

ANOLLM: LARGE LANGUAGE MODELS FOR TABULAR ANOMALY DETECTION

Anonymous authors

Paper under double-blind review

ABSTRACT

We introduce AnoLLM, a novel framework that leverages large language models (LLMs) for unsupervised tabular anomaly detection. By converting tabular data into a standardized text format, we further adapt a pre-trained LLM with this serialized data, and assign anomaly scores based on the negative log likelihood generated by the LLM. Unlike traditional methods that can require extensive feature engineering, and often lose textual information during data processing, AnoLLM preserves data integrity and streamlines the preprocessing required for tabular anomaly detection. This approach can effectively handle mixed-type data, especially those containing textual features. Our empirical results indicate that AnoLLM delivers the best performance on six benchmark datasets with mixed feature types. Additionally, across 30 datasets from the ODDS library, which are predominantly numerical, **AnoLLM performs on par with top performing baselines.**

1 INTRODUCTION

Anomaly detection (AD) seeks to examine specific data points to identify rare, or specious occurrences that deviate from established behavior patterns. AD has been applied to a wide range of applications spanning computer vision (Chandola et al., 2009), natural language processing (NLP) (Schölkopf et al., 2001) and tabular data (Hawkins, 1980). Among these, tabular anomaly detection is particularly crucial, as tabular data is a fundamental format in machine learning which has been used extensively in applications related to cyber-attack prevention (Landauer et al., 2023), detecting fraudulent financial transactions (Dornadula & Geetha, 2019), and diagnosing medical conditions (Fernando et al., 2021).

Conversely, as large language models (LLMs) have shown remarkable performance on various NLP tasks (Radford et al., 2018; Chung et al., 2024; Wei et al., 2022), researchers are keen to understand the capabilities of LLMs applied to other modalities, including tabular data (Fang et al., 2024). While previous studies have demonstrated the proficiency of LLMs in prediction (Dinh et al., 2022), table understanding (Sui et al., 2024), and data generation (Borisov et al., 2023), their effectiveness in handling the task of tabular anomaly detection remains largely unexplored.

Applying LLMs to tabular anomaly detection presents several challenges. First, tabular data is inherently structured in two dimensions, which does not align well with the linear and sequential nature of LLM inputs. Second, unlike traditional tabular classification tasks, unsupervised anomaly detection lacks labels, making the in-context learning framework unfeasible. Third, while existing studies demonstrate that LLMs can be fine-tuned as generative models for tabular data, naively using output probabilities as anomaly scores introduces length-bias issues as the token-level probabilities are aggregated.

In this paper we introduce AnoLLM, a novel framework of using LLMs for unsupervised tabular anomaly detection. AnoLLM is comprised of three phases. In the initial phase, we serialize each row of a tabular dataset into a standardized text format. During the training phase, a pretrained LLM is fine-tuned with the before mentioned serialized tabular data via next-token-prediction, where the LLM learns to be a sequential tabular data generator that models the underlying data distribution. In the inference phase, anomaly scores are determined using the negative log likelihood produced by the LLM for the test samples provided, with higher scores indicating greater surprise by the model when encountering the inputs. Since naively multiplying all token probabilities is prone to length-bias issue, we propose to normalize probabilities for each textual column individually, while

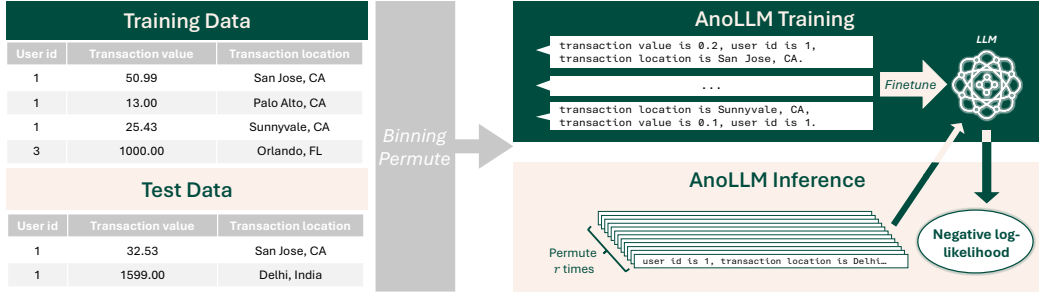


Figure 1: Overall framework of AnoLLM. During the preprocessing stage, numerical columns are binned into groups, and each data row is transformed into a natural language sequence with a randomly shuffled order of columns. In the training stage, a pretrained LLM is fine-tuned using the preprocessed tabular data. During inference, anomaly scores are determined by averaging the negative log-likelihood across r random permutations of the test data.

retaining the raw probabilities for numerical and categorical columns, as supported by theoretical validation.

AnoLLM has the following advantages over existing tabular anomaly detection approaches. First, existing approaches often discard feature names and categorical values when transforming data into vectors, leading to a potential loss of valuable information. Second, these methods rely heavily on feature engineering to handle mixed-type tabular data (Borisov et al., 2022), such as text data, which cannot be easily transformed into numerical and categorical features. For instance, system log anomaly detection requires extensive feature engineering to transform logs into real-valued vectors (He et al., 2017). Moreover, existing approaches often require the use of imputers to handle missing values effectively (Emmanuel et al., 2021). In contrast, using text representations as inputs offers a more flexible and comprehensive representation of data compared to traditional anomaly detection models.

Our contributions are summarized below.

- We propose AnoLLM, a novel framework for adapting LLMs for unsupervised tabular anomaly detection. To the best of our knowledge, AnoLLM is the first tabular AD approach that is capable of handling raw textual features without pre-processing.
- To mitigate the length-bias in the LLM’s output probabilities, we propose a novel method for computing anomaly scores based on LLM-generated probabilities, backed by theoretical validation.
- Empirically, we demonstrate that AnoLLM outperforms all existing methods, including four classical approaches and eight deep learning-based approaches, across six datasets containing mixed tabular data types.
- Despite known limitations of LLMs in performing basic arithmetic calculations (Shen et al., 2023; Lee et al., 2023a), our empirical evidence shows that the overall performance of AnoLLM matches best-performing baselines, k-nearest neighbors, internal contrastive learning, and diffusion time estimation, on 30 datasets in the ODDS library (Rayana, 2016), which primarily contain numerical features.

We also conduct extensive analysis on the key design choices in AnoLLM, including various feature binning methods and the impact of different LLM sizes. To the best of our knowledge, we are the first to apply LLMs to tabular anomaly detection, achieving state-of-the-art results.

2 ANoLLM: LARGE LANGUAGE MODELS FOR MIXED-TYPE TABULAR ANOMALY DETECTION

This section introduces the AnoLLM approach for detecting anomalies in tabular data, which comprises three main stages. An overview of the AnoLLM approach is shown in Figure 1. In the first stage, tabular data is serialized into text format. During the training stage, the serialized data is used

to fine-tune pretrained large language models through next token prediction. In the inference stage, anomaly scores are assigned based on the negative log-likelihood generated by the LLMs for the given test samples. Each stage is described in detail below.

