

Step-by-Step Data Cleaning Recommendations to Improve ML Prediction Accuracy

Sedir Mohammed
Hasso Plattner Institute
University of Potsdam
Germany
Sedir.Mohammed@hpi.de

Felix Naumann
Hasso Plattner Institute
University of Potsdam
Germany
Felix.Naumann@hpi.de

Hazar Harmouch
University of Amsterdam
The Netherlands
h.harmouch@uva.nl

ABSTRACT

Data quality is crucial in *machine learning* (ML) applications, as errors in the data can significantly impact the prediction accuracy of the underlying ML model. Therefore, data cleaning is an integral component of any ML pipeline. However, in practical scenarios, data cleaning incurs significant costs, as it often involves domain experts for configuring and executing the cleaning process. Thus, efficient resource allocation during data cleaning can enhance ML prediction accuracy while controlling expenses.

This paper presents COMET, a system designed to optimize data cleaning efforts for ML tasks. COMET gives step-by-step recommendations on which feature to clean next, maximizing the efficiency of data cleaning under resource constraints. We evaluated COMET across various datasets, ML algorithms, and data error types, demonstrating its robustness and adaptability. Our results show that COMET consistently outperforms feature importance-based, random, and another well-known cleaning method, achieving up to 52 and on average 5 percentage points higher ML prediction accuracy than the proposed baselines.

1 DATA CLEANING FOR ML

In an era dominated by data-driven strategies and rapid development in *machine learning* (ML), *data quality* has become a significant factor in the success of ML applications. The availability of vast and diverse datasets has empowered various domains, such as healthcare, biology and finance, to leverage the capabilities of ML algorithms, fostering significant improvements [30, 41]. However, the effectiveness of these algorithms depends on the quality of the input data during training and testing phases [11]. Real-world datasets often contain imperfections, inconsistencies, and inaccuracies that can significantly impact the prediction accuracy of ML models. Thus, data cleaning, which entails identifying and correcting data errors, is essential for reliable and accurate ML predictions.

Traditionally, data scientists performed cleaning during data acquisition, often without considering the underlying ML task [17]. However, the focus has shifted from isolated data cleaning strategies preceding the ML task (“Cleaning before ML”) to an integrated perspective that views data cleaning and the underlying ML task as a cohesive entity (“Cleaning for ML”) [27]. This development marks a new era in which data quality and ML outcomes are symbiotic components within the ML pipeline.

The growing paradigm of data-centric *Artificial Intelligence* (AI) emphasizes the role of data quality throughout the ML pipeline [39, 41]. Initially, developers and researchers focused on models to improve prediction accuracy. With data-centric AI, the focus

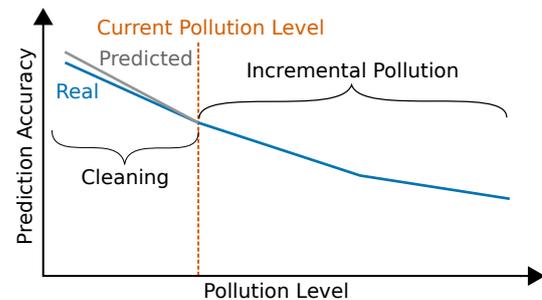


Figure 1: COMET incrementally pollutes features and based on the observed negative effect of the pollution on the prediction accuracy, it estimates the positive impact of data cleaning.

is on the data and its quality instead of the ML model, potentially leading to more robust and generalizable AI applications. However, this shift also presents challenges, like ensuring data quality, obtaining reliable annotations, handling missing values, and fostering data diversity. Especially in real-world scenarios where data cleaning is often associated with costs as this task is typically performed by users, respectively domain experts. These challenges motivate the need for strategic data preparation methods to ensure effective ML models and an efficient data cleaning process [30, 41].

Given the common scenario where a user possesses a (dirty) dataset and aims to deploy an ML model within a limited resource allocation for data cleaning [25], it is beneficial to determine the order of clean operations, and thus to maximize the improvement of the downstream ML task within a budget. To determine this order, we introduce the Cleaning Optimization and Model Enhancement Toolkit (COMET), a progressive approach for cleaning recommendations in such scenarios. Our approach provides the user with a step-by-step recommendation on which feature to clean next.

By analyzing the relationships between feature-wise data quality, cleaning costs, and their impact on ML outcomes, COMET allocates the cleaning budget to maximize improvements in prediction accuracy. To estimate how feature-wise data cleaning affects prediction accuracy, COMET incrementally injects additional errors into features (see Figure 1) and measures prediction accuracy after each pollution. It then interpolates a trend from these measurements to predict the effect of cleaning the respective feature. Considering the cleaning costs, COMET recommends the next feature to clean. The principle of incremental data pollution has shown promise in other fields [3]: In quantum computing, researchers amplify noise during calculations to back-calculate results in a noise-free environment.

© 2025 Copyright held by the owner/author(s). Published in Proceedings of the 28th International Conference on Extending Database Technology (EDBT), 25th March-28th March, 2025, ISBN 978-3-89318-099-8 on OpenProceedings.org. Distribution of this paper is permitted under the terms of the Creative Commons license CC-by-nc-nd 4.0.

The cleaning process within COMET can be executed manually or through automatic error correction algorithms [17]. Recently, *Automated Machine Learning* (AutoML) emerged to streamline and automate various ML stages, from data preprocessing, including data cleaning, to feature engineering, algorithm selection, and hyperparameter tuning [14]. The flexibility of COMET allows it to fit seamlessly into an AutoML pipeline. In critical domains, where the data cleaning process might require expert validation even when it is done automatically, COMET’s step-wise feature cleaning recommendations provide a controlled and transparent approach.

In this paper, “prediction accuracy” refers to metrics like accuracy or F1 score, even though COMET can also optimize other AI-related metrics, such as fairness or bias [39, 41].

To evaluate our approach, we introduce three baselines and a related approach. The first baseline uses feature importance from the (dirty) input data to guide feature-wise cleaning recommendations [24]. The second baseline acts on random feature selection – a contrastingly non-strategic method. The third baseline breaks COMET’s step-wise approach, using a ranked feature list from the first step for all subsequent cleaning operations. The fourth contender is *ActiveClean* [20], a well-known method aiming to optimize the cleaning process like COMET. Our empirical results show that COMET, on average, outperforms the baselines across diverse datasets, error types and ML algorithms. The results show the potential for improved model prediction, emphasizing the importance of an informed data cleaning strategy in resource-constrained scenarios.

Contribution. In summary, our research advances the understanding of data cleaning as a key facet of data-centric AI. It offers practitioners a principled approach to optimize data-cleaning efforts within budget constraints. As ML increasingly permeates various domains, insights from this study promise to enable more accurate and cost-efficient deployment of ML models on real-world, imperfect data. We specifically make the following main contributions:

- (1) We present COMET, an innovative approach that incrementally pollutes data to assess the effects of data cleaning on prediction accuracy and iteratively makes cleaning recommendations while considering cleaning costs.
- (2) We evaluate the performance of COMET by simulating the cleaning of various datasets and by comparing the prediction accuracies against two cleaning baselines and *ActiveClean* [20].
- (3) We further evaluate COMET using the CleanML benchmark [5, 22], which provides both dirty and clean versions of real-life datasets for comprehensive analysis.

Outline. Section 2 presents relevant related work. Section 3 details the functionality and various components of COMET. We present the data considered in the experiments, ML algorithms, error types, cleaning cost models, and other implementation details of COMET in Section 4. Section 5 shows the results of the experiments in which we compare COMET to the baselines and related work. Finally, Section 6 summarizes our findings and outlines future research directions.

2 RELATED WORK

Traditional error detection and cleaning methods, such as NA-DEEF [8] and KATARA [4], rely on user-defined rules or domain expertise, which can be costly and labor-intensive. As the complexity and volume of data increased, and ML methods became

more popular, the need for automated and adaptive data cleaning methods became more pressing. An example of ML-based *error detection* is HoloDetect [15], which learns error patterns in the data to generate further training data synthetically and trains an error detection model. Another approach is ED2 [29] which uses active learning and selects the entries for which the ML model is currently unsure for labelling by the user. There exist also dedicated approaches for automatic *data cleaning*. HoloClean [32] processes different cleaning signals in data, using probability theory to reconcile inconsistencies across various signals and repair data. Similarly, SCARE [40] uses ML techniques to repair data by maximizing data likelihood while introducing minimal changes.

The presented methods and a large amount of similar research focus solely on data without considering its downstream application. However, erroneous data may not be equally disruptive, as Foroni et al. [11] have explicitly demonstrated for the ML context as a downstream task. This aligns with our intuition in this paper. In the following, we discuss existing approaches that consider data cleaning in the context of ML applications.

ActiveClean treats data cleaning as a stochastic gradient descent problem [20]. It selects iteratively records for cleaning that are estimated to have the highest gradient of the loss function after cleaning the respective records. The gradient per record is estimated from the gradient of the dirty record and previous cleaning procedures. In contrast, COMET predicts the cleaning effect in each iteration by introducing additional errors into the current pollution state, making it adaptive by always estimating cleaning impacts based on the current state of the data. Unlike *ActiveClean*, which starts with a random sample due to a lack of gradient information, COMET provides well-founded recommendations from the first iteration. Despite these differences, *ActiveClean* is similar enough to COMET in its approach to serve as a baseline in our experiments.

CPClean adopts an incremental cleaning strategy similar to COMET, using a step-by-step process primarily informed by validation sets and counting queries [18]. This process persists until it determines that additional cleaning will no longer affect the prediction accuracy. However, *CPClean* is optimized specifically for nearest-neighbor models, lacking the model-agnostic flexibility that COMET offers.

