters simultaneously and efficiently in one training loop.

4.1 Parameterization

By introducing identity TF operators, given a space of TF types $\mathcal{S} = \{F_1, F_2, \dots, F_s\}$, a prototype $\mathcal{T} = \{T_1, T_2, \dots, T_s\}$ would be a permutation of \mathcal{S} . Hence, we can represent a prototype using a $s \times s$ permutation matrix $\alpha = \{\alpha_{ij}\}$, $\alpha_{ij} \in \{0, 1\}$ defined as follows.

$$\alpha_{ij} = \begin{cases} 1 & T_i = F_j \\ 0 & \text{Otherwise} \end{cases} \quad (16)$$

In other words, α_{ij} is 1 if the i -th TF type in the prototype is F_j , otherwise it is 0. Note that the permutation matrix has exactly one element of 1 each row and each column and 0s elsewhere, i.e., the sum of each row and each column is exactly 1. Hence, we have $\sum_i \alpha_{ij} = 1$ and $\sum_j \alpha_{ij} = 1$.

Example 4.1. Assume that the space consists of three TF types $\mathcal{S} = \{F_1, F_2, F_3\}$. Consider the prototype $\mathcal{T} = \{F_2, F_3, F_1\}$. We have $\alpha_{12} = 1$, $\alpha_{23} = 1$, $\alpha_{31} = 1$ and 0s elsewhere. Hence, This prototype can be represented using the following α matrix.

$$\alpha = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{bmatrix}$$

To generate a pipeline, in addition to the prototype, we also need to select a TF operator for each TF type. Let $F_i = \{f_{i1}, f_{i2}, \dots, f_{im}\}$, where f_{ij} is the j -th TF operator for the TF type F_i . Then, we can still use Equation 1 to define the β matrix, which represents the result of operator selection. An α matrix and a β matrix together uniquely defines a data preprocessing pipeline. The data after i steps transformation can be computed as:

$$x_i = \sum_{j=1}^s \sum_{k=1}^m \alpha_{ij} \beta_{jk} f_{jk}(x_{i-1}) \quad (17)$$

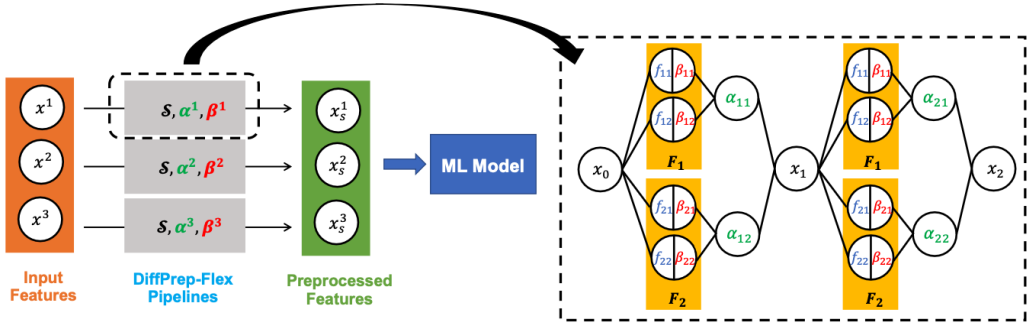


Fig. 4. DiffPrep-Flex Architecture

Note that in Equation 17, only the TF type F_j selected for the i -th step has $\alpha_{ij} = 1$ and only the TF operator selected for F_j has $\beta_{jk} = 1$. Therefore, the summation is exactly the output of the selected TF operator for the i -th TF type in the prototype.

Similar to Equation 3 and 4, we can also compute the training loss $L_{train}(\alpha^1, \dots, \alpha^c, \beta^1, \dots, \beta^c, w)$ and validation loss $L_{val}(\alpha^1, \dots, \alpha^c, \beta^1, \dots, \beta^c, w)$ using Equation 17, which will be parameterized by α, β matrices. We can also rewrite the DPPS problem statement using α, β matrices as:

$$\min_{\alpha^1, \dots, \alpha^c, \beta^1, \dots, \beta^c} L_{val}(\alpha^1, \dots, \alpha^c, \beta^1, \dots, \beta^c, w^*) \quad (18)$$

$$\text{s.t. } w^* = \operatorname{argmin}_w L_{train}(\alpha^1, \dots, \alpha^c, \beta^1, \dots, \beta^c, w) \quad (19)$$

Architecture. Figure 4 shows the architecture of DiffPrep with a flexible prototype. For each input feature, DiffPrep will learn a set of α and β parameters that defines a data preprocessing pipeline. The α parameters define the prototype and the β parameters define the results of operator selection. The final output of DiffPrep-Flex pipelines will be the preprocessed features parameterized by α, β parameters, which will be fed into ML models.

4.2 Differentiable Relaxation

We can reuse the method described in Section 3.2 to relax β parameters. For α matrix, similar to relaxing β matrix, we can relax α to be continuous variables $\alpha_{ij} \in [0, 1]$, but we need to retain the constraints between parameters in the original binary matrix ($\sum_i \alpha_{ij} = 1, \sum_j \alpha_{ij} = 1$). In other words, α matrix requires to be a non-negative squared matrix with both rows and columns summing up to 1, which is so-called a *doubly stochastic matrix* (DSM). To enforce that α is a DSM, we generate it using *Sinkhorn normalization*. Sinkhorn [38, 39] showed that any non-negative square matrix can be converted into a DSM by repeatedly and alternatively normalizing its rows and columns. Cruz et al. [36] introduced *Sinkhorn Layer* that converts CNN predictions to a DSM using Sinkhorn normalizations. Following Adams et al. [3], we define Sinkhorn normalization over any squared matrix X as:

$$C_{ij}(X) = \frac{X_{ij}}{\sum_i X_{ij}}, R_{ij}(X) = \frac{X_{ij}}{\sum_j X_{ij}}, S(X) = \lim_{l \rightarrow \infty} S^l(X)$$

$$S^l(X) = \begin{cases} X_{ij} & l = 1 \\ R(C(S^{l-1}(X))) & l > 1 \end{cases}$$

where R and C are the row and column normalization. $S(X)$ is the result of Sinkhorn normalization, which is a DSM.