2.1 PROBLEM FORMULATION

We focus on **uncontaminated**, unsupervised anomaly detection using tabular data, represented as $\mathbf{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$, where $\mathbf{x}_i \in \mathcal{X}$ is the i^{th} row of the data. The objective is to train an anomaly detector, $f: \mathcal{X} \mapsto \mathbb{R}$, solely on the **normal** training samples \mathbf{X} . The anomaly detector assigns scores to each sample as $f(\mathbf{x})$, with higher scores indicating a higher likelihood of being an anomaly.

During the test stage, we are provided with a labeled test set, $\mathbf{D}' = \{(\mathbf{x}'_1, y'_1), \dots, (\mathbf{x}'_n, y'_n)\}$, where each $\mathbf{x}'_i \in \mathcal{X}$ and $y'_i \in \{0, 1\}$ with anomalies being labeled as $y = 1$. After applying the anomaly detector f to each test sample, we obtain anomaly scores $\{f(\mathbf{x}'_1), \dots, f(\mathbf{x}'_n)\}$. A threshold, or probability region, can be employed to classify whether a sample is an anomaly based on these scores.

Suppose that we have a tabular dataset with n rows and d columns (features). We denote the column (feature) names as $\mathbf{c} = [c_1, \dots, c_d]$, with each c_i being a natural language text sequence such as “age” or “job description”. Let the value of each entry be $\mathbf{x}_{i,j}$ with $i \in \{1, \dots, n\}$ and $j \in \{1, \dots, d\}$. We note that $\mathbf{x}_{i,j}$ can represent different types of features, such as numerical, categorical or textual features.

2.2 SERIALIZATION OF TABULAR DATA

For serialization, we follow previous tabular data serialization methods (Hegselmann et al., 2023; Borisov et al., 2023) by simply concatenating column names and column values. Specifically, we use the following template: *column name is column value*. Formally, the j^{th} column of the i^{th} sample \mathbf{x}_i is transformed into the text sequence $t_j(\mathbf{x}_i) = “c_j \text{ is } E(\mathbf{x}_{i,j})”$, where $E(\cdot)$ is a text encoder that pre-process the feature into the text form, which is introduced below.

Feature preprocessing: Since LLMs only take text as inputs, we use a text encoder $E(\cdot)$ to transform every row of data into text format. Specifically, for textual columns and categorical columns, we directly use their original form as they are already represented in natural language. For numerical columns, we do standard rescaling and then round them to single-digit decimals. This operation effectively reduces the number of digits when the inputs are very large, or have high precision. Formally, we apply the following affine transformation to each element $\mathbf{x}_{i,j}$.

$$E(x_{i,j}) = \frac{1}{10} \left\lfloor 10 \times \frac{\mathbf{x}_{i,j} - m_j}{z_j} \right\rfloor, \text{ if } c_j \text{ is a numerical column.} \quad (1)$$

Here, m_j and z_j are the mean and the standard deviation of the j^{th} column in the training set, i.e. $m_j = \text{mean}(\mathbf{x}_{1,j}, \mathbf{x}_{2,j}, \dots, \mathbf{x}_{n,j})$, and $z_j = \text{std}(\mathbf{x}_{1,j}, \mathbf{x}_{2,j}, \dots, \mathbf{x}_{n,j})$. $\lfloor \cdot \rfloor$ is a rounding-to-integer operator. This affine transformation normalizes the numerical values and converts them into a decimal number with a single digital place. This operation effectively converts large or long-digit numbers into decimal numbers containing only a few digits. This step is crucial because we obtain the anomaly scores by computing negative log-likelihood, which is sensitive to token lengths.

This standard rescaling operation can be viewed as a feature binning method since most of the numerical values fall into a range centered around the mean 0.0, containing a finite number of single-digit decimal numbers. In this approach, we categorize numerical values and allow AnoLLM to treat these numerical attributes as categorical features. The categories are represented by single-digit decimals that retain information related to magnitude. We choose to use standard rescaling for our binning method since it outperformed other feature binning methods during experimentation, such as equal-width binning and percentile binning, as shown in Section 3.3.

Handling missing values and column names: For all missing feature values, the text encoder E treats them as a separate category and maps them to the word “Unknown.” For tabular data with missing column names, we manually assign names in alphabetical order as follows:

“A”, “B”, ..., “Z”, “AA”, “AB”, ..., “AZ”, “BA”, “BB”, etc. As we will demonstrate in the experiments, AnoLLM remains effective even if the column names are not provided.

Random column permutations: Since there is no inherent ordering among columns, we permute the d serialized text sequences corresponding to each column using a randomly selected permutation $\pi \sim S_d$, where S_d is a symmetric group of d elements, containing all $d!$ possible permutations. The resulting text sequence representing the row \mathbf{x}_i is as follows:

$$\mathbf{T}_c(\pi, \mathbf{x}_i) = \underbrace{c_{\pi(1)} \text{ is } E(\mathbf{x}_{i,\pi(1)})}_{t_{\pi(1)}(\mathbf{x}_i)}, \underbrace{c_{\pi(2)} \text{ is } E(\mathbf{x}_{i,\pi(2)})}_{t_{\pi(2)}(\mathbf{x}_i)}, \dots, \underbrace{c_{\pi(d)} \text{ is } E(\mathbf{x}_{i,\pi(d)})}_{t_{\pi(d)}(\mathbf{x}_i)}. \quad (2)$$

which contains d permuted columns names and their corresponding values on the i^{th} row. During training, we randomly select a permutation $\pi \sim S_d$ at each gradient step to ensure that the LLMs do not depend on the feature order. For inference, we fix the permutation for all test samples to maintain consistency.

2.3 FINE-TUNING LARGE LANGUAGE MODELS

Finally, we describe the fine-tuning process of a pretrained LLM with serialized tabular data $\mathcal{T} = \{\mathbf{T}_c(\pi, \mathbf{x}_i) \mid \pi \in S_d, i \in \{1, \dots, n\}\}$. For each $\mathbf{s} \in \mathcal{T}$, the serialized tabular data is first transformed into a sequence of tokens, i.e. $\text{Tokenize}(\mathbf{s}) = (w_1, \dots, w_{l(\mathbf{s})})$, of length $l(\mathbf{s})$ using the tokenizer of the given LLM. We then fine-tune the pretrained LLM in an auto-regressive manner using the causal language modeling loss:

$$\mathcal{L}_\theta = \mathbb{E}_{\mathbf{s} \in \mathcal{T}} [\ell_{clm}(\theta, \mathbf{s})] = \mathbb{E}_{\mathbf{s} \in \mathcal{T}} \left[- \sum_{k=1}^{l(\mathbf{s})} \log p_\theta(w_k | w_1, \dots, w_{k-1}) \right], \quad (3)$$

where θ is the parameters of the LLM. By fine-tuning a pretrained LLM with causal language modeling loss, it can learn to predict column values based on the columns that have already been seen. Given that pretrained LLMs have already acquired semantic understanding of the features, it can leverage its extensive contextual knowledge to predict the subsequent column values. For example, the features age of vehicle and vehicle price have clear coherence since a vehicle price decreases as it ages. Furthermore, by representing numerical feature values with decimal numbers, the pretrained LLMs can utilize their basic arithmetic abilities to encode quantity of numerical values (Hanna et al., 2024; Stolfo et al., 2023). This fine-tuning process not only reinforces the LLMs’ comprehension of the specific formatting specified by our serialization method but also enhances their ability to learn the dependencies between different features.