BoostClean approaches data cleaning as a boosting problem, generating a cleaning program applicable to training or test records, while actively conducting the cleaning [19]. The primary criterion for cleaning decisions is the model’s test accuracy. Although *BoostClean* estimates the impact of data cleaning similarly to our approach, it requires fully clean labels in the test data and pre-defined detection and cleaning functions by a domain expert. In contrast, COMET operates without prerequisites on the cleaning state of training and test data and autonomously estimates the impact of cleaning without relying on pre-defined rules or functions.

AutoSklearn is a complete AutoML system that automates data cleaning by constructing and evaluating ML pipelines with model validation accuracy [10]. It focuses on imputation when cleaning errors based on mean, median or most frequent values. Neutatz et al. [28] extended the system with additional data cleaning methods for outlier detection and advanced imputation strategies. Although these AutoML frameworks offer comprehensive

solutions, their data cleaning functionalities are tailored to handle specific error types and require pre-defined detection and cleaning routines.

DiffPrep, similar to COMET aims to optimize data cleaning but focuses on feature-wise transformations during pre-processing while considering the impact on the underlying ML model [21]. DiffPrep determines the order of transformation types (e.g., cleaning missing values first, then outliers) and selects the appropriate cleaning method (e.g., mean imputation for missing values, Z-score for outliers) for each feature, tailored to the ML model. Unlike COMET, DiffPrep does not clean incrementally; it establishes a final transformation pipeline for all features without guiding the user step-by-step, which limits direct comparability with COMET.

3 CLEANING RECOMMENDATION

COMET provides cost-aware feature-wise cleaning recommendations in a progressive fashion, i.e., offering step-by-step guidance on which features should the next cleaning effort be spent on. These cleaning recommendations are always error type and ML algorithm-specific. However, COMET can handle arbitrary error types and ML algorithms, without being restricted to a specific ML task. In this manuscript, we configured COMET for four specific error types (see Section 3.4) and focused on classification. The actual cleaning can be performed by either a data cleaning algorithm, such as HoloClean [32], SCARE [40], or, more importantly, a dedicated domain expert. At no point during the process does COMET require information about the actual pollution level of the individual features, nor which entries are actually erroneous. Figure 2 shows the architecture of COMET, consisting of three modules: *Polluter* (Section 3.1), *Estimator* (Section 3.2) and *Recommender* (Section 3.3). COMET initially requires (dirty) data as input, which serves as the basis for generating cleaning recommendations for the subsequent cleaning step. In the following subsections, we provide a detailed description of each module and its functionality.

In the rest of the paper, we refer to the number of data entries to be cleaned in each iteration as the *cleaning step*. We also use the term *Cleaner* to denote the data cleaning methods, encompassing both the algorithm-based and the human-based techniques.

In the description of each module, we assume a specific error type. However, the process can be applied separately to different error types to identify the optimal feature and error type combination.

3.1 Incremental data pollution

The first module of COMET, the *Polluter*, incrementally pollutes the input data given a specific error type for which COMET should make a cleaning recommendation in the current iteration. The *Polluter* introduces additional erroneous entries for each feature, quantified by a *pollution step*, which is the counterpart to a cleaning step. Additionally, we use the notion of the pollution level, which denotes the percentage of cells that should be additionally polluted. Since selecting the entries to be polluted may also have a potentially different effect on the prediction accuracy, the *Polluter* selects randomly multiple combinations of entries per feature to be polluted. So, the functionality of the *Polluter* is represented mathematically by the function:

$$\text{Polluter}(d, f, \text{Err}, \rho) = d'_{f,\rho,c} \quad (1)$$

where d represents the input data, f is the considered feature to be polluted, Err is the error type, ρ signifies the pollution level for the feature f , and $d'_{f,\rho,c}$ represents a polluted data state for the pollution level ρ and the combination c .

Example. Suppose we have a feature “Income” in a dataset d . If the error type Err is missing values and the pollution level ρ is set to 1%, the *Polluter* introduces missing values to 1% of the entries in the “Income” column. This results in a polluted data state $d'_{\text{Income},1\%,c}$, where c represents one specific combination of polluted entries.

By default, COMET assumes all features contain dirty entries and must be cleaned until the *Cleaner* marks them as clean. To produce the next cleaning recommendation, the *Polluter* performs two further pollution steps for each feature (see Figure 1). Each new pollution level corresponds to an additionally polluted input data state. Pollution is applied separately to the train and the test data to prevent information leakage between them.

It is worth noting that, given that the *Polluter* lacks knowledge about which specific entries in a feature are erroneous, it may overwrite already dirty entries with artificial errors. Thus, the targeted pollution level might not be fully achieved after a pollution step. Using the hypergeometric distribution, we estimate that the probability of selecting clean entries to pollute remains high when the number of already dirty entries is small relative to the total.

3.2 Cleaning impact estimation

The *Estimator* receives from the *Polluter* feature by feature the different versions of the data, including d and a set D'_f of $d'_{f,\rho,c}$. The *Estimator* operates in two steps:

Step 1: Pollution Effect Measurement. We evaluate the influence of incremental feature-wise pollution on the prediction accuracy for the considered pollution levels ρ . Since the prediction accuracy depends on the used ML algorithm, COMET expects an ML algorithm as an additional input in the *Estimator*. This phase is expressed as

$$E_1(d, D'_f, f, \text{MLA}) = A(f) \quad (2)$$

where MLA is the chosen ML algorithm, and $A(f)$ is a set of measured prediction accuracies per pollution step for the feature f .

Example. Given the data states $D'_f = \{d'_{\text{Income},1\%,c}, d'_{\text{Income},2\%,c}\}$ and the ML algorithm MLA (e.g., SVM), the *Estimator* computes the prediction accuracy $A(\text{Income}) = \{0.85, 0.82\}$ for the respective pollution levels $\rho = 1\%, 2\%$.

Step 2: Predictive Model Construction. In the second step, denoted as E_2 , the *Estimator* utilizes the prediction accuracies measured from the incrementally polluted input data per feature. Based on the measured prediction accuracies, the *Estimator* trains a Bayesian regression model to predict the prediction accuracy for the next cleaning step (see the line denoted as “predicted” in Figure 1). The regression-based prediction is formulated as

$$E_2(A(f)) = (P_{\text{next}}(f), U(f)) \quad (3)$$

where $P_{\text{next}}(f)$ is the predicted accuracy for the feature f and $U(f)$ is the uncertainty of the prediction. The Bayesian regression model quantifies the prediction uncertainty, which the *Recommender* uses later for feature ranking and selection.

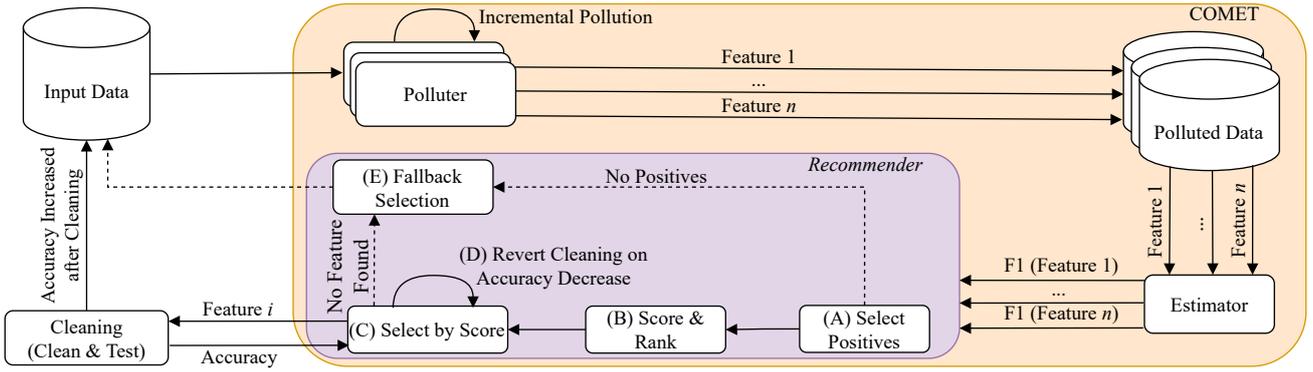


Figure 2: COMET workflow for an individual error type: (1) *Polluter*: Introduce further pollution; (2) *Estimator*: Evaluate pollution/cleaning effects on ML model accuracy; (3) *Recommender*: Propose feature-wise cleaning strategies based on scoring.

3.3 Optimal feature selection

The last module, the *Recommender*, selects the next feature for the *Cleaner* to clean (see Figure 2). This decision is based on the predicted accuracy per feature, factoring in the cleaning costs. Initially, the *Recommender* considers only features with a positive predicted accuracy (Figure 2 – (A) Select Positives). These are features where the *Estimator* predicts that the prediction accuracy improves after cleaning by one cleaning step. It then scores and ranks these features, using a score that balances the potential accuracy gain against its associated uncertainty and cleaning cost (Figure 2 – (B) Score & Rank). We define this score for a feature f as

$$\text{Score}(f) = \frac{P_{\text{next}}(f) - U(f)}{C(f)} \quad (4)$$

where $P_{\text{next}}(f)$ is the predicted accuracy gain for a feature f , $U(f)$ is the uncertainty associated with this prediction, and $C(f)$ is the cleaning cost function for the feature f . The uncertainty $U(f)$ is calculated as the difference between the upper and lower bounds of the confidence interval of the Bayesian regression model. This ranking prioritizes features that provide the most significant prediction accuracy improvement per unit of cleaning cost, while also considering the confidence level of these predictions. Thus, COMET recommends the most cost-effective feature f for cleaning in the current iteration (Figure 2 – (C) Select by Score).

Example. For the feature “Income”, the predicted accuracy improvement is $P_{\text{next}}(\text{Income}) = 0.88$, the uncertainty is $U(\text{Income}) = 0.02$, and the cleaning cost is $C(\text{Income}) = 1$. The score is computed as: $\text{Score}(\text{Income}) = \frac{0.88 - 0.02}{1} = 0.86$.