Inspired by these work, we define α matrix as $\alpha = S(\theta)$, where $\theta = \{\theta_{ij}\} \in \mathbb{R}_+^{s \times s}$ are non-negative underlying parameters. Note that this normalization is a differentiable process, which

Algorithm 3 DiffPrep-Flex

Require: Space of TF types and operators S , training set D_{train} , validation set D_{val}

Ensure: Optimal prototype parameters α , pipeline parameters β , and model parameters w

- 1: Initialize θ , τ and w
- 2: **while** *not converged* **do**
- 3: Fit TF operators on the transformed training data
- 4: Forward Propagation: Compute $L_{train}(\alpha, \beta, w^+)$, $L_{train}(\alpha, \beta, w^-)$, $L_{val}(\alpha, \beta, w')$
- 5: Backward Propagation: Compute $\nabla_{\alpha} L_{train}(\alpha, \beta, w^+)$, $\nabla_{\alpha} L_{train}(\alpha, \beta, w^-)$, $\nabla_{\alpha} L_{val}(\alpha, \beta, w')$,
 $\nabla_{\beta} L_{train}(\alpha, \beta, w^+)$, $\nabla_{\beta} L_{train}(\alpha, \beta, w^-)$, $\nabla_{\beta} L_{val}(\alpha, \beta, w')$
- 6: Compute $\nabla_{\tau} L_{val}(\alpha, \beta, w')$, $\nabla_{\theta} L_{val}(\alpha, \beta, w')$
- 7: Update τ , θ : $\tau = \tau - \eta_1 \nabla_{\tau} L_{val}(\alpha, \beta, w')$, $\theta = \theta - \eta_1 \nabla_{\theta} L_{val}(\alpha, \beta, w')$
- 8: Update w : $w = w - \eta_2 \nabla_w L_{train}(\alpha, \beta, w)$
- 9: **return** $\alpha(\theta)$, $\beta(\tau)$, w

allows us to compute the gradients w.r.t the underlying parameters. By enforcing these constraints, we can interpret α_{ij} (Equation 16) as the probability that F_j is the i -th transformation type in the prototype, and β_{ij} (Equation 1) as the probability that f_{ij} is selected for F_i . In this sense, Equation 17 will return the expected value of the transformed data, which can be used to compute the expected training and validation loss.

4.3 Bi-level Optimization

Now we can use gradient descent to solve the bi-level optimization (Equation 18 and 19) and find the optimal prototype and pipeline represented by α and β parameters, similar to Algorithm 2.

Algorithm. Algorithm 3 shows the pseudocode of DiffPrep with flexible prototypes. Compared to Algorithm 2, the input only requires a space of TF types and TF operators, and does not need a pre-defined prototype. At each iteration, we compute the gradients of the validation loss with respect to both α and β using the method described in Section 3.3 (Line 4 - 6). Then we update the prototype parameters, pipeline parameters and model parameters (Line 7 - 8) until convergence.

5 EXPERIMENTS

We conduct extensive experiments to evaluate the effectiveness and efficiency of DIFFPREP. All our experiments were performed on a single machine with a 2.20GHz Intel Xeon(R) Gold 5120 CPU. The source code of our experiments is available at <https://github.com/chu-data-lab/DiffPrep>.

5.1 Experimental Setup

Datasets. We evaluate our approaches on 18 diverse real-world datasets from the OpenML data repository [41]. These datasets are frequently used in the AutoML and data cleaning literature [5, 20, 23, 24, 28]. Table 2 shows the size and number of label classes of each dataset, as well as the number of missing values and outliers detected using the Z-score method (i.e., values more than 3 standard deviations away from the mean of the column are considered outliers). We randomly split each dataset into training/validation/test sets by a ratio of 60%/20%/20%.

Model. For all preprocessing methods considered in the experiments, we use Logistic Regression as the downstream ML model, which is a classical differentiable model commonly used in practice and frequently considered in the literature [10, 22, 29, 40].

Search Space. We consider four TF types in our experiments, including missing value imputation, normalization, outlier removal and discretization. For each TF type, we select several TF operators that are widely used in practice and provided by scikit-learn [31] as shown in Table 1. For operators

that have parameters, we discretize them with different parameters. For example, for the Z-score method, we consider Z-score(2), Z-score(3) and Z-score(4) as 3 different TF operators. We also add an identity operator to each TF type which amounts to a skipping operator. Since most TF operators provided by scikit-learn require the input data to contain no missing values, we do not add the identity operator to the missing value imputation. We enforce the missing value imputation to be the first TF type in the prototype so that missing values are imputed in the first step. In addition, many operators only accept numerical input data. Therefore, after missing value imputation, we transform all categorical features into numerical features using one-hot encoding.

Methods Compared. We compare the following data preprocessing approaches.

- *DiffPrep-Fix (DP-Fix)*. This is our approach with a pre-defined fixed transformation order. We use the order $\{missing\ value\ imputation, normalization, outlier\ removal, discretization\}$. Our method learns the best β parameters that define a pipeline for each feature. Note that different features can have different β parameters and thus can be preprocessed differently.
- *DiffPrep-Flex (DP-Flex)*. This is our approach with a flexible transformation order. It will automatically learn α parameters that define the transformation order and β parameters that define a pipeline under the learned order for each feature.
- *Default (DEF)*. This is the default data preprocessing pipeline that is commonly used in practice and adopted in AutoML frameworks such as H2O [23]. It first imputes numerical missing values with the mean value of the column and categorical missing values with the most frequent value of the column. Then, it normalizes each feature using standardization. Note that this default pipeline does not include outlier removal and discretization, which is the default setup for many existing AutoML frameworks [4, 23]. All features use the same pipeline.
- *RandomSearch (RS)*. This approach searches for a pipeline by training ML models with randomly sampled pipelines N times and selecting the one with the best validation accuracy. We use the same transformation order as DiffPrep-Fix and sample one TF operator for each TF type from the space (Table 1) to generate a sampled pipeline. All features are preprocessed using the same pipeline. We set the number of trials $N = 20$.
- *Auto-Sklearn (AS)* [10]. This is an open-source AutoML package with automatic data preprocessing. Its built-in data preprocessor performs missing value imputation, removing low-variance features and normalization for numerical features, and performs category shift, missing value imputation, minority coalescence and one-hot encoding for categorical features. The built-in data preprocessor already contains operators like category coalescence, category shift, TF-IDF encoding that are not included in our space. It uses Bayesian optimization to search the optimal pipeline that will be used to preprocess all the features. The time limit for search is set to 1 hour and default settings are used for other configurations. AS also provides a feature processor that performs feature engineering operators, such as feature extraction and feature embedding, which are out of the scope of data preprocessing. Since we focus on data preprocessing and feature engineering is not performed for all other methods, for a fair comparison, we turn off the feature processor and only use the data processor. We will report the results of AS with its feature processor later in Section 5.6, where we will show that our methods can also be combined with a feature processor.
- *Learn2Clean (LC)* [5]. This is a reinforcement-learning-based data cleaning and preparation method proposed in the literature. It uses Q-learning to select the optimal preprocessing pipeline that maximizes the ML model performance. The TF types, operators and order are all flexible, but it only generates one pipeline to preprocess all the features in the same way. We modify the search space of TF operators to be the same as ours (Table 1).