2.4 CALCULATION OF ANOMALY SCORES

Anomaly scores are calculated using the negative log-likelihood of the fine-tuned LLM, as this score reflects the (negative) probability of a sample being generated by the LLM. A higher negative log-likelihood indicates greater surprise from the LLM when predicting the inputs. Assuming that the majority of the training samples are normal, the fine-tuned LLM should predict normal inputs more accurately and assign higher probabilities to them. Since we found that directly computing the negative log-likelihood for textual features, which vary greatly in length, could introduce a length bias, we employ different methods to calculate anomaly scores.

Case 1: Datasets comprised of only numerical and categorical columns: We first discuss how to obtain probabilities of all categories in the categorical column c_j . Given the prefix sequence of the test sample \mathbf{x}'_i , “ c_1 is $E(\mathbf{x}'_{i,1})$, c_2 is $E(\mathbf{x}'_{i,2})$, ..., c_j is”, our objective is to obtain the conditional probabilities of all possible categories, given this prefix.

Theorem 1. *Given a prefix token sequence s_0 , assume that there are only q possible subsequent token sequences, which are $s^{(1)}, s^{(2)}, \dots, s^{(q)}$. Let each sequence $s^{(i)}$ have probability $P(s^{(i)} | s_0)$ to be sampled. Let each sequence $s^{(i)}$ consisting of l_i tokens, denoted as $w_1^{(i)}, w_2^{(i)}, \dots, w_{l_i}^{(i)}$. Then the optimal language model p_{θ^*} that minimizes the causal language modeling loss in Eqn. 3 should align with the true distribution to satisfy the following equation:*

$$\text{for all } i \in \{1, 2, \dots, q\}, p_{\theta^*}(s^{(i)} | s_0) = \prod_{k=1}^{l_i} p_{\theta^*}(w_k^{(i)} | s_0, w_1^{(i)}, \dots, w_{k-1}^{(i)}) = P(s^{(i)} | s_0). \quad (4)$$

When we instantiate $s^{(1)}, s^{(2)}, \dots, s^{(q)}$ with q class of one category and $P(s^{(i)}|s_0)$ with their true conditional probabilities given the previously observed features, the above theorem shows that the product of all token probability resembles the probability of being in the corresponding category even though categorical names may have different lengths for the optimal LLMs. As LLMs have been fine-tuned with extensive tabular data using the same template as in Eqn. 2, they should inherently recognize when to generate categorical data, the LLMs should meet the optimality conditions required for the above theorem to be applicable. Since we treat numerical columns as categorical, the theorem should also apply to numerical columns.

Note that we apply the same template, “ c_1 is $E(\mathbf{x}'_{i,1})$, c_2 is $E(\mathbf{x}'_{i,2})$, \dots , c_j is ”, across all test samples. Therefore, the token probabilities associated with the feature names, such as “ c_j is ”, should be consistent across all test samples. As a result, we can compute anomaly scores by simply summing the log probabilities of every token in the sequence.

Next, given that the fine-tuned LLM has learned to generate tabular data *with arbitrary column orders*, we run the LLM inference r times with different column orders in the test stage to reduce variance. Specifically, we compute anomaly scores using the following formula: given test samples $\mathbf{X}' = \{\mathbf{x}'_1, \dots, \mathbf{x}'_m\}$, the anomaly score of the sample \mathbf{x}'_i is given as follows:

$$\text{score}(\mathbf{x}'_i) = \frac{1}{r} \sum_{j=1}^r \ell_{clm}(\theta, \mathbf{T}_c(\pi_j, \mathbf{x}'_i)) = \frac{1}{r} \sum_{j=1}^r \sum_{k=1}^{l(\mathbf{T}_c(\pi_j, \mathbf{x}'_i))} -\log p_\theta(w'_k | w'_1, \dots, w'_{k-1}), \quad (5)$$

where $\pi_1, \pi_2, \dots, \pi_r \sim S_d$ are different permutations over $\{1, 2, \dots, d\}$, $(w_1, \dots, w_{l(\mathbf{T}_c(\pi_j, \mathbf{x}'_i))}) = \text{Tokenize}(\mathbf{T}_c(\pi_j, \mathbf{x}'_i))$ is the token sequence of the serialized tabular data \mathbf{x}'_i with permutation π_j , and this sequence has length $l(\mathbf{T}_c(\pi_j, \mathbf{x}'_i))$. For consistency, we apply the same permutations $\pi_1, \pi_2, \dots, \pi_r$ to all test samples. The impact of the number of permutations r is analyzed in the experimental section.

Case 2: Datasets containing textual columns: For datasets containing textual features, we observe that Eqn. 5 is heavily influenced by text lengths. Longer texts naturally yield higher anomaly scores due to the summation over more words. To address this, we normalize each textual column by its length separately and sum the scores up. Specifically, we compute anomaly scores using the formula, $\text{score}(\mathbf{x}'_i) = \frac{1}{r} \sum_{j=1}^r \sum_{k=1}^d \ell_{nor}(\theta, \pi_j, t_{\pi_j(k)}(\mathbf{x}'_i))$, with ℓ_{nor} defined below:

$$\ell_{nor}(\theta, \pi_j, t_{\pi_j(k)}(\mathbf{x}'_i)) = \frac{-1}{g(c_{\pi_j(k)})} \log p_\theta \left(E(\mathbf{x}'_{i, \pi_j(k)}) \mid \text{“} t_{\pi_j(1)}(\mathbf{x}'_i), \dots, t_{\pi_j(k-1)}(\mathbf{x}'_i), c_{\pi_j(k)} \text{ is ”} \right). \quad (6)$$

When $c_{\pi_j(k)}$ is a textual column, we set $g(c_{\pi_j(k)})$ to the number of tokens in $E(\mathbf{x}'_{i, \pi_j(k)})$. When $c_{\pi_j(k)}$ is not a textual column, we set $g(c_{\pi_j(k)}) = 1$ so that it is not normalized by the number of tokens. Different from Eqn. 5, we compute the negative log-likelihood only when predicting column values $E(\mathbf{x}'_{i, \pi_j(k)})$, and normalize the loss only for textual columns. We exclude the log-likelihood of the column name, “ $c_{\pi_j(k)}$ is ”, since the fine-tuned LLM recognizes it is as part of the serialization template and therefore should not be included when counting the text length. Normalization is done separately for each column to ensure that columns with longer texts do not disproportionately influence the anomaly scores.