As part of the cleaning recommendation, COMET provides the *Cleaner* with details on which entries of the recommended feature were temporarily polluted by the *Polluter* (D'_f). These entries, which have led to the recommendation, are potentially partially dirty in the input data d_f (feature f of the input data) – the *Cleaner* should first clean them. If the already dirty entries in d_f are fewer than required for a cleaning step, the *Cleaner* must clean additional random entries from this feature. After cleaning, the *Recommender* evaluates the impact on prediction accuracy and considers the cleaning step as successful if the accuracy increases – a new iteration of COMET starts (Figure 2 – *Cleaner*, Clean & Test). If the prediction accuracy decreases after cleaning, COMET reverts the data to its pre-cleaning state, retaining the cleaned data in a *cleaning buffer* (step (D) in Figure 2).

The *Recommender* then moves to the next highest-ranked feature for cleaning (Figure 2 – (C) Select by Score). If the feature to be cleaned is in the cleaning buffer, the *Recommender* removes it and applies the changes to data.

In rare cases where all ranked features result in a decrease in prediction accuracy after cleaning, or none are predicted to improve accuracy, the *Recommender* switches to a fallback strategy (Figure 2 – (E) Fallback Selection). This strategy focuses on cleaning the feature that previously archived highest F1 score after cleaning. If a feature is deemed completely cleaned, the *Recommender* moves on to the next most important feature in the ranking.

Once a feature is cleaned by one cleaning step, the *Recommender* compares the actual increase in prediction accuracy to the one predicted by COMET, identifying any discrepancies in the predictive modelling for the considered feature. The *Estimator* then adjusts the prediction to improve its accuracy, which involves calculating the mean of the measured discrepancies and then modifying the prediction by adding or subtracting this mean value. Even if the *Recommender* evaluates the cleaning as inefficient, and COMET restores the pre-cleaning state, or activates the fallback strategy, the *Estimator* adjusts the prediction model for that feature.

The *Recommender* based decision-making, denoted as R , is formalized as

$$R(\{P_{\text{next}}(f_i), C(f_i), U(f_i)\}_{i=1}^n) = f_{\text{rec}} \quad (5)$$

where the *Recommender* R expects the predicted accuracy, the cleaning cost, and the uncertainty for each considered feature $f_1 \dots f_n$ to recommend a feature f_{rec} for cleaning the considered error type by one cleaning step.

3.4 Error types

COMET is conceptually an error type-agnostic approach. However, we configured COMET for a specific set of error types Err , which we detail in the following. We consider error types that emulate common data errors, such as incomplete or incorrectly data, scale inconsistencies, and random noise. For the data pollution, the *Polluter* uses the JENGA framework [35] to randomly sample a specific number of records x from a feature, based on the pollution level, and introduce the respective error type. The following presents the considered error types, including the respective pollution process.

Missing values. Missing values are a common problem in many real-world data [9, 34]. These can appear as empty entries or be represented by placeholders, such as “NaN”. However, hidden missing values may use unconventional placeholders [31], such as representing a missing date with 1900-01-01. To pollute the data, the *Polluter* replaces the randomly selected records of a feature with a placeholder.

Gaussian noise. Numerical values can be noisy, e.g., due to erroneous sensors, external interference, or incorrect user input. To pollute the data for a given feature, the *Polluter* randomly samples x records and adds Gaussian noise to each selected data point, generated with a mean of zero and a randomly chosen standard deviation within [1-5].

Categorical shift. Categorical shift is the counterpart to Gaussian noise for categorical variables, where incorrect categories are assigned values. During the pollution, the *Polluter* samples x records from a feature and swap their categories with other categories in the feature.

Scaling. We assume scaling in numerical values could occur due to incorrect conversion of units, such as from *cm* to *m*. To pollute the data, the *Polluter* randomly increases the scaling by 10, 100, or 1000 in the selected rows of the considered feature.

4 EXPERIMENTAL SETUP

This section details the methods and datasets we use to demonstrate the performance of COMET in giving valuable recommendations for iterative feature-wise cleaning. In the rest of the paper, we refer to a unique combination of a dataset, an ML algorithm, and a specified error type in the data as a *configuration*. Each configuration is a distinct experimental scenario, allowing us to study the effects of various data conditions, algorithms, and errors on COMET.

4.1 Pollution and cleaning setup

To validate the recommendation-based cleaning, we use datasets that have both dirty and cleaned versions (ground truth) and other datasets that we artificially pre-polluted to establish a ground truth. We introduce the datasets in Section 4.3. Each setting of pollution levels across features is referred to as a *pre-pollution setting*. These settings affect only features, so we do not add any errors to the labels in our experiments. Since the pre-pollution setting is at dataset level, we keep the same pre-pollution settings across different configurations involving the same dataset. Based on discussions with ML experts, both training and test data are equally polluted, reflecting common real-world scenarios.

Given the potential variations in data quality per feature in real-world datasets and the significant influence of the pre-pollution level distribution on the prediction quality [11], it is essential to consider diverse pre-pollution settings. Thus, we sample for each pre-pollution setting the pollution level per feature using an exponential distribution to ensure a wide-ranging representation of pollution level distribution. The cleaning and pollution step is set at 1% of the total data size of the train or test data to ensure a consistent impact across experiments. We consider two scenarios for the pre-pollution: In the first scenario, we pollute the data according to the pre-pollution level, with only one error type. In the second scenario, we randomly select an error type for each pollution step of a feature during pre-pollution. To conduct the pre-pollution, we use the same pollution methods used by the *Polluter*.

4.2 Cleaning cost models

For the two scenarios, where we consider either single or multiple errors, we use different cleaning cost models per feature. In the first scenario, applied within the single-error context, we assume a *constant* cost function for each feature: each cleaning step incurs a uniform cost of one unit. In this way, we maintain comparability and avoid favoring any particular feature. However, in practice, different error types might have varying cleaning costs, thus in the multi-error scenario, we go beyond *constant* cost function by incorporating *linear* and *one-shot* cost functions associated with different error types. Note that our assignment of cost functions to error types is an example and can be adapted as needed.

Constant cost. We apply this cost to both *categorical shifts* and *scaling* errors. Categorical shifts can be identified by detecting frequent correlations through FD discovery algorithms or association rule mining, with deviations corrected via imputation. Similarly, scaling errors identified by outlier detection are also corrected using imputation methods. In the experiments, we assign the cost of one unit per cleaning step.

One-shot cost. A one-shot cost function implies a higher initial cleaning cost that does not recur in subsequent steps. This model is used for cleaning *missing values*, starting with identifying missing values and then performing data imputation once for the entire column. In the experiments, we assign costs of two units for the first cleaning step and zero for all further cleaning steps.

Linear cost. In this model, each cleaning step of a particular error type costs incrementally more than the previous one. We assume a linear cost function for the cleaning of *Gaussian noise*. Initial detection involves estimating noise distribution and identifying strong outliers, which becomes gradually harder as subtle deviations are harder to detect. In the experiments, we assume an increase of one unit per cleaning step performed, with initial costs of one unit.

For both scenarios, we limited the cleaning budget to 50 units, balancing practicality and experimental rigor. We are aware that the effectiveness of such budgets differs notably across various pre-pollution settings: it may be adequate for thoroughly cleaning datasets with fewer dirty features or lower pollution levels, or, conversely, it may only allow partial cleaning for more dirty features or overall higher pollution levels. This variation mirrors real-world data cleaning challenges, where the extent of pollution significantly impacts the resources needed for effective cleaning.

4.3 Datasets

We use seven different datasets commonly used for classification tasks. Three of them have both clean and dirty versions available and include at least one of the error types considered in this paper. These datasets are part of the benchmark provided by CleanML [5, 22]. Additionally, we established three distinct pre-pollution settings for four further datasets from the UCI Machine Learning Repository and Kaggle. We present the used datasets in the following; in addition, Table 1 shows an overview of the characteristics of these datasets.

Datasets used with pre-pollution. The *Contraceptive Method Choice* (CMC) dataset contains a subset of the 1987 Indonesian National Contraceptive Prevalence Survey [23]. The classification task is to predict the current contraceptive method of the women surveyed. The *EEG Eye State* (EEG) dataset contains data from EEG measurements with the Emotiv EEG neuroheadset [33]. The classification task is to predict the eye state (closed or open). The

“categorical shift” is not applicable to this dataset, since it contains only numerical variables. The *Telco customer churn* (Churn) dataset from IBM consists of fictional customer data from a telecommunications company [37]. The classification task is to predict whether the customer terminated their relationship with the company as an active customer in the last month. Finally, the *South German Credit* (S-Credit) dataset contains bank data from a southern German bank from 1973 to 1975 [13, 16, 26]. we use in our experiments the version from [13]. The classification task is to predict whether a bank customer will comply with the terms of the contract or not.

Datasets provided by CleanML. The *Airbnb* dataset contains hotel information from the top 10 tourist destinations and US metropolitan areas [1, 5, 22]. The CleanML authors provided the data by scraping the Airbnb.com website. The classification task is to predict whether the rating of a hotel is 5 or not. In the context of this dataset, we consider scaling errors. The *Credit* dataset consists of credit data, whereby the classification task is to predict whether a client will be in financial distress in the next two years [5, 7, 22]. In the context of this dataset, we consider missing values and scaling errors. The *Titanic* dataset contains passenger records [5, 22, 38]. The classification task is to determine whether a passenger survives. In the context of this dataset, we consider missing values.

Name	# Rows	# Cat.	# Num.	# Class
<u>Datasets used with Pre-pollution</u>				
CMC	1,473	7	2	3
Churn	7,032	16	3	2
EEG	14,980	0	14	2
S-Credit	1,000	17	3	2
<u>Datasets provided by CleanML</u>				
Airbnb (scaling errors)	26,288	3	37	2
Credit (scaling errors)	11,985	0	10	2
Titanic (missing values)	891	6	2	2

Table 1: Overview of our used datasets.