- *BoostClean (BC)* [20]. This is an automatic data cleaning approach that uses boosting to select an ensemble of data cleaning operators (or pipelines) from a pre-defined candidate set to maximize validation accuracy. BoostClean combines the ML models trained on different transformed data induced by different pipelines, thus it needs to train multiple models on different candidate pipelines. BoostClean does not support feature-wise pipelines. To generate candidate pipelines, we use the same transformation order as DiffPrep-Fix and randomly sample 50 pipelines from the space (Table 1). BoostClean is designed for binary classification but can be adapted for multi-class classification by breaking it down into multiple binary tasks using the one-vs-all method. We adopt this method in our experiments. We set its ensemble size $B = 5$, the best setup it reported.

Training Process. We use an SGD optimizer to optimize the model parameters and use an Adam optimizer to optimize the pipeline parameters. The learning rate is tuned using the validation set. The batch size is 512 and the number of training epochs is 1000. We keep track of the validation loss during training and report the result at the epoch with the minimum validation loss. The training and evaluation are implemented using PyTorch [30], which utilizes parallelism by default.

Evaluation Metrics. Our goal is to automatically and efficiently compose a data preprocessing pipeline to maximize downstream ML model performance. Therefore, we use the test accuracy of the model and the end-to-end running time as the evaluation metrics.

5.2 Performance Comparison

Table 2. Comparison of model test accuracy on 18 real-world datasets using different data preprocessing pipelines. DIFFPREP (DP-Fix and DP-Flex combined) achieves the best test accuracy on 15 out of 18 datasets.

Dataset	Data Characteristics					Test Accuracy						
	#Ex.	#Feat.	#classes	#MVs	#Out.	DEF	RS	AS	LC	BC	DP-Fix	DP-Flex
abalone	4177	9	28	0	200	0.24	0.243	0.216	0.186	0.168	0.238	0.255
ada_prior	4562	15	2	88	423	0.848	0.844	0.853	0.816	0.848	0.854	0.846
avila	20867	11	12	0	4458	0.553	0.598	0.615	0.597	0.585	0.638	0.63
connect-4	67557	43	3	0	45873	0.659	0.671	0.667	0.658	0.69	0.732	0.701
eeg	14980	15	2	0	209	0.589	0.658	0.657	0.641	0.659	0.678	0.677
google	9367	9	2	1639	109	0.586	0.627	0.664	0.549	0.616	0.645	0.641
house	1460	81	2	6965	617	0.928	0.938	0.945	0.812	0.928	0.932	0.945
jungle_chess	44819	7	3	0	0	0.668	0.669	0.678	0.676	0.667	0.682	0.682
micro	20000	21	5	0	8122	0.564	0.579	0.584	0.582	0.561	0.586	0.588
mozilla4	15545	6	2	0	290	0.855	0.922	0.931	0.854	0.93	0.923	0.922
obesity	2111	17	7	0	25	0.775	0.841	0.737	0.723	0.652	0.893	0.896
page-blocks	5473	11	5	0	1011	0.942	0.959	0.969	0.92	0.951	0.957	0.967
pbseq	1945	19	2	1445	99	0.71	0.73	0.712	0.704	0.72	0.725	0.743
pol	15000	49	2	0	8754	0.884	0.879	0.877	0.737	0.903	0.904	0.919
run_or_walk	88588	7	2	0	8548	0.719	0.829	0.851	0.728	0.835	0.907	0.917
shuttle	58000	10	7	0	5341	0.964	0.996	0.998	0.997	0.997	0.998	0.997
uscensus	32561	15	2	4262	2812	0.848	0.84	0.851	0.786	0.848	0.857	0.852
wall-robot-nav	5456	25	4	0	1871	0.697	0.872	0.869	0.69	0.9	0.898	0.914

Accuracy Comparison. Table 2 shows the test accuracy of the ML model on 18 real-world datasets with different preprocessing methods. We have the following observations from Table 2:

- (1) Different preprocessing pipelines can lead to significantly different model performances. For example, on wall-robot-nav, using different pipelines, the accuracy differs by more than 20%.
- (2) DIFFPREP (DP-Fix and DP-Flex combined) achieves the best test accuracy on 15 out of 18 datasets. Our methods surpass the best baseline method by more than 1% on 9 datasets. In particular, on run_or_walk, obesity and connect-4, our methods outperform the best baseline method by 6.6%, 5.5% and 4.2%, respectively. This shows the effectiveness of DIFFPREP and the significant

improvement of model performance by using a larger search space. The reasons for gains vary by datasets. For example, connect-4 may benefit from cleaning outliers properly as it has a great number of outliers. On this dataset, we found DP-Fix selects Z-score methods for some features, while selecting skip operators for some other features. This shows the benefits of DiffPrep by using feature-wise pipelines, where we can clean true outliers in some features without affecting other features that may not contain true outliers. We will explain the contribution of each component of DiffPrep in more detail in the ablation study (Section 5.4).