3 EXPERIMENTS

Datasets: Since popular anomaly detection benchmarks, such as ADBench (Han et al., 2022) and the ODDS library (Rayana, 2016), mainly consist of numerical features, we manually collect six datasets that contain mixed types of features. The six datasets are derived from ODDS library (Rayana, 2016), the fraud dataset benchmarks (Grover et al., 2022) and Kaggle. The dataset statistics are described in Table 1. To demonstrate the ability of AnoLLM to accommodate numerical columns, we also evaluate the approach on 30 datasets from the ODDS library, which are mainly composed of numerical features. The ODDS library is collected from various domains, such as chemistry, healthcare, and astronautics. Due to space constraints, we include dataset statistics with respect to the ODDS library in Table 5 of the Appendix.

Datasets	# Samples	# text	# num	# category	# anomaly (%)
Fake job posts (Grover et al., 2022)	17,880	5	3	8	866 (4.84%)
Fraud ecommerce (Grover et al., 2022)	151,112	0	1	6	14,151 (9.36%)
Lymphography (Rayana, 2016)	148	0	3	15	6 (4.05%)
Seismic (Rayana, 2016)	2,584	0	14	4	170 (6.58%)
Vehicle insurance (Kaggle)	15,420	0	8	24	923 (5.99%)
20 newsgroup	11,905	1	0	0	591 (4.96%)

Table 1: Dataset statistics for the six datasets from the mixed-type benchmark. # text, # num, and # category stand for the numbers of textual, numerical, and categorical features in each dataset.

Baselines: We compare against 11 prominent methods in the field of tabular anomaly detection. For classical approaches, we compare against IForest (Liu et al., 2008), PCA (Shyu et al., 2003), KNN (Ramaswamy et al., 2000) and ECOD (Li et al., 2022). These are classical approaches that are still competitive against deep-learning-based methods (Han et al., 2022). For deep-learning based approaches, we compare against DeepSVDD (Ruff et al., 2018), RCA (Liu et al., 2021), and self-supervised learning based approaches including SLAD (Xu et al., 2023b), GOAD (Bergman & Hoshen, 2020), NeuTral (Qiu et al., 2021), ICL (Shenkar & Wolf, 2022), DTE (Livernoche et al., 2024) and REPEN (Pang et al., 2018). All methods are implemented with identical dataset partitioning.

Implementation details: We use SmoLLM-135M and SmoLLM-360M (Allal et al., 2024) as our backbone LLM since they represent state-of-the-art small models among open-weights LLMs. We choose small models (135M and 360M parameters) since the experimental results in Section 3.4 suggest that increasing model sizes does not provide much improvements. Fine-tuning is conducted for 2,000 steps with an AdamW optimizer (Loshchilov & Hutter, 2019) with learning rate 5×10^{-5} across all datasets¹ as the training loss converges uniformly. Batch sizes are adjusted for each dataset to accommodate the varying lengths of serialized data. During inference, we select the number of permutations $r = 21$ since further increasing r does not result in any observed improvement. Fine-tuning and inference are performed on seven Nvidia A100 40GB GPUs hosted on Amazon EC2 P4 Instances. Detailed hyperparameters are shown in Table F of the Appendix.

For baseline implementation, we use PyOD library (Zhao et al., 2019) for shallow methods and DeepOD library (Xu et al., 2023a) for deep-learning-based approaches. We standardize all numerical features to have a mean of zero and a standard deviation of one as AnoLLM does. For categorical features, we group all rare classes with less than 1% of samples together and use one-hot encoding². For textual columns, we use averaged word2vec embeddings (Mikolov et al., 2013) across all words in each column. The word2vec embeddings are 300-dimensional and are trained from the Google news dataset³. For each method, we used the best-performing set of hyperparameters reported in their original paper. For others not specified, we use the default hyperparameters as suggested by DeepOD and PyOD toolkits.

Evaluation protocols: Following prior works (Shenkar & Wolf, 2022; Xu et al., 2023b), we conduct experiments in an uncontaminated, unsupervised setting. The training set consists of a random sample of 50% from the pool of normal examples, with the test set comprising the remaining normal examples, along with all anomalies. We randomly split each dataset using 5 different random seeds and reported the averaged results. We employ Area Under the Receiver Operating Characteristic Curve (AUC-ROC) as our primary evaluation metric. Other metrics, e.g. area under the precision-recall curves (AUC-PRC) and F1 scores, reflect similar trends in our experimental results and are shown in Section H of the Appendix. Moreover, a runtime comparison is shown in Section F of the Appendix.

3.1 MAIN RESULTS ON MIXED-TYPE TABULAR DATASETS

Table 2 shows the overall performance of different methods on the six mixed-type datasets. As shown in the table, AnoLLM consistently delivers the best results on all datasets. Specifically, AnoLLM with the SmoLLM-360M backbone exhibits at least a 6.8% improvement over the baseline

¹For the fake job post dataset, the model is fine-tuned for 10,000 steps due to a longer convergence time.

²For SLAD, we use ordinal encoding for categorical features as it does not support one-hot encoding.

³<https://code.google.com/archive/p/word2vec/>

Methods \ Datasets	Fake job posts	Fraud ecommerce	Lymphography	Seismic	Vehicle insurance	20news groups	Average
Classical methods							
Iforest	0.755	0.501	0.587	0.692	0.496	0.623	0.609
PCA	0.724	0.647	0.811	0.692	0.509	0.623	0.668
KNN	0.636	1	0.838	0.738	0.524	0.605	0.724
ECOD	0.512	0.755	0.762	0.692	0.509	0.62	0.642
Deep learning based methods							
DeepSVDD	0.561	1	0.87	0.713	0.505	0.597	0.708
RCA	0.629	1	0.907	0.727	0.531	0.546	0.723
SLAD	0.603	0.998	0.948	0.714	0.556	0.64	0.743
GOAD	0.566	0.998	0.849	0.717	0.512	0.63	0.712
NeuTral	0.548	1	0.851	0.681	0.507	0.658	0.708
ICL	0.699	1	0.769	0.719	0.501	0.671	0.727
DTE	0.548	1	0.909	0.714	0.512	0.6	0.714
REPEN	0.653	1	0.8	0.724	0.513	0.574	0.711
AnoLLM							
SmolLM-135M	0.800	1	0.968	0.712	0.569	0.766	0.803
SmolLM-360M	0.814	1	0.995	0.746	0.555	0.752	0.810

Table 2: AUC-ROC scores for all methods on the six datasets containing mixed types of features.

methods. This performance is especially notable on datasets with textual columns, such as the fake job posts and 20 newsgroups datasets, where AnoLLM substantially outperforms other methods. This highlights the advantage of utilizing LLMs for anomaly detection, as they can process raw text inputs effectively and do not require the extensive feature engineering for textual inputs that is necessary for baseline methods.