4.4 ML algorithms

While COMET is not limited to a specific ML task, here we focus on classification. To demonstrate the superiority of COMET over state-of-the-art baselines, we selected four diverse ML algorithms: Support Vector Machine (SVM) [6], a k -nearest neighbors classifier (KNN) [2], a multi-layer perceptron (MLP) [36] and a Gradient Boosting classifier (GB) [12]. We performed a 10-sampled random hyperparameter optimization for each configuration and pre-pollution setting. With this, we simulate a real-world scenario wherein users working with dirty data aim for the highest prediction accuracy given the dataset’s current state.

4.5 Evaluation metrics and baselines

As COMET makes recommendations for individual cleaning steps, we compare the prediction accuracy per step to evaluate the performance of COMET against four baselines: *random recommendations* (RR), *feature importance-based recommendations* (FIR), light version of COMET (CL) and *ActiveClean* (AC). As a metric for the prediction accuracy, we use the well-known F1 score.

RR adopts a randomized approach, randomly selecting a feature in each cleaning step among those that have been marked

to be cleaned. For our evaluation, this process is repeated five times for each pre-pollution setting, and we averaged the resulting F1 scores per cleaning step. FIR uses Shapley values [24] to rank features by importance within each configuration and pre-pollution setting, selecting the highest-ranked yet polluted feature for cleaning in each step until it is fully cleaned, then moving to the next. CL is a simplified version of COMET: it applies the idea of COMET once to generate a ranked list of features based on the *Estimator*’s output. Similar to FIR, the *Recommender* selects the feature with the highest rank at each cleaning step and continues until that feature is fully cleaned. Similar to COMET, the *Recommender* in CL uses the same cleaning step and reverting and fallback strategies. The authors of AC have focused on ML algorithms with a convex loss function. Specifically, they used SVM (in the following denoted as *AC-SVM*), linear regression (LIR) and logistic regression (LOR) [20]. Thus, we focus on these ML algorithms when comparing COMET with AC. AC relates the cleaning procedure to the number of records that can be cleaned. We adapt AC to align with our concept of cleaning on a feature level, ensuring a comparable experimental setting.

Additionally, we introduce a *local optimum* (Oracle), which shows the optimal feature for cleaning in each step based on the gain in F1 score relative to the associated cleaning costs. The cleaning order generated by the Oracle does not always guarantee the highest accuracy at every cleaning step compared to COMET or the other baselines. A single divergent decision on feature cleaning by COMET or the baselines can lead to divergent cleaning orders that can reach, at some point, a higher accuracy than the Oracle. However, Oracle performs much better than any other approach on average and thus can be treated as an upper bound.

5 EXPERIMENTAL RESULTS

This section delves into the findings from our extensive evaluation of COMET, conducted across various datasets, error types, ML algorithms, and baselines. We designed these experiments to answer the following research questions (RQ) under two scenarios:

- (1) How effective is COMET in multi-error types and diverse cost functions scenario?
- (2) How effective is COMET compared to baselines FIR, RR, and CL?
- (3) How effective is COMET compared to related work AC [20]?
- (4) How does COMET performance vary among different error types and ML algorithms?
- (5) How accurate are the predictions produced by the *Estimator*?
- (6) How efficient is COMET in terms of runtime?

For the comparison with the respective baselines (RQ 1-RQ 4), we calculate the F1 score differences (shown as *F1 advantage* in the plots) for each cleaning step and the considered baseline we average these differences across the pre-pollution setting (for the datasets used with pre-pollution). A positive F1 advantage indicates that COMET outperforms the corresponding baseline.

5.1 Comparison to baselines for multiple error types and diverse cost functions

For the first research question (RQ 1), we compare the performance of COMET with FIR, RR, CL (see Figure 3) and AC (see Figure 4), considering multiple error types present in the data/per feature and various cost functions, as introduced in Section 4.2, for cleaning. As the CleanML datasets contain only single error

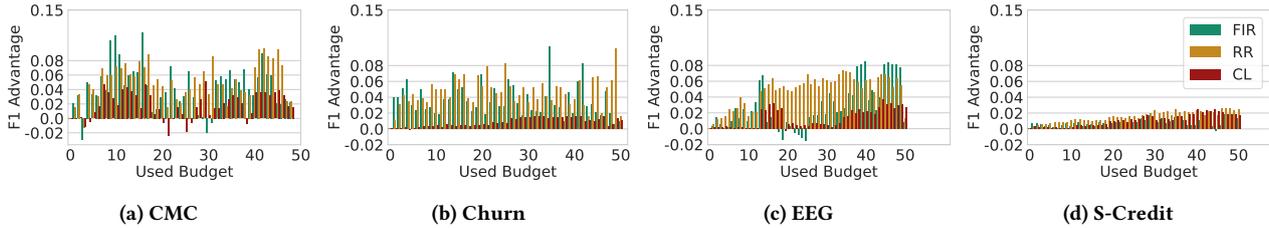


Figure 3: Comparison of COMET with the baselines for SVM across multiple error types and cost functions.

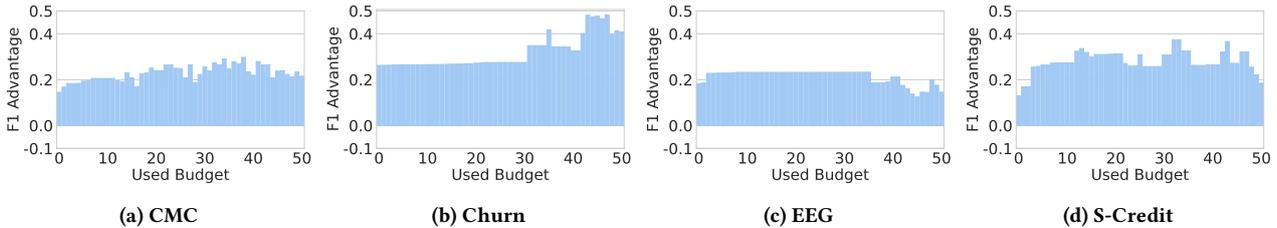


Figure 4: Comparison of COMET with AC for LIR across multiple error types and cost functions.

types, we excluded them and focused on the datasets we pre-polluted (see Table 1). Since our experimental setup consists of many configurations and pre-pollution settings, we focus on several specific configurations: those using SVM as the ML algorithm for FIR, RR, and CL, and those using LIR for the comparison with AC. Our analysis showed that COMET achieved the best results with SVM compared to FIR, RR, and CL, and the best results with LIR in comparison to AC.

Comparison to FIR, RR and CL. In the comparison with FIR, RR and CL, the results show that COMET outperforms the baselines in this scenario. COMET achieves an F1 score difference of up to 11%*pt* for CMC. COMET is also consistently superior throughout the cleaning process for Churn, EEG and S-Credit, although the differences vary. For example, COMET achieves a difference of 10%*pt* in the F1 score compared to FIR in Churn with an invested budget of 35, while for the next cleaning step, the difference drops to 3%*pt*. Overall, COMET maintains a positive difference across all datasets. The difference for S-Credit is smaller. Our analysis shows that the average difference between the F1 score in dirty and fully cleaned state is -1.5%*pt*, limiting the potential for clear superiority.

Considering different error types within a feature and introducing various cost functions expands the search space for potential cleaning steps. FIR and RR more frequently make suboptimal decisions. The different cost functions reinforce these findings, especially for linear cleaning costs, where repeated poor decisions are heavily penalized by consuming more budget. In contrast, COMET accurately estimates the impact of cleaning combinations of features and error types, maintaining its superiority.

Comparison to AC. In the comparison to AC, Figure 4 shows that COMET consistently outperforms AC throughout the entire cleaning process. In most cases, COMET achieves a difference of 20%*pt* in the F1 score across all datasets. For Churn (Figure 4(b)), COMET even achieves a maximum F1 score difference of nearly 50%*pt*. Overall, the differences are less erratic than in the single-error and constant cost function comparisons (Figure 8), which is attributed to the influence of the cost functions. Since AC cleans on a per-record basis, different error types are corrected across multiple features during each cleaning step. For comparison, we

propagate for longer periods the F1 scores achieved from previously utilized budget units until an actual F1 score is measured for the current unit.

5.2 Comparison to FIR, RR, and CL for a single error type

In RQ 1, we examined the performance of COMET in a multi-error setting, considering various cost functions. To assess the impact of individual error types on COMET’s performance compared to FIR, RR, and CL, we now focus on single-error scenarios as part of the second research question (RQ 2). The results are shown in Figures 5 and 6. Here, we assume constant costs for each cleaning step, regardless of the error type, to ensure comparability. An early stop in the cleaning process indicates that the dataset is fully cleaned. Again, due to the extensive configurations and pre-pollution settings, we focus on one ML algorithm: MLP. As COMET performs worse with MLP than other algorithms, this experiment highlights its worst-case performance. However, the results for the other algorithms are similar enough that we do not go into further detail for them.

Categorical shift. Considering categorical shift errors (see Figure 5a), COMET achieves a higher F1 score throughout the cleaning process compared to FIR, RR, and CL except for a few outliers (EEG, being numerical, does not contain categorical shift errors). COMET shows the most significant advantage compared to the baselines in the S-Credit dataset. Here, COMET achieves a F1 score advantage for the baselines of up to 0.08 (which is equivalent to an advantage of 8 *percentage points* (%*pt*)). The higher the budget invested, the greater the advantage between COMET and the baselines.