- (3) DiffPrep-Fix (DP-Fix) and DiffPrep-Flex (DP-Flex) respectively outperform the best baseline methods on 11 and 13 datasets. Comparing DiffPrep-Flex with DiffPrep-Fix, DiffPrep-Flex outperforms DiffPrep-Fix on 9 out of 18 datasets, but the two methods differ by less than 2% on 17 datasets. DiffPrep-Flex does not significantly improve upon DiffPrep-Fix despite the larger search space, because the fixed order we used for DiffPrep-Fix is already the optimal order for many datasets. Hence, changing the order in DiffPrep-Flex may not lead to further improvement. Also, as the gradient descent is not always guaranteed to find the global optimum, it is possible to find a less optimal solution, especially for DiffPrep-Flex which has a harder optimization problem due to its flexibility. However, DiffPrep-Fix with default order is insufficient as it may not work well for all datasets. As we will show in the ablation study, if the order is not set properly, DiffPrep-Fix may perform worse, while DiffPrep-Flex can avoid such cases. In addition, the default order will be invalid if users add custom TF types. The benefit of DiffPrep-Flex is that it does not require and is not affected by the pre-defined transformation order.
- (4) RandomSearch (RS) performs better than Default on 15 out of 18 datasets. Especially on wall-robot-nav and run_or_walk, the accuracy is improved by 17.5% and 11%, respectively. This indicates that using the same default preprocessing pipeline for all datasets is generally not a good strategy, despite its wide adoption in practice.
- (5) BoostClean (BC) performs worse than DiffPrep (DP-Fix and DP-Flex combined) on 17 out of 18 datasets. Comparing BC with RandomSearch (RS), on binary classification tasks, BoostClean (BC) outperforms RS on 6 out of 9 datasets. This is because BoostClean uses an ensemble of models and statistical boosting to improve model performance. Also, BoostClean has a larger candidate set with 30 additional pipelines compared to RandomSearch, increasing its chance of finding a better pipeline. However, on datasets with more than 2 classes, BoostClean only outperforms RandomSearch on 3 out of 9 datasets. Especially on obesity and abalone, BoostClean is 18.9% and 7.5% worse than RandomSearch. This is because BoostClean uses the one-vs-all method to handle multi-class classification, which may cause class imbalance issue and thus lead to worse model performances when the number of classes is large.
- (6) Although Auto-Sklearn and RandomSearch use different optimization methods, they perform close to each other on most datasets, where the difference is less than 2% on 14 out of 18 datasets. This is because they use a similar search space. In comparison, our methods consider a larger space with feature-wise pipelines and flexible order, which further improves the performance.
- (7) Learn2Clean (LC) on average performs worst and it outperforms Default only on 6 datasets. Although Learn2Clean has a larger search space than other baseline methods, the Q-learning method is not effective to find a good pipeline and thus it performs worse than other methods.

Running Time Comparison. To compare running time, we categorize all the datasets into different bins based on the number of rows of each dataset (e.g., 0-10K, 10-20K) and we report the average end-to-end running time of different methods over datasets in each bin as shown in Figure 5 (a). As expected, Default takes the shortest time since it only trains the model once with the fixed pipeline. RandomSearch (RS) is about 20 times slower than Default, since it simply trains the model 20 times with 20 randomly sampled pipelines. Auto-Sklearn (AS) takes about 1 hour constantly as it keeps

searching for the best result within the given 1 hour time limit. BoostClean (BC) is about 50 times slower than Default on binary class datasets as it needs to train a model on every candidate pipeline, but it is much slower on multi-class datasets as it uses the one-vs-all method and needs to rerun the whole algorithm for each subclass. Additionally, while both RandomSearch and BoostClean can utilize parallelism to decrease running time by training multiple models simultaneously on multiple machines, this study only considers the single machine setup, which is more common for general users. Overall, DiffPrep-Fix takes about half of the time of RandomSearch, which amounts to 10 times slower than Default and it is generally faster than Learn2Clean. DiffPrep-Flex takes similar time as RandomSearch and is generally faster than Auto-Sklearn. Also, we can see the running time of our methods grows linearly with the size of the dataset, showing their scalability.

Note that although DiffPrep only requires training the model once, the additional overhead comes from the forward/backward propagation through the pipelines. Recall that DiffPrep uses “dynamic” pipelines that are changed at every iteration. Therefore, we need to reinvoke TF operators to transform data at every iteration. The ratio between our method and the Default method depends on the time it takes to propagate through the pipelines and through the model. Our experiments use logistic regression, in which propagation through the model is relatively fast. However, as the model sizes increase, the ratio between our methods and the Default will go down. For example, Figure 5 (b) shows that with a larger ML model (two-layer neural network with 100 neurons), DiffPrep-Fix is 2-3 times slower and DiffPrep-Flex is 6-7 times slower than Default, both of which are much faster than other baseline methods.

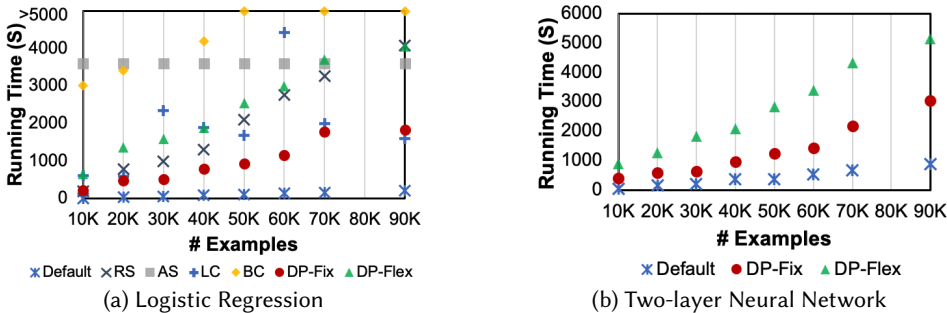


Fig. 5. Comparison of end-to-end running time of different preprocessing methods.

5.3 Sensitivity Analysis

Non-linear ML models. One concern for our methods is that the approximate gradient computation (e.g., Equation 11) may not work well for non-linear ML models. To understand the impact of end ML models, we test our methods with a non-linear ML model: a two-layer neural network with ReLU activation and 100 neurons in the hidden layer. As shown in Table 3, compared with other methods, DIFFPREP (DP-Fix and DP-Flex combined) achieves the best test accuracy on 10 out of 18 datasets. Especially on pbseq and avila, our method outperforms the best baseline method by 3.1% and 3%, respectively. This shows that our method is still effective with non-linear models in practice. This corroborates with findings in DARTS [25], which uses similar approximate gradient computation for neural network architecture search and empirically works for non-linear models.

Impact of validation size on overfitting. Our methods aim to minimize validation loss, but this could lead to model overfitting on the validation set, especially when the validation set is small. To understand how the validation size affects the performance, we first hold out 20% data as the test set and then randomly split the remaining data into validation/training sets with ratios ranging from 1% to 99%. We repeat it 5 times and run DiffPrep-Fix with each split. Figure 6 shows the

Table 3. Comparison of model test accuracy using a two-layer neural network.