3.2 RESULTS ON THE ODDS BENCHMARK

Figure 2 presents the overall performance of these different methods on 30 tasks in the ODDS library. The results for each dataset and standard error are provided in Table 10 and Table 11 in Appendix. A critical difference diagram showing statistical significance is included as Figure 7 in the Appendix. We note that in ODDS benchmarks, over 98.5% of columns are numerical, 0% of columns are textual, and only 10 out of 30 datasets contain human-understandable column names, which diminishes the strengths of LLMs. However, by using pseudo feature names and feature binning, AnoLLM with SmolLM-135M backbone performs comparably to the competitive baselines, KNN, ICL, and DTE. These results highlight AnoLLM’s ability to effectively handle numerical data and its robustness to missing column names. These findings are consistent with previous studies on tabular anomaly detection which indicates that KNN, ICL, and DTE yield top performance (Livernoche et al., 2024).




Figure 2: Averaged AUC-ROC scores with standard error bars for all methods over 30 datasets in ODDS. Colour scheme: yellow (shallow methods), purple (deep learning methods), dark blue (ours).

3.3 EFFECTS OF FEATURE BINNING

In this experiment, we systematically study different options for feature binning strategies for numerical columns. Specifically, we choose to use the 30 datasets in ODDS library since they are primarily comprised of numerical columns. All methods use the same pretrained LLM, smolLM-135M. We compare the standard rescaling method as in Eqn. 1 with four alternative methods listed below.

Equal-width binning: The range of each column is divided into 10 intervals, each with the same width. The resulting bins are labeled with single-digit decimal numbers, ranging from 0.1 to 1.0.

Quantile binning: All numerical columns are divided into 10 bins with each bin contain approximately 10% of samples. The bins are labeled by percentiles, from “the 10th percentile” to “the 100th percentile”.

Language: This method follows the same approach as equal-width binning, but the bin labels are replaced with 10 adjectives in natural language, representing magnitudes: “minimal”, “slight”, “moderate”, “noticeable”, “considerable”, “significant”, “substantial”, “major”, “extensive”, and “maximum”.

No binning: Numerical features are used in their raw and unaltered form.

Table 3 shows the averaged performance of each method over 30 datasets in the ODDS library. The results show that the standard rescaling method delivers the best performance, followed by equal-width binning and its language-based variant. In contrast, both no binning and quantile binning show comparatively lower performance, highlighting the importance of linear rescaling and rounding. This is evident as linear rescaling methods, such as equal-width binning and standard rescaling, outperform approaches that rely on raw numerical values, which may include long digit numbers that negatively impact overall performance. Moreover, replacing bin labels with adjectives has minimal effect on overall performance. Detailed comparison of each dataset can be seen in Table 16 in Appendix.

Methods	AUC-ROC
Equal-width	0.865
Quantile	0.752
Language	0.863
No binning	0.800
Standard	0.884

Table 3: Comparison on different binning methods for numerical columns on ODDS library.

3.4 IMPACT OF LLM SIZES

In this experiment, we discuss the impact of LLM sizes to the AnoLLM framework. We compare LLMs in the `SmallLM` family as they share the same pre-training corpus and tokenizer, with the only difference being the number of model parameters⁴. Due to GPU memory constraints, we use a LoRA adapter (Hu et al., 2022) for the 1.7B model instead of full fine-tuning.

Table 4 shows the averaged performance for all LLM sizes on the mixed-type benchmark and the ODDS library. Detailed performance on each individual dataset is provided in Table 17 in the Appendix. The results suggest that using the 1.7B model does not provide much performance boost. This could be because larger models are trained on text data that are not relevant to our tabular tasks. As a result, the increased model capacity and contextual knowledge do not lead to any improvements. Therefore, we recommend that practitioners use smaller models, which offer a better balance of efficiency and effectiveness.

LLM sizes	Mix-typed	ODDS
135M	0.803	0.884
360M	0.811	0.865
1.7B	0.812	0.861

Table 4: Comparison of different LLM sizes across the six mixed-type tabular benchmarks and the ODDS benchmark, with the numbers representing AUC-ROC scores.

4 RELATED WORKS

Unsupervised Anomaly Detection for Tabular data: Tabular anomaly detection is a long-standing problems and number of approaches, including classic methods and deep-learning based methods, are proposed over past several decades. One line of classical methods compute anomaly scores based on density, e.g. kernel density estimators (Latecki et al., 2007), local density (Breunig et al., 2000), and gaussian mixture models (Yang et al., 2009). Other methods include isolation forest (Liu et al., 2008), empirical-cumulative-distribution-based outlier detector (ECOD) (Li et al., 2022), One-class support vector machine (Schölkopf et al., 1999), and k-nearest neighbors

⁴We note that `SmallLM-1.7B` uses 1T tokens and 135M and 360M models only use 600M tokens for pre-training.

(KNN) (Ramaswamy et al., 2000). Notably, KNN remains a strong baseline even when numerous deep learning based approaches have been proposed (Livernoche et al., 2024).

Deep-learning-based approaches can be categorized into two groups, margin-based approaches and self-supervised learning (SSL) based approaches. Margin-based approaches, e.g. DeepSVDD (Ruff et al., 2018) and DROCC (Goyal et al., 2020), employ neural networks to map normal data into hyperspace with minimal volume. SSL-based approaches typically define pseudo-tasks where normal data is expected to perform better than abnormal data. As a result, these methods often incorporate SSL loss functions into their objectives. For example, NeuTral (Qiu et al., 2021) and REPEN (Pang et al., 2018) employ contrastive loss to bring the embeddings of similar data closer, while distancing the embeddings of dissimilar data. SLAD (Xu et al., 2023b) and ICL (Shenkar & Wolf, 2022) both construct pseudo tasks by splitting the input vectors. Please refer to Section C in Appendix for more details. We note that all the above methods are not able to handle raw texts and focus on numerical features only. Their applications to categorical data are not thoroughly discussed (Taha & Hadi, 2019).

LLMs for Anomaly Detection: With the rising popularity of LLMs, recent studies have explored their application in various anomaly detection (AD) tasks. In time-series AD, Liu et al. (2024a) demonstrate that LLMs can deliver accurate and interpretable results through well-designed prompts. Alnegheimish et al. (2024) investigate the zero-shot performance of LLMs in time-series AD. For image data, several studies (Gu et al., 2024; Cao et al., 2023; Zhu et al., 2024; Yang et al., 2024c) examine the performance of visual-linguistic models in visual and video AD across different settings. Additionally, Jin et al. (2024) and Liu et al. (2024b) show that large visual-linguistic models can effectively detect fake news using multi-modal inputs. For system log data, several studies (Han et al., 2023; Yamanaka et al., 2024; Lee et al., 2023b; Hadadi et al., 2024) have fine-tuned pretrained LLMs to detect anomalies in system logs.