Figure 7 shows an example of the cleaning process of S-Credit considering one pre-pollution setting. Throughout the cleaning budget range, COMET maintains up to a 11%*pt* higher F1 score over the considered baselines. Notably, between budgets 6 and 32, COMET’s F1 score even surpasses the Oracle, underscoring the effectiveness of its cleaning recommendations. However, the baselines and even COMET show a fluctuating behavior. Cleaning categorical shift errors causes the cleaned feature distributions to become more aligned with their true categories, which might

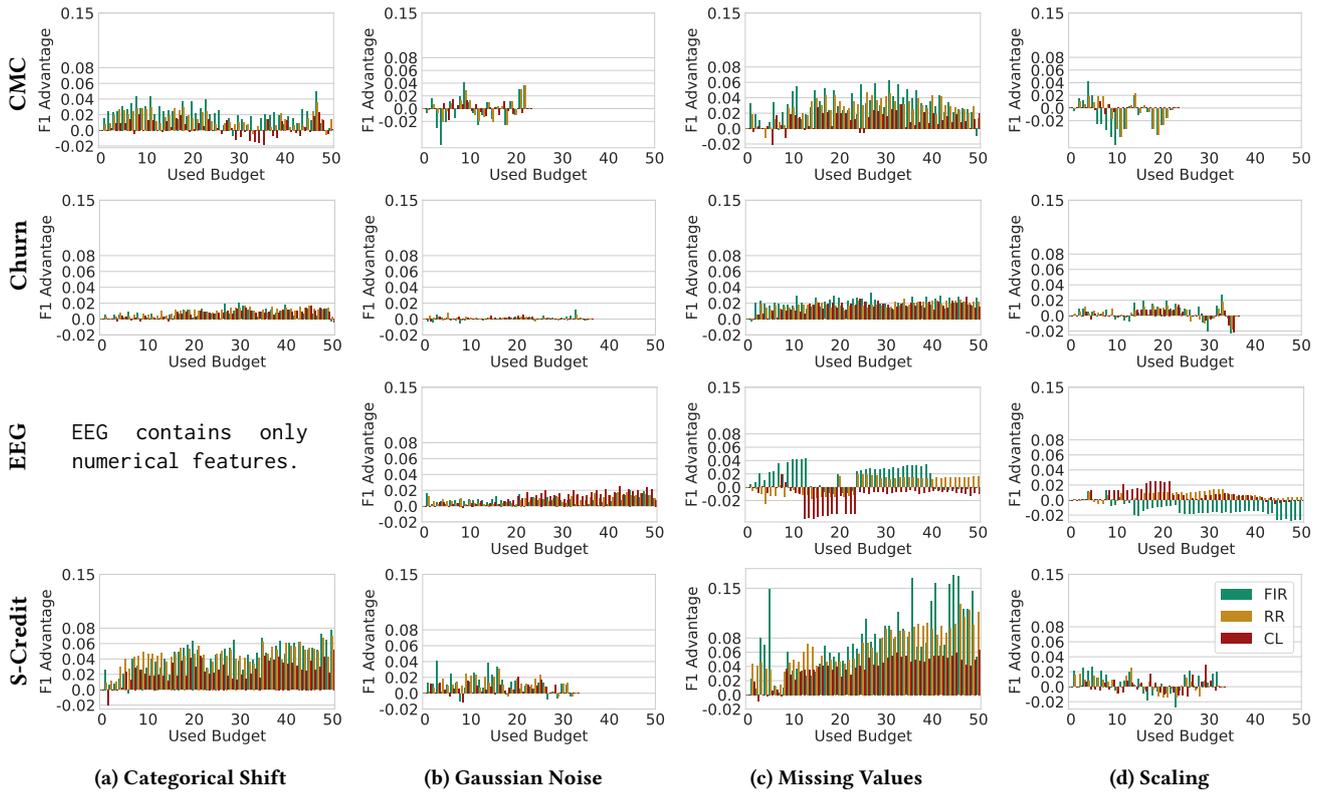


Figure 5: Comparison of COMET with the baselines FIR, RR and CL for MLP across error types.

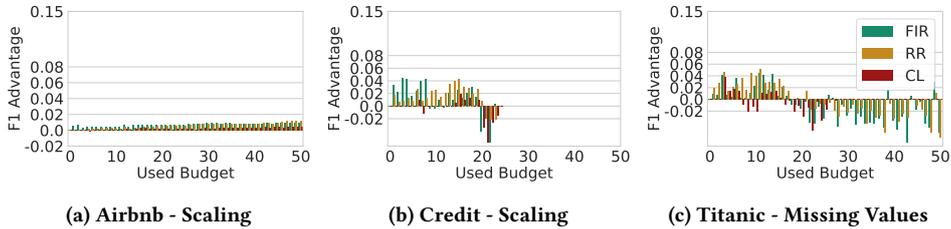


Figure 6: Comparison of COMET with the baselines FIR, RR and CL for MLP across error types, for datasets from CleanML.

cause the model to adjust how it weighs other features. This adjustment can result in unpredictable, temporary dips in performance.

The horizontal “cleaned” line in the figure represents the scenario where the dataset is completely clean, independent of the budget, showing that Oracle and COMET achieve higher F1 scores than complete cleaning for budgets exceeding 6.

Gaussian noise. COMET demonstrates superior performance compared to the baselines in handling Gaussian noise, albeit with a slightly less pronounced advantage than observed for categorical shifts (see Figure 5b). In most cases, COMET’s recommendations lead to F1 scores up to 4%pt higher than FIR, RR and CL. EEG shows, especially for CL, an increasing trend in the F1 score difference with a higher budget. For CL, the initial feature ranking appears beneficial during the first few iterations, resulting in a smaller advantage for COMET over CL. However, this ranking becomes outdated in later iterations and no longer represents an optimized cleaning order – the advantage from COMET increases. In the Churn dataset, the difference between COMET and the baselines is minimal, with COMET achieving only

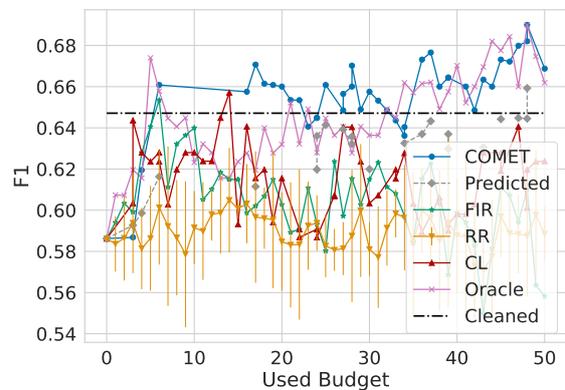


Figure 7: Example of COMET’s advantage over FIR and RR, considering S-Credit with categorical shift errors and MLP.

slightly higher F1 scores due to the small difference (max 1.5%*pt*) between the dirty and cleaned states. Therefore, no significant advantage of COMET over the baselines can develop during the cleaning process. The CMC dataset shows an outlier of COMET. Here, the baselines FIR and RR briefly achieve an advantage of up to 5%*pt* over COMET.

Missing values. COMET outperforms the baselines in most experiments, achieving a F1 score advantage in most cleaning steps, as shown in Figure 5c and 6c. In S-Credit, the F1 score difference again increases with a higher budget, with COMET occasionally leading by over 15%*pt*. This advantage is also evident in other datasets. In Churn, the advantage similarly increases, while CMC shows a bell curve trend: the F1 score advantage increases up to 6%*pt* for half the budget, then gradually decreases, meaning that the baselines compensate previous wrong decisions in later iterations. The EEG results show a sudden drop in the advantage between COMET and FIR, respectively between COMET and CL, from a budget of 12. FIR maintains its advantage until a budget of 22, after which COMET recovers and surpasses FIR, while CL continues to maintain its superiority. However, from a budget of 40, the difference between COMET and FIR narrows again, mainly due to a specific pre-pollution setting in EEG. This pattern can be explained by the fact that each step of the cleaning process builds upon the previous one. As a result, there is a risk that COMET may recommend a suboptimal feature at certain stages, causing prediction accuracy to stagnate temporarily. This stagnation can last for several cleaning steps before a noticeable improvement is observed. This phenomenon also occurs in other experiments, where the considered baseline or AC is superior, as the incremental nature of the cleaning process sometimes leads to delayed improvements in accuracy. The results for Titanic represent an outlier in terms of COMET’s performance compared to the baselines. COMET outperforms the baselines up to a budget of 17, but the baselines are superior afterward, with an advantage of up to 5%*pt*.

While the baselines perform better in a few budget-cases, COMET remains the superior choice in most cleaning scenarios.

Scaling. The trends of the individual methods for scaling errors are similar to those of Gaussian noise (see Figure 5d). Overall, COMET again shows an advantage over the baselines, although this advantage is smaller than categorical shifts and missing values (see Figures 6a and 6b). When considering scaling errors, COMET achieves an F1 score advantage of up to 4%*pt* compared to FIR, RR and CL. Similar to the previously considered results, a trend can be seen for Airbnb: the difference between COMET and the baselines increases as the invested budget increases. However, there are also some fluctuations. In CMC, the baselines FIR and RR generally outperform COMET, and in EEG, FIR outperforms COMET from a budget of 11, though COMET keeps its superiority over RR and CL.

5.3 Comparison to AC for a single error type

We compared COMET’s performance with ActiveClean ([20], RQ 3), focusing on AC-SVM among the many configurations and pre-pollution settings considered, in a single-error scenario. Similar to RQ 2, we also assume constant costs for each cleaning step in RQ 3, regardless of the error type. Figures 8 and 9 show COMET’s performance per budget, comparing it to AC for individual error types with AC-SVM. The overall trends are consistent with those observed for other ML algorithms (LIR, LOR).

