

# TabularMark: Watermarking Tabular Datasets for Machine Learning

Yihao Zheng  
Zhejiang University  
Hangzhou, China  
zhengyihao@zju.edu.cn

Haocheng Xia  
University of Illinois  
Urbana-Champaign  
Urbana-Champaign, United States  
hxia7@illinois.edu

Junyuan Pang  
Zhejiang University  
Hangzhou, China  
junyuanpang@zju.edu.cn

Jinfei Liu  
Zhejiang University  
Hangzhou, China  
jinfeiliu@zju.edu.cn

Kui Ren  
Zhejiang University  
Hangzhou, China  
kuiren@zju.edu.cn

Lingyang Chu  
McMaster University  
Hamilton, Canada  
chul9@mcmaster.ca

Yang Cao  
Tokyo Institute of Technology  
Tokyo, Japan  
cao@c.titech.ac.jp

Li Xiong  
Emory University  
Atlanta, United States  
lxiong@emory.edu

## ABSTRACT

Watermarking is broadly utilized to protect ownership of shared data while preserving data utility. However, existing watermarking methods for tabular datasets fall short on the desired properties (detectability, non-intrusiveness, and robustness) and only preserve data utility from the perspective of data statistics, ignoring the performance of downstream ML models trained on the datasets. Can we watermark tabular datasets without significantly compromising their utility for training ML models while preventing attackers from training usable ML models on attacked datasets?

In this paper, we propose a hypothesis testing-based watermarking scheme, TabularMark. Data noise partitioning is utilized for data perturbation during embedding, which is adaptable for numerical and categorical attributes while preserving the data utility. For detection, a custom-threshold one proportion z-test is employed, which can reliably determine the presence of the watermark. Experiments on real-world and synthetic datasets demonstrate the superiority of TabularMark in detectability, non-intrusiveness, and robustness.

## CCS CONCEPTS

• Security and privacy → Database and storage security.

## KEYWORDS

Data Ownership; Watermark; Tabular Dataset; Machine Learning

## ACM Reference Format:

Yihao Zheng, Haocheng Xia, Junyuan Pang, Jinfei Liu, Kui Ren, Lingyang Chu, Yang Cao, and Li Xiong. 2024. TabularMark: Watermarking Tabular Datasets for Machine Learning. In *Proceedings of the 2024 ACM SIGSAC Conference on Computer and Communications Security (CCS '24)*, October 14–18, 2024, Salt Lake City, UT, USA. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3658644.3690373>

## 1 INTRODUCTION

The proliferation of machine learning (ML) has brought significant benefits to a wide range of industries, such as healthcare [26, 40, 49], retail [12, 34, 45], and finance [16, 24, 36]. Structured tabular datasets (or relational data) are prevalent and used across many sectors. However, the nature of datasets allows for near-zero cost replication [42], making them susceptible to unauthorized copying and use. There is a pressing need for a robust mechanism to assert and protect the ownership of such datasets.

Watermarking is a widely adopted technique for asserting ownership and preventing unauthorized usage of shared data and have been widely employed on multimedia data such as images [3, 46, 55] and audios [53, 54], and relational data [27]. A watermark is embedded into original data usually through subtle perturbations. Data owners can extract the watermark from suspicious data to claim ownership. Typically, there are several minimum desired properties of watermarking. 1) detectability: it can be reliably detected, with the aid of some secret information; 2) non-intrusiveness: it should not alter the data in a way that degrades its quality or usability; 3) robustness: it should be resilient to manipulations.

Tabular data poses unique challenges due to its data characteristics considering the desired properties above. 1) Tabular data typically consists of precise values with each entry carrying significant and specific information. There is little to no perceptual redundancy compared to multimedia data, which makes it less flexible for designing a watermark that satisfies both detectability and non-intrusiveness, since even minor changes may significantly impact data integrity or usability. 2) The mixture of different data

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).  
CCS '24, October 14–18, 2024, Salt Lake City, UT, USA

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.  
ACM ISBN 979-8-4007-0636-3/24/10  
<https://doi.org/10.1145/3658644.3690373>

types, including categorical and numerical, may require different or more complex watermarking strategies. 3) Tabular data can undergo a variety of data manipulations like insertions, deletions, and foreign key replacements, and a watermark must be resilient to such operations without being easily removed.

Watermarking schemes [3, 46, 53–55] that mainly focus on multimedia are well-studied in the literature but are difficult to port over for tabular datasets, due to the dependency of intrinsic patterns or semantic information for a specific multimedia type. In order to extend watermarking outside multimedia, many watermarking schemes [2, 23, 25, 28, 30, 32, 43, 44] have been proposed specifically for relational data. However, they are either not generally applicable to all types of tabular data, or fall short in one or more of the above properties. [2, 25] embed watermarks by modifying the least significant bits (LSBs), causing them to be inapplicable to categorical attributes. For example, embedding bits into encoded categorical attributes (e.g., 0-6) can cause significant distortion and may exceed the value range (e.g., modifying 0110 to 0111), violating *non-intrusiveness*. Schemes proposed in [43, 44] opt to embed watermark bits into the statistics of the data. However, they necessitate the use of primary keys in the partitioning algorithms. Primary keys are not essential in tabular datasets, and replacing the original primary key with a new column will prevent the correct extraction of the watermarking information, violating *robustness*. In recent years, studies in [23, 31] focus on reversible watermarking schemes based on histogram shifting. These methods are limited to integer attributes and unsuitable for floating-point attributes [15, 35, 39]. They also lack a theoretical guarantee on the false positive rate for watermark detection, which weakens the *detectability*. Furthermore, nearly all existing methods [2, 25, 28, 30, 32, 43, 44] primarily measure data utility (for non-intrusiveness) on the basic statistics of data, such as mean and variance for query tasks. Given the current prevalent use of tabular data for building ML models, the added watermark should almost not affect the utility of the downstream models, which we refer to as ML utility.