For tabular data, a concurrent study (Li et al., 2024) examines the zero-shot performance of LLMs. Their approach differs from ours in two ways. First, they focus on zero-shot performance and report weaker results on the ODDS benchmark, whereas we fine-tune LLMs on tabular data, achieving better performance. Second, their method prompts LLMs to detect anomalies for each dimension separately and then aggregates the scores. Their performance lags behind ECOD, which also computes per-feature anomaly scores but does so by computing empirical cumulative distribution. In contrast, our methods are able to leverage contextual dependencies between features and achieve better results than ECOD. Additionally, Biester et al. (2024) and Park (2024) utilize LLMs as agents to generate domain-specific contexts or to format data, such as for financial data processing and data cleaning. Their methods require additional modules and, in some instances, need human intervention to process the outputs produced by the LLMs.

LLMs for Tabular Data: Recent advances in LLMs have spurred extensive exploration of their use in various tabular data tasks, including prediction (Dinh et al., 2022; Hegselmann et al., 2023; Manikandan et al., 2023), synthesis (Borisov et al., 2023; Zhang et al., 2023), feature engineering (Han et al., 2024), and table understanding (Sui et al., 2024). Additionally, recent studies have investigated pre-training foundation models specifically for tabular data (Zhu et al., 2023; Yan et al., 2024; Ye et al., 2024; Yang et al., 2024b). For a comprehensive overview, we refer readers to a survey paper (Fang et al., 2024).

5 CONCLUSION AND FUTURE WORK

We introduce AnoLLM, a novel framework for adapting pretrained LLMs to unsupervised anomaly detection for tabular data. AnoLLM is a robust tabular anomaly detection method that can ingest raw textual features without pre-processing. Empirical evaluations demonstrate that AnoLLM achieves state-of-the-art performance across six benchmark datasets containing mixed feature types, and matches the performance of the leading methods, KNN, ICL, and DTE, on 30 datasets predominantly containing numerical attributes.

One future direction for this work could be reducing the computational overhead associated with AnoLLM, given that LLMs are much less efficient compared to traditional anomaly detection methods. Furthermore, considering AnoLLM’s superior capability in modeling categorical and textual features, an interesting direction for future work could involve utilizing LLMs to extract salient

representations from such features, which could then be leveraged by classical anomaly detection methods to improve efficiency. Furthermore, since LLMs are known to exhibit limitations in numerical reasoning, it is of interest to enhance AnoLLM’s performance on numerical data. One potential improvement could involve the development of specialized tokenization strategies and encoding tailored to numerical attributes.

Finally, as building foundation models for tabular data is an ongoing trend (Van Breugel & Van Der Schaar, 2024; Gardner et al., 2024), AnoLLM represents an initial step in this direction by demonstrating that LLMs can be effectively adapted for tabular anomaly detection. While numerous other tabular tasks remain to be explored, an important avenue for future work is the creation of a general-purpose tabular foundation model capable of addressing a wide range of tabular data challenges.

REFERENCES

- Loubna Ben Allal, Anton Lozhkov, Elie Bakouch, Leandro von Werra, and Thomas Wolf. SmolLM - blazingly fast and remarkably powerful, 2024.
- Sarah Alnegheimish, Linh Nguyen, Laure Berti-Equille, and Kalyan Veeramachaneni. Large language models can be zero-shot anomaly detectors for time series? *arXiv preprint arXiv:2405.14755*, 2024.
- Liron Bergman and Yedid Hoshen. Classification-based anomaly detection for general data. In *International conference on learning representations*, 2020.
- Fabian Biester, Mohamed Abdelaal, and Daniel Del Gaudio. LLMclean: Context-aware tabular data cleaning via LLM-generated odfs. *arXiv preprint arXiv:2404.18681*, 2024.
- Vadim Borisov, Tobias Leemann, Kathrin Seßler, Johannes Haug, Martin Pawelczyk, and Gjergji Kasneci. Deep neural networks and tabular data: A survey. *IEEE transactions on neural networks and learning systems*, 2022.
- Vadim Borisov, Kathrin Seßler, Tobias Leemann, Martin Pawelczyk, and Gjergji Kasneci. Language models are realistic tabular data generators. In *International conference on learning representations*, 2023.
- Markus M Breunig, Hans-Peter Kriegel, Raymond T Ng, and Jörg Sander. Lof: identifying density-based local outliers. In *Proceedings of the 2000 ACM SIGMOD international conference on Management of data*, pp. 93–104, 2000.
- Yunkang Cao, Xiaohao Xu, Chen Sun, Xiaonan Huang, and Weiming Shen. Towards generic anomaly detection and understanding: Large-scale visual-linguistic model (gpt-4v) takes the lead. *arXiv preprint arXiv:2311.02782*, 2023.
- Varun Chandola, Arindam Banerjee, and Vipin Kumar. Anomaly detection: A survey. *ACM computing surveys (CSUR)*, 41(3):1–58, 2009.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. Scaling instruction-finetuned language models. *Journal of Machine Learning Research*, 25(70):1–53, 2024.
- Tuan Dinh, Yuchen Zeng, Ruisu Zhang, Ziqian Lin, Michael Gira, Shashank Rajput, Jy-yong Sohn, Dimitris Papailiopoulos, and Kangwook Lee. Lift: Language-interfaced fine-tuning for non-language machine learning tasks. *Advances in Neural Information Processing Systems*, 35:11763–11784, 2022.
- Vaishnavi Nath Dornadula and Sa Geetha. Credit card fraud detection using machine learning algorithms. *Procedia computer science*, 165:631–641, 2019.
- Tlameo Emmanuel, Thabiso Maupong, Dimane Mpoeleng, Thabo Semong, Banyatsang Mphago, and Oteng Tabona. A survey on missing data in machine learning. *Journal of Big data*, 8:1–37, 2021.