Our implementation of AC is based on the code published by the authors¹, which contains only the basic functions of AC; we extended it by the component of gradient-based selection of records. AC’s approach also includes an error detection component, which we skip in the experiments. Furthermore, AC always assumes record-wise data cleaning, while COMET assumes feature-wise cleaning. To integrate AC into our setting, we first pre-train the respective ML model with the records that are already clean according to the pre-pollution settings. We then follow AC’s approach by selecting a sample of records according to the gradients and cleaning the records across all features. The sample corresponds to the size of a cleaning step. Though the number of cleaned entries per iteration or budget may differ, the discrepancies are minor due to our assumed equal error distribution.

Categorical shift. Except for a few minor deviations in CMC, COMET generally outperforms AC for all considered datasets, as Figure 8a shows. The F1 score *advantage* reaches up to 40%*pt* in certain datasets like Churn. The fluctuating F1 score advantages highlight that the performance of AC is quite erratic. Our analysis shows that, for example in S-Credit, the F1 score can drop by up to 30%*pt* after a cleaning step, only to recover after a few further cleaning steps. In contrast, COMET shows a steady increase, with F1 scores improving by up to 1%*pt*, and far lower variance.

Gaussian noise. COMET also outperforms AC when dealing with Gaussian noise, although the advantage per used budget are smaller (Figure 8b). Apart from CMC, COMET is almost always preferred over AC. The highest difference can be seen with S-Credit, where COMET achieves a gap of up to 40%*pt*, though this advantage narrows with a higher invested budget. In CMC, however, AC periodically outperforms COMET, particularly with a budget between 5 and 13, where AC holds a 10%*pt* advantage, mainly again due to sudden F1 score changes.

Missing values. When the data includes missing values, COMET outperforms the AC in most cleaning steps (see Figures 8c and 9c). Notably, COMET shows the most consistent performance across all cleaning steps in the S-Credit dataset, maintaining a performance gap of 20%*pt* to 35%*pt* throughout the cleaning process. For the Churn dataset, COMET leads AC by up to 40%*pt*. However, in the EEG and CMC datasets, COMET’s advantage over AC is less pronounced, peaking at around 20%*pt*. In EEG, AC frequently outperforms COMET due to one pre-pollution setting. For the Titanic dataset, performance differences vary abruptly (Figure 9c). Once more, no consistent pattern relative to the cleaning budget emerges, even when using COMET throughout the cleaning steps.

Scaling. The performance of COMET in addressing scaling errors mirrors its effectiveness with missing values. In the EEG dataset, AC slightly outperforms COMET, a trend also occasionally observed in the CMC dataset where AC demonstrates superiority. In the credit dataset, COMET and AC perform similarly, except in two cases where COMET excels. However, COMET distinctly outperforms AC in these outlier scenarios. For other datasets, COMET consistently maintains its advantage over AC throughout the cleaning steps. The Airbnb dataset shows a pattern similar to S-Credit, with COMET steadily leading AC, except for a few exceptions.

¹https://www.dropbox.com/sh/r2vv252m5lnqpm/AAAMj0WRaZX9EKH_8dLOHQpla?e=3&dl=0

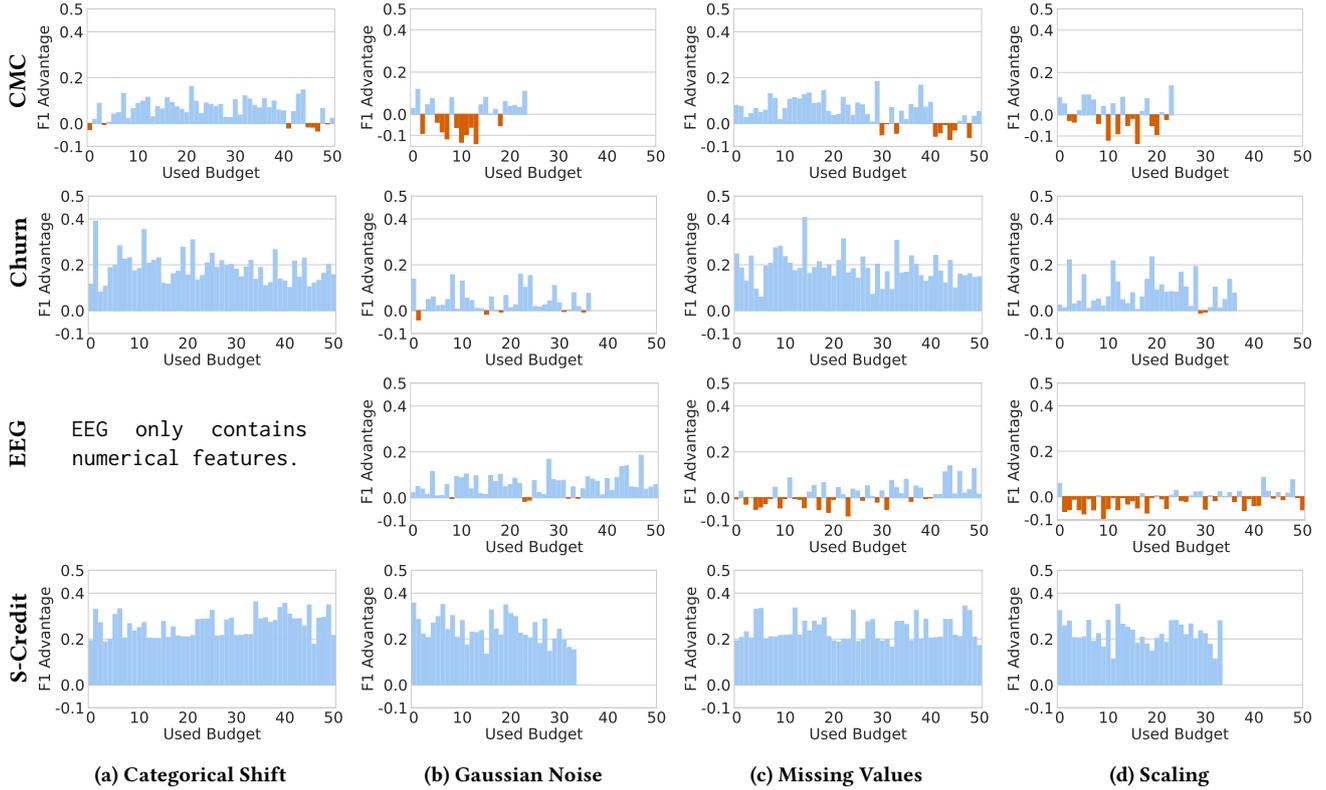


Figure 8: Comparison of COMET with AC for AC-SVM across error types.

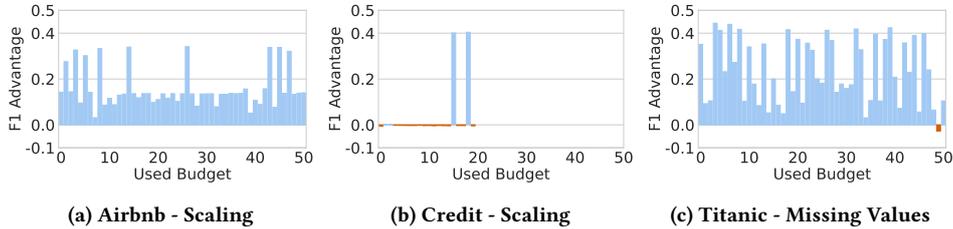


Figure 9: Comparison of COMET with AC for AC-SVM across error types, for datasets from CleanML.

5.4 Overall performance per Baseline

For our fourth research question (RQ 4), we examine the performance of COMET across our considered error types respectively ML algorithms (see Figure 10a and b). The experiments help to understand the adaptability and robustness of our approach regardless of a specific error type and invested budget.

Performance grouped by ML algorithm. Figure 10a summarizes the F1 score differences grouped by ML algorithm, respectively aggregated across all error types. Here, we include our experiment results from the multi-error and single-error scenario. Each bar shows the mean advantage between FIR, RR, and AC compared to COMET, highlighting its good performance independent of error type and invested budget.

It should be noted again that we tested only the ML algorithms LIR, LOR and AC-SVM in connection with AC, while we tested only GB, KNN, MLP and SVM in connection with FIR, RR, and CL. The results align with trends from the previous sections, showing that COMET’s advantage over AC is significantly greater than over FIR, RR, and CL (see Figure 10a). In particular, LIR achieves

the highest average F1 score advantage of 24%pt. However, the results for LOR and AC-SVM also emphasize the superiority of COMET with an average F1 score advantage between 12%pt and 15pt. The differences between COMET compared to FIR, RR, and CL are smaller, with SVM showing the highest average difference of 3%pt and the other ML algorithms (MLP, GB, KNN) ranging between 1%pt and 2%pt.

Performance grouped by error type. For our third research question (RQ 4), we examine COMET’s performance compared to the baselines, considering the mean F1 scores differences, but now grouped by error type, respectively aggregated across the ML algorithms (Figure 10b). This analysis shows how COMET performs per error type, regardless of the ML algorithm and budget. Consequently, this perspective only includes the results from the single-error scenario.

COMET consistently outperforms the baselines across all error types, with minor variations. It achieves the highest advantage of 6%pt in categorical shift errors, followed by 5%pt in missing value errors. For Gaussian noise and scaling errors, the advantage is slightly lower, between 3%pt and 4%pt. Again, the results are

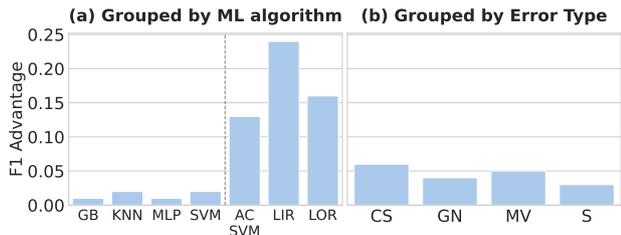


Figure 10: Overall performance of COMET. (CS – Categorical Shift, GN – Gaussian Noise, MV – Missing Values, S – Scaling).

consistent with the trends from the previous sections, which show that COMET performs slightly better compared to the baselines in the presence of categorical shifts or missing values than in the presence of Gaussian noise or scaling errors.