Dataset	DEF	RS	AS	LC	BC	DP-Fix	DP-Flex
abalone	0.251	0.266	0.251	0.271	0.162	0.246	0.273
ada_prior	0.829	0.845	0.855	0.815	0.829	0.834	0.84
avila	0.82	0.914	0.864	0.824	0.928	0.946	0.958
connect-4	0.802	0.799	0.804	0.795	0.8	0.807	0.806
eeg	0.906	0.955	0.933	0.834	0.955	0.924	0.939
google	0.576	0.657	0.667	0.573	0.657	0.677	0.666
house	0.928	0.935	0.921	0.928	0.935	0.908	0.928
jungle_chess	0.854	0.856	0.852	0.825	0.842	0.853	0.861
micro	0.626	0.626	0.63	0.617	0.615	0.628	0.626
mozilla4	0.916	0.934	0.932	0.925	0.933	0.935	0.934
obesity	0.91	0.95	0.941	0.815	0.927	0.941	0.943
page-blocks	0.966	0.967	0.975	0.983	0.959	0.969	0.959
pbseq	0.743	0.743	0.748	0.697	0.743	0.779	0.779
pol	0.989	0.967	0.988	0.99	0.988	0.981	0.984
run_or_walk	0.985	0.986	0.99	0.984	0.983	0.99	0.991
shuttle	0.999	0.998	0.999	0.999	0.999	1	0.999
uscensus	0.856	0.845	0.854	0.79	0.856	0.854	0.853
wall-robot-nav	0.898	0.961	0.962	0.92	0.957	0.965	0.962

mean and standard error of the test accuracy for some datasets. When the validation is extremely small (1%), although the training set is large, the test accuracy is generally bad. This is because the pipeline parameters optimizing on the validation set make our method overfit the validation set. From 1% to 25%, the test accuracy improves significantly as overfitting is reduced. From 25% to 50%, the gain becomes insignificant on most datasets because the validation set is large enough to prevent overfitting. However, as we further increase the validation size, the test accuracy drops on most datasets. Especially, when the validation set becomes extremely large (99%), the accuracy is generally bad as the ML model overfits the small training set. Therefore, we can conclude that using either a too small validation set or a too small training set can result in overfitting issues and using 60:20 train/val split (i.e., 25% ratio in the plot) is good enough in practice. The results on other datasets and on DiffPrep-Flex reveal similar findings and thus are omitted here.

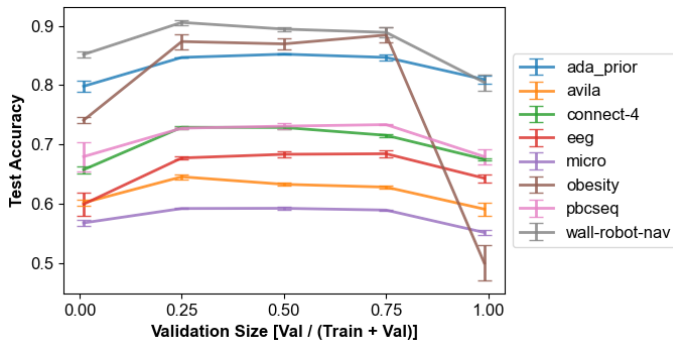


Fig. 6. Using either a too small validation set or a too small training set can result in poor test accuracy due to overfitting. 60: 20 train/val split is generally good enough in practice.

5.4 Ablation Study

Contribution of feature-wise pipelines. To study the benefits of using feature-wise pipelines, we turn off feature-wise pipelines on DiffPrep-Fix by letting all features share the same β parameters.

Table 4. Results of Ablation Study

Dataset	DiffPrep-Fix	DiffPrep-Flex	DiffPrep-Fix (w/o feature-wise)	DiffPrep-Fix (worst order)	DiffPrep-Fix (train-opt)
abalone	0.238	0.255	0.237	0.238	0.243
ada_prior	0.854	0.846	0.851	0.848	0.849
avila	0.638	0.63	0.613	0.637	0.637
connect-4	0.732	0.701	0.688	0.719	0.676
eeg	0.678	0.677	0.662	0.648	0.675
google	0.645	0.641	0.621	0.601	0.64
house	0.932	0.945	0.942	0.914	0.938
jungle_chess	0.682	0.682	0.683	0.682	0.681
micro	0.586	0.588	0.576	0.586	0.584
mozilla4	0.923	0.922	0.921	0.854	0.923
obesity	0.893	0.896	0.879	0.863	0.9
page-blocks	0.957	0.967	0.952	0.945	0.964
pbseq	0.725	0.743	0.722	0.71	0.722
pol	0.904	0.919	0.888	0.882	0.896
run_or_walk	0.907	0.917	0.862	0.888	0.898
shuttle	0.998	0.997	0.998	0.998	0.996
uscensus	0.857	0.852	0.853	0.858	0.857
wall-robot-nav	0.898	0.914	0.905	0.882	0.894

This way, our method can still learn the optimal β parameters but all features will be processed in the same way. The result is shown in Table 4 as DiffPrep-Fix (w/o feature-wise). Without feature-wise pipelines, the test accuracy drops on 14 out of 18 datasets. Especially on run_or_walk and connect-4, the accuracy is reduced by 4.5% and 4.4%, respectively. This shows that using feature-wise pipelines can indeed improve the performance as it significantly enlarges the search space. Note that compared with the baseline methods in Table 2, even without feature-wise pipelines, our method still achieves the highest accuracy on 7 out of 18 datasets. This is because our method uses continuous parameters, which amounts to using a combination of different TF operators for each TF type and thus results in a larger space than choosing a single TF operator for each TF type.

Benefits of flexible order. Compared with DiffPrep-Fix that uses a pre-defined transformation order, DiffPrep-Flex can search for the optimal order automatically with no need of user input. This is essential for less experienced users as they may not know how to set the order properly. To understand the benefit of using flexible order, we run DiffPrep-Fix with different transformation orders and we show the performance with the worst order for each dataset in Table 4 as DiffPrep-Fix (worst order). With a bad order, the performance of DiffPrep-Fix drops significantly: on mozilla4, for example, the performance is reduced by 6.9%. Compared to DiffPrep-Fix (worst order), DiffPrep-Flex performs better on 13 out of 18 datasets. This suggests that for inexperienced users, DiffPrep-Flex is a better method as it can avoid performance degradation due to bad pre-defined order.

Alternative Optimization Objective. To understand the benefits of bi-level optimization, we experiment with an alternative, one-level optimization objective, where both pipeline parameters (α , β) and model parameters (w) are learned to minimize the training loss. This optimization problem can be solved by computing the gradient of the training loss w.r.t each parameter and updating α , β , w simultaneously using gradient descent. We run DiffPrep-Fix using this alternative objective and the results are shown in Table 4 as DiffPrep-Fix (train-opt). Compared with this alternative method, our original method (DiffPrep-Fix) performs better on 12 out of 18 datasets and has the same accuracy on 2 datasets. We hypothesize that this is because optimizing pipelines and models jointly on training data makes it easier to overfit the training data and thus reduces the test accuracy. In addition, although our method achieves better accuracy in most cases, the difference

between the two methods is less than 1% on 17 out of 18 datasets. This is because even though the alternative objective only involves the training set, we tune hyperparameters (e.g., learning rates) using the validation set and adopt early stopping to mitigate the overfitting. Similar findings are also observed in DARTS [25], which uses a validation set to select network architecture.