- Xi Fang, Weijie Xu, Fiona Anting Tan, Jiani Zhang, Ziqing Hu, Yanjun Jane Qi, Scott Nickleach, Diego Socolinsky, Srinivasan Sengamedu, Christos Faloutsos, et al. Large language models (llms) on tabular data: Prediction, generation, and understanding-a survey. *Transactions on Machine Learning Research(TMLR)*, 2024.
- Tharindu Fernando, Harshala Gammulle, Simon Denman, Sridha Sridharan, and Clinton Fookes. Deep learning for medical anomaly detection—a survey. *ACM Computing Surveys (CSUR)*, 54(7): 1–37, 2021.
- Joshua P Gardner, Juan Carlos Perdomo, and Ludwig Schmidt. Large scale transfer learning for tabular data via language modeling. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL <https://openreview.net/forum?id=WH5blx5tZ1>.
- Sachin Goyal, Aditi Raghunathan, Moksh Jain, Harsha Vardhan Simhadri, and Prateek Jain. Drocc: Deep robust one-class classification. In *International conference on machine learning*, pp. 3711–3721. PMLR, 2020.
- Prince Grover, Julia Xu, Justin Tittelfitz, Anqi Cheng, Zheng Li, Jakub Zablocki, Jianbo Liu, and Hao Zhou. Fraud dataset benchmark and applications. *arXiv preprint arXiv:2208.14417*, 2022.
- Zhaopeng Gu, Bingke Zhu, Guibo Zhu, Yingying Chen, Ming Tang, and Jinqiao Wang. Anomalygpt: Detecting industrial anomalies using large vision-language models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 1932–1940, 2024.
- Fatemeh Hadadi, Qinghua Xu, Domenico Bianculli, and Lionel Briand. Anomaly detection on unstable logs with gpt models. *arXiv preprint arXiv:2406.07467*, 2024.
- Songqiao Han, Xiyang Hu, Hailiang Huang, Minqi Jiang, and Yue Zhao. Adbench: Anomaly detection benchmark. *Advances in Neural Information Processing Systems*, 35:32142–32159, 2022.
- Sungwon Han, Jinsung Yoon, Sercan Ö. Arik, and Tomas Pfister. Large language models can automatically engineer features for few-shot tabular learning. In *ICML*, 2024. URL <https://openreview.net/forum?id=fRG45xL1WT>.
- Xiao Han, Shuhan Yuan, and Mohamed Trabelsi. Loggpt: Log anomaly detection via gpt. In *2023 IEEE International Conference on Big Data (BigData)*, pp. 1117–1122. IEEE, 2023.
- Michael Hanna, Ollie Liu, and Alexandre Variengien. How does gpt-2 compute greater-than?: Interpreting mathematical abilities in a pre-trained language model. *Advances in Neural Information Processing Systems*, 36, 2024.
- Douglas M Hawkins. *Identification of outliers*, volume 11. Springer, 1980.
- Pinjia He, Jieming Zhu, Zibin Zheng, and Michael R Lyu. Drain: An online log parsing approach with fixed depth tree. In *2017 IEEE international conference on web services (ICWS)*, pp. 33–40. IEEE, 2017.
- Stefan Hegselmann, Alejandro Buendia, Hunter Lang, Monica Agrawal, Xiaoyi Jiang, and David Sontag. Tabllm: Few-shot classification of tabular data with large language models. In *International Conference on Artificial Intelligence and Statistics*, pp. 5549–5581. PMLR, 2023.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. In *International conference on learning representations*, 2022.
- Ruihan Jin, Ruibo Fu, Zhengqi Wen, Shuai Zhang, Yukun Liu, and Jianhua Tao. Fake news detection and manipulation reasoning via large vision-language models. *arXiv preprint arXiv:2407.02042*, 2024.
- Max Landauer, Sebastian Onder, Florian Skopik, and Markus Wurzenberger. Deep learning for anomaly detection in log data: A survey. *Machine Learning with Applications*, 12:100470, 2023.

- Longin Jan Latecki, Aleksandar Lazarevic, and Dragoljub Pokrajac. Outlier detection with kernel density functions. In *International workshop on machine learning and data mining in pattern recognition*, pp. 61–75. Springer, 2007.
- Nayoung Lee, Kartik Sreenivasan, Jason D Lee, Kangwook Lee, and Dimitris Papailiopoulos. Teaching arithmetic to small transformers. *arXiv preprint arXiv:2307.03381*, 2023a.
- Yukyung Lee, Jina Kim, and Pilsung Kang. Lanobert: System log anomaly detection based on bert masked language model. *Applied Soft Computing*, 146:110689, 2023b.
- Aodong Li, Yunhan Zhao, Chen Qiu, Marius Kloft, Padhraic Smyth, Maja Rudolph, and Stephan Mandt. Anomaly detection of tabular data using llms. *arXiv preprint arXiv:2406.16308*, 2024.
- Zheng Li, Yue Zhao, Xiyang Hu, Nicola Botta, Cezar Ionescu, and George H Chen. Ecod: Unsupervised outlier detection using empirical cumulative distribution functions. *IEEE Transactions on Knowledge and Data Engineering*, 35(12):12181–12193, 2022.
- Boyang Liu, Ding Wang, Kaixiang Lin, Pang-Ning Tan, and Jiayu Zhou. Rca: A deep collaborative autoencoder approach for anomaly detection. In *IJCAI: proceedings of the conference*, volume 2021, pp. 1505. NIH Public Access, 2021.
- Fei Tony Liu, Kai Ming Ting, and Zhi-Hua Zhou. Isolation forest. In *2008 eighth ieee international conference on data mining*, pp. 413–422. IEEE, 2008.
- Jun Liu, Chaoyun Zhang, Jiaxu Qian, Minghua Ma, Si Qin, Chetan Bansal, Qingwei Lin, Saravan Rajmohan, and Dongmei Zhang. Large language models can deliver accurate and interpretable time series anomaly detection. *arXiv preprint arXiv:2405.15370*, 2024a.
- Xuannan Liu, Pei Pei Li, Huaibo Huang, Zekun Li, Xing Cui, jiahao.liang, lixiong Qin, Weihong Deng, and Zhaofeng He. FKA-owl: Advancing multimodal fake news detection through knowledge-augmented LVLMS. In *ACM Multimedia 2024*, 2024b. URL <https://openreview.net/forum?id=9UGFbtt44F>
- Victor Livernoche, Vineet Jain, Yashar Hezaveh, and Siamak Ravanbakhsh. On diffusion modeling for anomaly detection. In *International conference on learning representations*, 2024.
- Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International conference on learning representations*, 2019.
- Hariharan Manikandan, Yiding Jiang, and J Zico Kolter. Language models are weak learners. *Advances in Neural Information Processing Systems*, 36:50907–50931, 2023.
- Tomas Mikolov, Kai Chen, G.s Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. *Proceedings of Workshop at ICLR*, 2013, 01 2013.
- Guansong Pang, Longbing Cao, Ling Chen, and Huan Liu. Learning representations of ultrahigh-dimensional data for random distance-based outlier detection. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, pp. 2041–2050, 2018.
- Taejin Park. Enhancing anomaly detection in financial markets with an llm-based multi-agent framework. *arXiv preprint arXiv:2403.19735*, 2024.
- Chen Qiu, Timo Pfommer, Marius Kloft, Stephan Mandt, and Maja Rudolph. Neural transformation learning for deep anomaly detection beyond images. In *International conference on machine learning*, pp. 8703–8714. PMLR, 2021.
- Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. Improving language understanding by generative pre-training. 2018.
- Sridhar Ramaswamy, Rajeev Rastogi, and Kyuseok Shim. Efficient algorithms for mining outliers from large data sets. In *Proceedings of the 2000 ACM SIGMOD international conference on Management of data*, pp. 427–438, 2000.
- Shebuti Rayana. Odds library. <https://odds.cs.stonybrook.edu>, 2016.