Conclusion. The experiments conducted to address RQ1-RQ4 demonstrate that, with few exceptions, COMET outperforms the baselines FIR, RR, and CL in the cleaning process, achieving up to a 52%pt and, on average, 5%pt advantage. COMET optimizes the cleaning process compared to these baselines by giving the *Cleaner* feature-wise cleaning recommendations.

COMET’s superior performance in enhancing prediction accuracy stems from two key factors. First, COMET optimizes directly for a specific prediction metric, the F1 score, ensuring alignment to improve model performance. This targeted focus allows for more precise prediction quality improvements. Second, the prediction model of COMET is adaptive. It assesses the impact of cleaning per iteration by incrementally polluting the data to derive the effect of cleaning a particular feature. COMET adapts its predictive model individually to the circumstances per iteration and feature. This method contrasts with FIR and CL, which rely on static or outdated information from previous dataset states, potentially leading to misguided cleaning efforts. Likewise, AC, although it adapts using gradients of dirty records and prior average feature changes, also relies on outdated information, leading to suboptimal cleaning decisions. The same applies trivially to RR, which randomly selects features for the next cleaning step.

5.5 Prediction accuracy

COMET predicts the feature-wise impact of data cleaning on ML prediction accuracy. This section focuses on how closely the estimated F1 scores per cleaning step aligns with the actual outcomes for the features that COMET recommended for cleaning (RQ 5).

We assess the predictive accuracy using the Mean Absolute Error (MAE) between predicted and the actual F1 scores after cleaning, focusing on the predictions that the *Recommender* finally used for cleaning. Figure 11 presents the MAE of COMET grouped by error type and ML Algorithm across all datasets.

The MAE ranges from 0.0007 to 0.05, demonstrating COMET’s robust predictive model and ability to guide informed cleaning decisions. However, there are variations across ML algorithms and error types. For instance, the MAE for MLP ranges from 0.007 for scaling errors to 0.015 for categorical shifts. Compared to the overall minimum MAE of 0.0007 (by GB and categorical shift), COMET’s predictions are less accurate when MLP is used. COMET behaves similarly with AC-SVM, where the MAE is between 0.009 (Gaussian noise) and 0.015 (missing values).

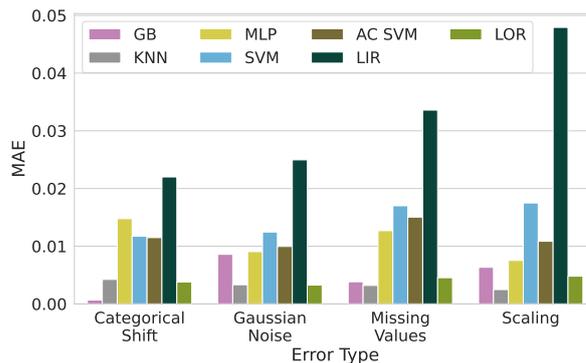


Figure 11: MAE of COMET’s predictions.

Figure 7 shows that sudden jumps in F1 scores, especially for categorical shifts, lead to deviations between predicted and actual F1 scores. In Figure 7, initial predictions match actual F1 scores, but sudden jumps at budgets 4 and 24 cause significant deviations. The *Recommender* compensates occasional inaccuracies by temporarily restoring a feature’s cleaning step, optimizing the process so that COMET outperforms the baselines.

Figure 11 also shows that KNN achieves the lowest MAE for all error types except categorical shift, with MAE ranging from 0.002 (missing values) and to 0.003 (scaling): COMET makes particularly accurate predictions when KNN is used, aligning with its superior performance over FIR and RR in most cases.

Conversely, LIR exhibits the highest MAE, varying between 0.02 and 0.05. Despite this, COMET still significantly outperforms AC, with F1 score advantages of up to 50%pt, due to AC’s unpredictable behavior and the restoring strategy of the *Recommender*. While COMET shows variable predictive accuracy across different algorithms, it consistently guides experts, respectively the *Cleaner*, towards more effective cleaning decisions.

5.6 Runtime to produce a recommendation

COMET recommends to the *Cleaner* the next feature to clean until they have completely cleaned the dataset or the cleaning budget is spent. After each recommended cleaning, they wait for the next suggestion before proceeding. In our sixth research question (RQ 6), we examine this waiting time, i.e., the runtime of COMET per iteration. We measured runtimes on a Slurm-managed compute cluster with two AMD EPYC 7742 processors, 64 cores, and 512 GB RAM, allocating up to 35 GB RAM and 40 CPU cores per job. We use Scikit-learn 1.1.3 for ML model training and testing.

Figure 12 shows average runtimes per ML algorithm and error type across all datasets and pre-pollution settings, measured during the first iteration when the full extent of polluted features are considered, thus leading to the highest runtime. COMET shows the shortest runtimes for Gaussian noise (GN) and scaling errors (S) across all ML algorithms, with median runtimes between 230 and 330 seconds (s) and averages between 220 and 500s. The Airbnb dataset is an outlier (in the context of GN and S) due to its high number of numerical features and its number of rows (see Table 1). Since COMET has to analyze the effect of cleaning each feature, the runtime scales with the number of features.

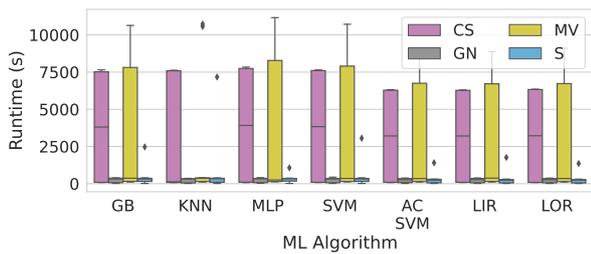


Figure 12: COMET’s runtimes. (CS – Categorical Shift, GN – Gaussian Noise, MV – Missing Values, S – Scaling).

Categorical shift (CS) and missing values (MV) errors exhibit significantly higher runtimes across all ML algorithms, particularly in the upper quartiles, with medians between 125 and 3919s. The churn dataset largely contributes to these high runtimes for CS and MV. Excluding the churn dataset, CS and MV runtimes are typically under 400s. These longer runtimes are likely due to one-hot encoding, which increases the number of features and extends ML model training time.

6 CONCLUSION

Data errors can significantly impact the prediction accuracy of machine learning (ML) models. However, data cleaning in real-world scenarios is often associated with costs and expertise, usually requiring domain experts to configure and execute the cleaning process.

In this work, we introduced COMET, a system that optimizes data cleaning efforts by aligning them with the underlying ML tasks. COMET recommends which feature to clean next, balancing cleaning costs against the potential for enhanced ML prediction accuracy. We evaluated COMET across seven diverse datasets, four ML algorithms, and four types of data errors, benchmarking it against a well-known method and other baselines. Our results show COMET achieves up to 52 and, on average, 5 percentage points higher F1 scores, indicating more efficient data cleaning.

Looking ahead, several directions for future research emerge. One direction involves extending COMET to capture and recommend multiple features to clean within one iteration. Another direction would be to extend COMET for other ML tasks, such as regression and clustering. It will also be valuable to add further error types, such as inconsistent representations or duplicates [26].

Acknowledgements

This research was performed partially in the context of the KITQAR project, supported by Denkfabrik Digitale Arbeitsgemeinschaft im Bundesministerium für Arbeit und Soziales (BMAS).

REFERENCES

- [1] Airbnb 2023. *Airbnb dataset* (Last accessed: 2023-11-19). https://www.dropbox.com/s/nerfrhbrseev928/CleanML-datasets-2020.zip?dl=0&file_subpath=%2Fdata%2FTitanic%2Fmissing_values
- [2] Naomi S Altman. 1992. An introduction to kernel and nearest-neighbor nonparametric regression. *The American Statistician* 46, 3 (1992), 175–185.
- [3] Kenneth Chang. 2023. Quantum Computing Advance Begins New Era, IBM Says. *The New York Times* (2023). <https://www.nytimes.com/2023/06/14/science/ibm-quantum-computing.html>
- [4] Xu Chu, Mourad Ouzzani, John Morcos, Ihab F. Ilyas, Paolo Papotti, Nan Tang, and Yin Ye. 2015. KATARA: Reliable Data Cleaning with Knowledge Bases and Crowdsourcing. *PVLDB* 8, 12 (2015), 1952–1955. <https://doi.org/10.14778/2824032.2824109>
- [5] CleanML 2023. *CleanML/DatasetDescriptions.pdf* at master · chu-data-lab/CleanML (Last accessed: 2023-11-19). <https://github.com/chu-data-lab/CleanML/blob/master/DatasetDescriptions.pdf>