5.5 Case study with Synthetic Datasets

Traditional data cleaning or preprocessing methods focus on improving the data quality without considering downstream ML models. However, in the context of data cleaning for ML, prior works [24, 27] have shown that this may not improve and can even downgrade the ML model performance. This motivates us to take ML models into consideration and use the model accuracy as guidance to select preprocessing pipelines. To justify this motivation, we randomly inject 10% missing values into three datasets. Since the ground truth of the synthetic datasets is known, we can evaluate the quality of missing value imputation using the root mean squared error (RMSE) between imputed values and ground truth values. Table 5 shows the model accuracy and the quality of missing value imputation using different preprocessing methods. Note that we omit data quality for BoostClean (BC) as it uses an ensemble of ML models trained on different imputed data. The results show that DiffPrep consistently achieves the best model accuracy on all synthetic datasets. However, the imputation methods selected by our methods are not necessarily the method with the best data quality. For example, on wal-robot-nav, the imputation quality of DiffPrep-Flex is the worst, however, it achieves better accuracy than other baseline methods. This indicates that selecting operators by data quality purely may not lead to the best model performance. Instead, we need to jointly consider other operators in the pipeline and the ML model to achieve the best result.

Table 5. Comparison of model accuracy and data cleaning quality on synthetic datasets with injected errors.

Dataset	Model Accuracy							Imputation Quality (RMSE)					
	DEF	RS	AS	LC	BC	DP-Fix	DP-Flex	DEF	RS	AS	LC	DP-Fix	DP-Flex
avila	0.521	0.577	0.575	0.536	0.546	0.630	0.614	1.043	1.05	1.043	1.05	0.696	0.8
eeg	0.577	0.646	0.637	0.521	0.641	0.661	0.663	1787.601	1788.85	1787.601	1788.85	2662.249	2000.434
wal-robot-nav	0.66	0.863	0.818	0.646	0.862	0.892	0.888	1.25	0.965	1.25	1.312	1.273	1.427

5.6 Beyond Data Preprocessing

While DiffPrep focuses on optimizing data preprocessing-related parameters, it can be easily integrated with AutoML methods that target other components in the ML pipeline, such as feature selection, feature embedding and hyperparameter tuning. In previous experiments, we have combined DiffPrep with learning rate tuning. In this example, we demonstrate how DiffPrep can work with feature extraction. Specifically, we train a random forest classifier on the raw input data as a features extractor, where we can use the tree node embeddings and predicted probability as the extracted features. We append the extracted features to the original data and feed the data into DiffPrep for data preprocessing and model training. Table 6 shows the changes in the number of features and test accuracy of DP-Fix with and without the feature extractor. The accuracy significantly improved on 16 out of 18 datasets due to feature extraction. For reference, we report the performance change of Auto-Sklearn (AS) with its feature processor turned on, which performs feature extraction and selection on top of data preprocessing. The accuracy is significantly improved on 11 out of 18 datasets. Note that the performance on some datasets become worse, as enabling feature processor can make the search harder due to the increase in training time and search space.

6 RELATED WORK

Automated Machine Learning. The traditional workflow of developing ML models requires significant domain knowledge and human effort, especially for data preprocessing and model

Table 6. Test accuracy combined with a feature preprocessor.

Dataset	DP-Fix w/o Feat vs. w/ Feat		AS w/o Feat vs. w/ Feat	
	# Feats Change	Test Accuracy Change	# Feats Change	Test Accuracy Change
abalone	9→424	0.238→0.26	9→851	0.216→0.237
ada_prior	15→408	0.854→0.848	15→120	0.853→0.846
avila	11→410	0.638→0.994	11→23760	0.615→0.94
connect-4	43→445	0.732→0.813	43→10454	0.667→0.808
eeg	15→333	0.678→0.93	15→679	0.657→0.91
google	9→388	0.645→0.673	9→174	0.664→0.674
house	81→472	0.932→0.935	81→47	0.945→0.928
jungle_chess	7→409	0.682→0.809	7→2106	0.678→0.855
micro	21→425	0.586→0.605	21→1771	0.584→0.627
mozilla4	6→407	0.923→0.947	6→416	0.931→0.94
obesity	17→420	0.893→0.943	17→629	0.737→0.886
page-blocks	11→391	0.957→0.967	11→176	0.969→0.969
pbccseq	19→415	0.725→0.807	19→103	0.712→0.692
pol	49→442	0.904→0.98	49→623	0.877→0.884
run_or_walk	7→401	0.907→0.989	7→12942	0.851→0.992
shuttle	10→394	0.998→0.999	10→220	0.998→0.925
uscensus	15→416	0.857→0.851	15→279	0.851→0.85
wall-robot-nav	25→428	0.898→0.994	25→1956	0.869→0.759

training, which can be time-consuming and expensive. Automated machine learning, also known as AutoML, aims to reduce the need of human involvement by automating the whole process of ML model development. Our work provides an effective and efficient solution to automate the process of data preprocessing.

Many existing AutoML tools enable automation in the process of data preprocessing and model training. Azure [4] is an AutoML system developed by Microsoft. It uses matrix factorization and Bayesian optimization to automatically select models and tune hyperparameters. However, for data preprocessing, it only considers tuning the normalization methods. For other transformations, it simply uses pre-defined default methods, such as mean imputation for missing value imputation. The order of transformations is fixed and different features are preprocessed in the same way. Auto-Sklearn [10] is an open-source AutoML package. It also relies on Bayesian optimization to automate data preprocessing and model training. Specifically, it adopts random-forest-based sequential model-based optimization to explore the search space. For data preprocessing, Auto-Sklearn considers tuning a larger set of transformations than Azure. However, it also uses a fixed order of transformation and preprocesses all features in the same way. In comparison, our methods allow flexible order of transformations and different features can use different pipelines. Instead of Bayesian optimization, we use bi-level optimization with gradient descent, which is faster and only needs to train the model once.

Data Cleaning for ML. Our work builds upon a line of data cleaning methods that incorporate signals from the downstream ML models into the design of cleaning objectives [27]. DiffML [17] is our closest work in this area since it adopts a similar idea of making ML pipelines differentiable so that preprocessing steps can be jointly learned with the ML model using backpropagation. However, DiffML considers each preprocessing step separately while we consider a pipeline of preprocessing steps and formulate the ordering using Sinkhorn. In addition, DiffML minimizes training loss with one-level optimization, while we minimize validation loss with bi-level optimization, which reduces the risk of overfitting. BoostClean [20] automatically selects cleaning algorithms from the search space via statistical boosting to maximize the ML model’s validation accuracy. It supports conditional cleaning operations defined by a combination of custom detection and repair functions.