- Lukas Ruff, Robert Vandermeulen, Nico Goernitz, Lucas Deecke, Shoaib Ahmed Siddiqui, Alexander Binder, Emmanuel Müller, and Marius Kloft. Deep one-class classification. In *International conference on machine learning*, pp. 4393–4402. PMLR, 2018.
- Bernhard Schölkopf, Robert C Williamson, Alex Smola, John Shawe-Taylor, and John Platt. Support vector method for novelty detection. *Advances in neural information processing systems*, 12, 1999.
- Bernhard Schölkopf, John C Platt, John Shawe-Taylor, Alex J Smola, and Robert C Williamson. Estimating the support of a high-dimensional distribution. *Neural computation*, 13(7):1443–1471, 2001.
- Ruoqi Shen, Sébastien Bubeck, Ronen Eldan, Yin Tat Lee, Yuanzhi Li, and Yi Zhang. Positional description matters for transformers arithmetic. *arXiv preprint arXiv:2311.14737*, 2023.
- Tom Shenkar and Lior Wolf. Anomaly detection for tabular data with internal contrastive learning. In *International conference on learning representations*, 2022.
- Mei-Ling Shyu, Shu-Ching Chen, Kanoksri Sarinnapakorn, and LiWu Chang. A novel anomaly detection scheme based on principal component classifier. In *Proceedings of the IEEE foundations and new directions of data mining workshop*, pp. 172–179. IEEE Press Piscataway, NJ, USA, 2003.
- Alessandro Stolfo, Yonatan Belinkov, and Mrinmaya Sachan. A mechanistic interpretation of arithmetic reasoning in language models using causal mediation analysis. *Empirical Methods in Natural Language Processing*, 2023.
- Yuan Sui, Mengyu Zhou, Mingjie Zhou, Shi Han, and Dongmei Zhang. Table meets llm: Can large language models understand structured table data? a benchmark and empirical study. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining*, pp. 645–654, 2024.
- Ayman Taha and Ali S Hadi. Anomaly detection methods for categorical data: A review. *ACM Computing Surveys (CSUR)*, 52(2):1–35, 2019.
- Boris Van Breugel and Mihaela Van Der Schaar. Position: Why tabular foundation models should be a research priority. In Ruslan Salakhutdinov, Zico Kolter, Katherine Heller, Adrian Weller, Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp (eds.), *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 48976–48993. PMLR, 21–27 Jul 2024. URL <https://proceedings.mlr.press/v235/van-breugel24a.html>.
- Hu Wang, Guansong Pang, Chunhua Shen, and Congbo Ma. Unsupervised representation learning by predicting random distances. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI’20*, 2021. ISBN 9780999241165.
- Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, et al. Emergent abilities of large language models. *Transactions on Machine Learning Research (TMLR)*, 2022.
- Hongzuo Xu, Guansong Pang, Yijie Wang, and Yongjun Wang. Deep isolation forest for anomaly detection. *IEEE Transactions on Knowledge and Data Engineering*, 35(12):12591–12604, 2023a.
- Hongzuo Xu, Yijie Wang, Juhui Wei, Songlei Jian, Yizhou Li, and Ning Liu. Fascinating supervisory signals and where to find them: Deep anomaly detection with scale learning. In *International Conference on Machine Learning*, pp. 38655–38673. PMLR, 2023b.
- Yuuki Yamanaka, Tomokatsu Takahashi, Takuya Minami, and Yoshiaki Nakajima. Logelectra: Self-supervised anomaly detection for unstructured logs. *arXiv preprint arXiv:2402.10397*, 2024.
- Jiahuan Yan, Bo Zheng, Hongxia Xu, Yiheng Zhu, Danny Chen, Jimeng Sun, Jian Wu, and Jintai Chen. Making pre-trained language models great on tabular prediction. In *International conference on learning representations*, 2024.

- Xingwei Yang, Longin Jan Latecki, and Dragoljub Pokrajac. Outlier detection with globally optimal exemplar-based gmm. In *Proceedings of the 2009 SIAM international conference on data mining*, pp. 145–154. SIAM, 2009.
- Yazheng Yang, Yuqi Wang, Guang Liu, Ledell Wu, and Qi Liu. Unitabe: A universal pretraining protocol for tabular foundation model in data science. In *The Twelfth International Conference on Learning Representations*, 2024a.
- Yazheng Yang, Yuqi Wang, Sankalok Sen, Lei Li, and Qi Liu. Unleashing the potential of large language models for predictive tabular tasks in data science. *arXiv preprint arXiv:2403.20208*, 2024b.
- Yuchen Yang, Kwonjoon Lee, Behzad Dariush, Yinzhi Cao, and Shao-Yuan Lo. Follow the rules: Reasoning for video anomaly detection with large language models. *arXiv preprint arXiv:2407.10299*, 2024c.
- Chao Ye, Guoshan Lu, Haobo Wang, Liyao Li, Sai Wu, Gang Chen, and Junbo Zhao. Towards cross-table masked pretraining for web data mining. In *Proceedings of the ACM on Web Conference 2024*, pp. 4449–4459, 2024.
- Jiaxin Yin, Yuanyuan Qiao, Zitang Zhou, Xiangchao Wang, and Jie Yang. Mcm: Masked cell modeling for anomaly detection in tabular data. In *The Twelfth International Conference on Learning Representations*, 2024.
- Tianping Zhang, Shaowen Wang, Shuicheng Yan, Li Jian, and Qian Liu. Generative table pre-training empowers models for tabular prediction. In Houda Bouamor, Juan Pino, and Kalika Bali (eds.), *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 14836–14854, Singapore, December 2023. Association for Computational Linguistics. URL <https://aclanthology.org/2023.emnlp-main.917>.
- Yue Zhao, Zain Nasrullah, and Zheng Li. Pyod: A python toolbox for scalable outlier detection. *Journal of Machine Learning Research*, 20(96):1–7, 2019. URL <http://jmlr.org/papers/v20/19-011.html>.
- Bingzhao Zhu, Xingjian Shi, Nick Erickson, Mu Li, George Karypis, and Mahsa Shoaran. Xtab: cross-table pretraining for tabular transformers. In *Proceedings of the 40th International Conference on Machine Learning, ICML’23*. JMLR.org, 2023.
- Jiaqi Zhu, Shaofeng Cai, Fang Deng, Beng Chin Ooi, and Junran Wu. Do llms understand visual anomalies? uncovering llm’s capabilities in zero-shot anomaly detection. In *Proceedings of the 32nd ACM International Conference on Multimedia, MM ’24*, pp. 48–57, New York, NY, USA, 2024. Association for Computing Machinery. ISBN 9798400706868. URL <https://doi.org/10.1145/3664647.3681190>.