- [6] Corinna Cortes and Vladimir Vapnik. 1995. Support-vector networks. *Machine learning* 20 (1995), 273–297.
- [7] Credit 2023. *Credit dataset* (Last accessed: 2023-11-19). https://www.dropbox.com/s/nerfrhbrseev928/CleanML-datasets-2020.zip?dl=0&file_subpath=%2Fdata%2FTitanic%2Fmissing_values
- [8] Michele Dallachiesa, Amr Ebaid, Ahmed Eldawy, Ahmed K. Elmagarmid, Ihab F. Ilyas, Mourad Ouzzani, and Nan Tang. 2013. NADEEF: a commodity data cleaning system. In *Proceedings of the International Conference on Management of Data (SIGMOD)*. ACM, 541–552. <https://doi.org/10.1145/2463676.2465327>
- [9] Tlameo Emmanuel, Thabiso M. Maupong, Dimane Mpoeleng, Thabo Semong, Banyatsang Mphago, and Oteng Tabona. 2021. A survey on missing data in machine learning. *J. Big Data* 8, 1 (2021), 140. <https://doi.org/10.1186/s40537-021-00516-9>
- [10] Matthias Feurer, Aaron Klein, Katharina Eggenberger, Jost Tobias Springenberg, Manuel Blum, and Frank Hutter. 2015. Efficient and Robust Automated Machine Learning. In *Proceedings of the International Conference on Neural Information Processing Systems (NeurIPS)*. 2962–2970. <https://proceedings.neurips.cc/paper/2015/hash/11d066287202fcd83f79975ec59a3a6-Abstract.html>
- [11] Daniele Foroni, Matteo Lissandrini, and Yannis Velegarakis. 2021. Estimating the extent of the effects of Data Quality through Observations. In *Proceedings of the International Conference on Data Engineering (ICDE)*. IEEE, 1913–1918. <https://doi.org/10.1109/ICDE51399.2021.00176>
- [12] Jerome H. Friedman. 2001. Greedy function approximation: A gradient boosting machine. *The Annals of Statistics* 29, 5 (2001), 1189 – 1232. <https://doi.org/10.1214/aos/1013203451>
- [13] Ulrike Groemping. 2019. South German credit data: Correcting a widely used data set. *Rep. Math., Phys. Chem., Berlin, Germany, Tech. Rep* 4 (2019), 2019.
- [14] Xin He, Kaiyong Zhao, and Xiaowen Chu. 2021. AutoML: A survey of the state-of-the-art. *Knowl. Based Syst.* 212 (2021), 106622. <https://doi.org/10.1016/j.knsys.2020.106622>
- [15] Alireza Heidari, Joshua McGrath, Ihab F. Ilyas, and Theodoros Rekatsinas. 2019. HoloDetect: Few-Shot Learning for Error Detection. In *Proceedings of the International Conference on Management of Data (SIGMOD)*. ACM, 829–846. <https://doi.org/10.1145/3299869.3319888>
- [16] Hans Hofmann. 1994. Statlog (German Credit Data). UCI Machine Learning Repository. DOI: <https://doi.org/10.24432/C5NC77>.
- [17] Ihab F. Ilyas and Theodoros Rekatsinas. 2022. Machine Learning and Data Cleaning: Which Serves the Other? *Journal on Data and Information Quality (JDIQ)* 14, 3 (2022), 13:1–13:11. <https://doi.org/10.1145/3506712>
- [18] Bojan Karlas, Peng Li, Renzhi Wu, Nezihe Merve Gürel, Xu Chu, Wentao Wu, and Ce Zhang. 2020. Nearest Neighbor Classifiers over Incomplete Information: From Certain Answers to Certain Predictions. *PVLDB* 14, 3 (2020), 255–267. <https://doi.org/10.5555/3430915.3442426>
- [19] Sanjay Krishnan, Michael J. Franklin, Ken Goldberg, and Eugene Wu. 2017. BoostClean: Automated Error Detection and Repair for Machine Learning. *CoRR* abs/1711.01299 (2017). arXiv:1711.01299 <http://arxiv.org/abs/1711.01299>
- [20] Sanjay Krishnan, Jiannan Wang, Eugene Wu, Michael J. Franklin, and Ken Goldberg. 2016. ActiveClean: Interactive Data Cleaning For Statistical Modeling. *PVLDB* 9, 12 (2016), 948–959. <https://doi.org/10.14778/2994509.2994514>
- [21] Peng Li, Zhiyi Chen, Xu Chu, and Kexin Rong. 2023. DiffPrep: Differentiable Data Preprocessing Pipeline Search for Learning over Tabular Data. *Proceedings of the International Conference on Management of Data (SIGMOD)* 1, 2 (2023), 183:1–183:26. <https://doi.org/10.1145/3589328>
- [22] 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 *Proceedings of the International Conference on Data Engineering (ICDE)*. IEEE, 13–24. <https://doi.org/10.1109/ICDE51399.2021.00009>
- [23] Tjen-Sien Lim. 1997. Contraceptive Method Choice. UCI Machine Learning Repository. DOI: <https://doi.org/10.24432/C59W2D>.
- [24] Scott M. Lundberg and Su-In Lee. 2017. A Unified Approach to Interpreting Model Predictions. In *Proceedings of the International Conference on Neural Information Processing Systems (NeurIPS)*. 4765–4774. <http://papers.nips.cc/paper/7062-a-unified-approach-to-interpreting-model-predictions.pdf>
- [25] Mark Mazumder, Colby R. Banbury, Xiaozhe Yao, Bojan Karlas, William Gaviria Rojas, Sudnya Frederick Damos, Greg Damos, Lynn He, Alicia Parrish, Hannah Rose Kirk, Jessica Quaye, Charvi Rastogi, Douwe Kiela, David Jurado, David Kanter, Rafael Mosquera, Will Cukierski, Juan Ciro, Lora Aroyo, Bilge Acun, Lingjiao Chen, Mehul Raj, Max Bartolo, Evan Sabri Eyuboglu, Amirata Ghorbani, Emmett D. Goodman, Addison Howard, Oana Inel, Tariq Kane, Christine R. Kirkpatrick, D. Sculley, Tzu-Sheng Kuo, Jonas W. Mueller, Tristan Thrush, Joaquin Vanschoren, Margaret Warren, Adina Williams, Serena Yeung, Newsha Ardalani, Praveen K. Paritosh, Ce Zhang, James Y. Zou, Carole-Jean Wu, Cody Coleman, Andrew Y. Ng, Peter Mattson, and Vijay Janapa Reddi. 2023. DataPerf: Benchmarks for Data-Centric AI Development. In *Proceedings of the International Conference on Neural Information Processing Systems (NeurIPS)*.
- [26] Sedir Mohammed, Lukas Budach, Moritz Feuerpfeil, Nina Ihde, Andrea Nathansen, Nele Noack, Hendrik Patzlaff, Felix Naumann, and Hazar Har-mouch. 2024. The Effects of Data Quality on Machine Learning Performance. <https://arxiv.org/abs/2207.14529>
- [27] Felix Neutatz, Binger Chen, Ziawasch Abedjan, and Eugene Wu. 2021. From Cleaning before ML to Cleaning for ML. *IEEE Data Engineering Bulletin* 44, 1 (2021), 24–41. <http://sites.computer.org/debul/A21mar/p24.pdf>

- [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* 22, 2 (2022), 121–130. <https://doi.org/10.1007/s13222-022-00413-2>
- [29] Felix Neutatz, Mohammad Mahdavi, and Ziawasch Abedjan. 2019. ED2: A Case for Active Learning in Error Detection. In *Proceedings of the International Conference on Information and Knowledge Management (CIKM)*. ACM, 2249–2252. <https://doi.org/10.1145/3357384.3358129>
- [30] Neoklis Polyzotis and Matei Zaharia. 2021. What can Data-Centric AI Learn from Data and ML Engineering? *CoRR* abs/2112.06439 (2021). arXiv:2112.06439 <https://arxiv.org/abs/2112.06439>
- [31] Abdulhakim A. Qahtan, Ahmed Elmagarmid, Raul Castro Fernandez, Mourad Ouzzani, and Nan Tang. 2018. FAHES: A Robust Disguised Missing Values Detector. In *Proceedings of the International Conference on Knowledge discovery and data mining (SIGKDD)*. ACM, New York, NY, USA, 2100–2109. <https://doi.org/10.1145/3219819.3220109>
- [32] Theodoros Rekatsinas, Xu Chu, Ihab F. Ilyas, and Christopher Ré. 2017. HoloClean: Holistic Data Repairs with Probabilistic Inference. *PVLDB* 10, 11 (2017), 1190–1201. <https://doi.org/10.14778/3137628.3137631>
- [33] Oliver Roesler. 2013. EEG Eye State. UCI Machine Learning Repository. DOI: <https://doi.org/10.24432/C57G7>.
- [34] Donald B Rubin. 1976. Inference and missing data. *Biometrika* 63, 3 (1976), 581–592.
- [35] S. Schelter, T. Rukat, and F. Biessmann. 2021. JENGA - A Framework to Study the Impact of Data Errors on the Predictions of Machine Learning Models. In *Proceedings of the International Conference on Extending Database Technology (EDBT)*. OpenProceedings.org, 529–534. <https://doi.org/10.5441/002/edbt.2021.63>
- [36] Donald F Specht. 1991. A general regression neural network. *IEEE Transactions on Neural Networks* 2, 6 (1991), 568–576.
- [37] Telco 2023. *Telco Customer Churn* (Last accessed: 2023-11-22). <https://www.kaggle.com/datasets/blastchar/telco-customer-churn>
- [38] Titanic 2023. Titanic dataset (Last accessed: 2023-11-19). https://www.dropbox.com/s/nerfrhrseev928/CleanML-datasets-2020.zip?dl=0&file_subpath=%2Fdata%2FTitanic%2Fmissing_values
- [39] Steven Euijong Whang, Yuji Roh, Hwanjun Song, and Jae-Gil Lee. 2023. Data collection and quality challenges in deep learning: a data-centric AI perspective. *VLDB Journal* 32, 4 (2023), 791–813. <https://doi.org/10.1007/s00778-022-00775-9>
- [40] Mohamed Yakout, Laure Berti-Équille, and Ahmed K. Elmagarmid. 2013. Don't be SCARED: use SCAlable Automatic REpairing with maximal likelihood and bounded changes. In *Proceedings of the International Conference on Management of Data (SIGMOD)*. ACM, 553–564. <https://doi.org/10.1145/2463676.2463706>
- [41] Daochen Zha, Zaid Pervaiz Bhat, Kwei-Herng Lai, Fan Yang, Zhimeng Jiang, Shaochen Zhong, and Xia Hu. 2023. Data-centric Artificial Intelligence: A Survey. *CoRR* abs/2303.10158 (2023). <https://doi.org/10.48550/arXiv.2303.10158>