However, only 192 to 1976 operators were used in the evaluation, which is much smaller than the number of pipelines evaluated in our search space. BoostClean does not consider feature-wise pipelines. AlphaClean [22] is a similar automatic data cleaning pipeline generation system that finds cleaning pipelines from data cleaning libraries to maximize user-defined quality measures. AlphaClean uses tree search algorithms and learns pruning heuristics to reduce the search space. Learn2Clean [5] uses reinforcement learning to select a sequence of data preprocessing operators such that the ML model performance is maximized. However, reinforcement learning needs to train and evaluate models many times to obtain rewards, which could lead to scalability challenges. ActiveClean [21] focuses on gradient-based models and prioritizes cleaning examples with higher gradients that are likely to have large impacts on the model. CPClean [18] quantifies the impact of data cleaning on ML models using the uncertainty of predictions and prioritizes cleaning examples that would lead to the maximum reduction on the uncertainty of predictions. Both ActiveClean and CPClean solve the problem of which examples should be cleaned, while our method is design to select appropriate preprocessing/cleaning operators. Data Cleaning and AutoML [28] investigated the impact of data cleaning for AutoML systems and found that data errors did not affect the AutoML performance significantly as AutoML can select robust models and adjust ML pipelines to handle data errors properly. This verifies our motivation to take downstream models into consideration when handling data errors rather than isolating data cleaning as a separate process.

Network Architecture Search. Network architecture search (NAS) aims to automate the design of neural network architectures. Although our work is not directly related to NAS, many of our ideas are inspired by DARTS [25], which is a differentiable NAS method. To search network architecture efficiently, DARTS first relaxes the categorical choice of network operators using continuous architecture parameters, which is similar to the pipeline parameters we used for choices of TF operators. It then formulates the architecture search as a bi-level optimization problem, which can be solved by updating the architecture and model parameters alternatively using gradient descent. Although DARTS and DiffPrep are similar in their methodology, there are two major differences. First, most architecture operators (e.g., convolution, ReLU) are differentiable and their gradients can be easily computed using automatic differentiation engines (e.g., PyTorch). However, the TF operators for data preprocessing can be complex and the gradients may not be easily computed. Second, the same type of TF operators rarely appears more than once in a pipeline and the order of TF operators (types) are usually flexible. In comparison, architecture operators usually repeat multiple times in an architecture but their order is usually fixed (e.g., DARTS uses ReLU-Conv-BN).

7 CONCLUSION

In this paper, we propose DIFFPREP, a method that can automatically and efficiently search the optimal data preprocessing pipeline for a given tabular dataset and an ML model such that the performance of the ML model is maximized. We formalize the problem of automatic data preprocessing as a bi-level optimization problem. We then relax the discrete search space using continuous parameters, which enables us to search optimal preprocessing pipelines efficiently using gradient descent. The experiments show that our method achieves the best test accuracy on 15 out of 18 real-world datasets and improves the model's test accuracy by up to 6.6 percentage point. We note that our proposed differentiable search strategy can only search for optimal data preprocessing pipelines when the end model is differentiable. We leave it to future work on how to search for the optimal data processing pipeline for non-differentiable end models, such as random forests.

ACKNOWLEDGMENTS

We thank the many members of the Georgia Tech Database Group for their valuable feedback on this work. This research was supported in part by Bosch Research North America.

REFERENCES

- [1] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. 2016. Tensorflow: a system for large-scale machine learning. In *Osdi*, Vol. 16. Savannah, GA, USA, 265–283.
- [2] Ziawasch Abedjan, Xu Chu, Dong Deng, Raul Castro Fernandez, Ihab F Ilyas, Mourad Ouzzani, Paolo Papotti, Michael Stonebraker, and Nan Tang. 2016. Detecting data errors: Where are we and what needs to be done? *Proceedings of the VLDB Endowment* 9, 12 (2016), 993–1004.
- [3] Ryan Prescott Adams and Richard S Zemel. 2011. Ranking via sinkhorn propagation. *arXiv preprint arXiv:1106.1925* (2011).
- [4] Microsoft Azure. [n. d.]. Azure AutoML. <https://azure.microsoft.com/en-us/services/machine-learning/automatedml/>. Accessed: April 12, 2023.
- [5] Laure Berti-Equille. 2019. Learn2clean: Optimizing the sequence of tasks for web data preparation. In *The World Wide Web Conference*. 2580–2586.
- [6] Riccardo Cappuzzo, Paolo Papotti, and Saravanan Thirumuruganathan. 2020. Creating embeddings of heterogeneous relational datasets for data integration tasks. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data*. 1335–1349.
- [7] Xu Chu, Ihab F Ilyas, Sanjay Krishnan, and Jiannan Wang. 2016. Data cleaning: Overview and emerging challenges. In *Proceedings of the 2016 International Conference on Management of Data*. ACM, 2201–2206.
- [8] Yuanzheng Ci, Chen Lin, Ming Sun, Boyu Chen, Hongwen Zhang, and Wanli Ouyang. 2021. Evolving search space for neural architecture search. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 6659–6669.
- [9] Yue Deng, Feng Bao, Youyong Kong, Zhiquan Ren, and Qionghai Dai. 2016. Deep direct reinforcement learning for financial signal representation and trading. *IEEE transactions on neural networks and learning systems* 28, 3 (2016), 653–664.
- [10] Matthias Feurer, Katharina Eggersperger, Stefan Falkner, Marius Lindauer, and Frank Hutter. 2020. Auto-Sklearn 2.0: Hands-free AutoML via Meta-Learning. *arXiv:2007.04074 [cs.LG]* (2020).
- [11] Matthias Feurer, Aaron Klein, Jost Eggersperger, Katharina Springenberg, Manuel Blum, and Frank Hutter. 2015. Efficient and Robust Automated Machine Learning. In *Advances in Neural Information Processing Systems 28 (2015)*. 2962–2970.
- [12] Luca Franceschi, Michele Donini, Paolo Frasconi, and Massimiliano Pontil. 2017. Forward and reverse gradient-based hyperparameter optimization. In *International Conference on Machine Learning*. PMLR, 1165–1173.
- [13] Salvador Garcia, Julián Luengo, and Francisco Herrera. 2015. *Data preprocessing in data mining*. Springer.
- [14] Joseph Giovannelli, Besim Bilalli, and Alberto Abelló Gamazo. 2021. Effective data pre-processing for AutoML. In *Proceedings of the 23rd International Workshop on Design, Optimization, Languages and Analytical Processing of Big Data (DOLAP): co-located with the 24th International Conference on Extending Database Technology and the 24th International Conference on Database Theory (EDBT/ICDT 2021): Nicosia, Cyprus, March 23, 2021*. CEUR-WS. org, 1–10.
- [15] Simon Hawkins, Hongxing He, Graham Williams, and Rohan Baxter. 2002. Outlier detection using replicator neural networks. In *International Conference on Data Warehousing and Knowledge Discovery*. Springer, 170–180.
- [16] Xin He, Kaiyong Zhao, and Xiaowen Chu. 2021. AutoML: A survey of the state-of-the-art. *Knowledge-Based Systems* 212 (2021), 106622.
- [17] Benjamin Hilprecht, Christian Hammacher, Eduardo Reis, Mohamed Abdelaal, and Carsten Binnig. 2022. DiffML: End-to-end Differentiable ML Pipelines. *arXiv preprint arXiv:2207.01269* (2022).
- [18] Bojan Karlaš, Peng Li, Renzhi Wu, Nezihe Merve Gürel, Xu Chu, Wentao Wu, and Ce Zhang. 2021. Nearest neighbor classifiers over incomplete information: From certain answers to certain predictions. In *PVLDB*, Vol. 14.
- [19] Konstantina Kourou, Themis P Exarchos, Konstantinos P Exarchos, Michalis V Karamouzis, and Dimitrios I Fotiadis. 2015. Machine learning applications in cancer prognosis and prediction. *Computational and structural biotechnology journal* 13 (2015), 8–17.
- [20] Sanjay Krishnan, Michael J Franklin, Ken Goldberg, and Eugene Wu. 2017. Boostclean: Automated error detection and repair for machine learning. *arXiv preprint arXiv:1711.01299* (2017).
- [21] Sanjay Krishnan, Jiannan Wang, Eugene Wu, Michael J Franklin, and Ken Goldberg. 2016. ActiveClean: interactive data cleaning for statistical modeling. *Proceedings of the VLDB Endowment* 9, 12 (2016), 948–959.
- [22] Sanjay Krishnan and Eugene Wu. 2019. Alphaclean: Automatic generation of data cleaning pipelines. *arXiv preprint arXiv:1904.11827* (2019).
- [23] Erin LeDell and Sebastian Poirier. 2020. H2o automl: Scalable automatic machine learning. In *Proceedings of the AutoML Workshop at ICMML*, Vol. 2020.
- [24] Peng Li, Xi Rao, Jennifer Blase, Yue Zhang, Xu Chu, and Ce Zhang. 2021. CleanML: a study for evaluating the impact of data cleaning on ml classification tasks. In *2021 ICDE*. IEEE, 13–24.

- [25] Hanxiao Liu, Karen Simonyan, and Yiming Yang. 2018. Darts: Differentiable architecture search. *arXiv preprint arXiv:1806.09055* (2018).
- [26] Risheng Liu, Jiabin Gao, Jin Zhang, Deyu Meng, and Zhouchen Lin. 2021. Investigating bi-level optimization for learning and vision from a unified perspective: A survey and beyond. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2021).
- [27] Felix Neutatz, Binger Chen, Ziawasch Abedjan, and Eugene Wu. 2021. From Cleaning before ML to Cleaning for ML. *Data Engineering* (2021), 24.
- [28] Felix Neutatz, Binger Chen, Yazan Alkhatib, Jingwen Ye, and Ziawasch Abedjan. 2022. Data Cleaning and AutoML: Would an optimizer choose to clean? *Datenbank-Spektrum* (2022), 1–10.
- [29] Randal S Olson and Jason H Moore. 2016. TPOT: A tree-based pipeline optimization tool for automating machine learning. In *Workshop on automatic machine learning*. PMLR, 66–74.
- [30] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. 2017. Automatic differentiation in PyTorch. (2017).
- [31] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* 12 (2011), 2825–2830.
- [32] Leif E Peterson. 2009. K-nearest neighbor. *Scholarpedia* 4, 2 (2009), 1883.
- [33] Alexandre Quemy. 2020. Two-stage optimization for machine learning workflow. *Information Systems* 92 (2020), 101483.
- [34] Jyoti Ramteke, Samarth Shah, Darshan Godhia, and Aadil Shaikh. 2016. Election result prediction using Twitter sentiment analysis. In *2016 international conference on inventive computation technologies (ICICT)*, Vol. 1. IEEE, 1–5.
- [35] Barna Saha and Divesh Srivastava. 2014. Data quality: The other face of big data. In *2014 ICDE*. IEEE, 1294–1297.
- [36] Rodrigo Santa Cruz, Basura Fernando, Anoop Cheria, and Stephen Gould. 2017. Deeppermnet: Visual permutation learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 3949–3957.
- [37] Zeyuan Shang, Emanuel Zraggen, Benedetto Buratti, Ferdinand Kossmann, Philipp Eichmann, Yeounoh Chung, Carsten Binnig, Eli Upfal, and Tim Kraska. 2019. Democratizing data science through interactive curation of ml pipelines. In *Proceedings of the 2019 International Conference on Management of Data*. 1171–1188.
- [38] Richard Sinkhorn. 1964. A relationship between arbitrary positive matrices and doubly stochastic matrices. *The annals of mathematical statistics* 35, 2 (1964), 876–879.
- [39] Richard Sinkhorn and Paul Knopp. 1967. Concerning nonnegative matrices and doubly stochastic matrices. *Pacific J. Math.* 21, 2 (1967), 343–348.
- [40] Chris Thornton, Frank Hutter, Holger H Hoos, and Kevin Leyton-Brown. 2013. Auto-WEKA: Combined selection and hyperparameter optimization of classification algorithms. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*. 847–855.
- [41] Joaquin Vanschoren, Jan N. van Rijn, Bernd Bischl, and Luis Torgo. 2013. OpenML: Networked Science in Machine Learning. *SIGKDD Explorations* 15, 2 (2013), 49–60. <https://doi.org/10.1145/2641190.2641198>
- [42] Zixuan Zhao and Raul Castro Fernandez. 2022. Leva: Boosting Machine Learning Performance with Relational Embedding Data Augmentation. In *Proceedings of the 2022 International Conference on Management of Data*. 1504–1517.

Received October 2022; revised January 2023; accepted February 2023