In this paper, we address the limitations of existing techniques by proposing a simple yet effective hypothesis testing-based watermarking scheme for tabular datasets, TabularMark. This scheme partitions data noise into two divisions and introduces designed perturbations to specific cells in the embedding phase. In the detection phase, the deviation distribution characteristics in the suspicious datasets are examined by hypothesis testing.

To ensure detectability, we utilize the one proportion z-test to detect the perturbations introduced in the embedding phase. As different datasets are assumed to be collected independently, the deviation between a non-watermarked dataset and the original dataset should be entirely random within the predefined range. Therefore, we use the one proportion z-test to analyze the characteristics of deviation distribution, which allows for reliable detection of watermarks in watermarked data with a statistically improbable rate of false positives. Additionally, we can adjust the z-test threshold to statistically limit the false positive rate.

To ensure non-intrusiveness, we control the distortion on the ML utility in the embedding phase by managing the number of perturbed cells called *key cells*. Due to randomness in deviation, the probability of the data deviation for a given cell falling within one division is 0.5. Consequently, as the number of key cells increases,

the probability of all of them falling within the chosen division exponentially decreases. Therefore, compared to the total number of cells, only very few key cells are required to embed a robust watermark. Because the deviation does not rely on a specific data type, the embedding and detection of TabularMark can be applied across different types of attributes. Whether numerical attributes or categorical attributes are perturbed, the ML utility of watermarked datasets remains almost unchanged thanks to the small number of perturbed cells.

To ensure robustness, we require the data owner to keep the relevant information from the watermark embedding phase confidential, such as the locations of key cells. Without knowing the key cells, attackers can only randomly perturb a large number of cells, much more than the number of key cells, in an attempt to reduce the z-score, an indicator of the z-test. Therefore, the ML utility of attacked datasets may be significantly reduced. Additionally, we employ the most significant bits (MSBs) of multiple attributes to match key cells in detected datasets to avoid primary key replacement attacks.

Experiments on commonly used real-world datasets validate that the watermark can be reliably detected in watermarked datasets. Additionally, the scheme is proven to be non-intrusive to the ML utility of watermarked datasets and robust against malicious attacks, including insertion, deletion, and alteration attacks. For example, the z-score on the Forest Cover Type increases by 18.6 after watermarking, far exceeding the threshold of 1.96. Even if attackers insert or delete up to 80% of the tuples from the watermarked Forest Cover Type dataset, the watermark can still be detected. Besides, if attackers successfully remove the watermark from the Forest Cover Type dataset by alteration attacks, the XGBoost model trained on the attacked dataset results in an average drop of 0.245 in the  $F_1$ -score compared to the model trained on the original dataset. However, the watermarked dataset only causes an average decrease of 0.001 in the  $F_1$ -score, thereby preserving the ML utility of the dataset. We briefly summarize our contributions as follows.

- We propose a simple yet effective hypothesis testing-based watermarking scheme TabularMark for tabular datasets using domain partition and one proportion z-test. To our best knowledge, this is the first study that utilizes random deviation to almost completely preserve the ML utility of tabular datasets.
- To enhance robustness, we employ a hypothesis testing method and multi-attribute matching to mitigate dependencies on data types and primary keys, respectively. Furthermore, we model the watermark removal and mathematically prove the hardness for attackers to remove the watermark.
- We demonstrate the detectability, non-intrusiveness, and robustness of TabularMark on popular real-world datasets for regression and classification tasks as well as explore the trade-offs among multiple hyperparameters of the scheme on synthetic datasets.

## 2 RELATED WORK

In this paper, we focus on protecting the ownership of static tabular datasets used for training ML models. The most related research is relational database watermarking [2, 23, 25, 28, 30, 32, 43, 44]. In this section, we review the progress in relational database watermarking

and discuss the corresponding limitations. Although there are other types of watermarking schemes, such as image watermarking [3, 46, 55], audio watermarking [53, 54], neural network watermarking [17, 38, 48], and large language model watermarking [14, 29], we do not discuss these related works in detail due to the significant differences in data structure from tabular datasets.

**Non-Reversible Database Watermarking.** The first watermarking scheme [2] for relational databases was proposed in 2002. They embedded watermarks by modifying the LSB (least significant bit) of certain attributes in some tuples and decided whether to embed a ‘0’ or ‘1’ in the LSB based on the parity of a hash value computed from the concatenated primary keys and private keys. Inspired by this work, subsequent studies [21, 50] improved on embedding a single bit by embedding multiple bits within the selected LSBs. However, these methods are less suitable for categorical attributes, as encoded categorical attributes are mostly small-range integers (e.g., 0-6), offering limited watermarking capacity and potentially causing undesirable distortion. Another type of relational database watermarking method opts to embed watermark bits into the statistics of the data. Sion et al. [44] used marker tuples to partition the tuples into different subsets, embedding the corresponding watermark bits in various subsets by modifying subset-related statistical information and utilizing a majority voting mechanism to enhance robustness. Shehab et al. [43] improved upon [44] by first optimizing the partitioning algorithm, avoiding the use of marker tuples for subset partitioning and instead using hash values based on primary keys and private keys, effectively resisting insertion and deletion attacks. Furthermore, to minimize distortion, they modeled the embedding of watermark bits as an optimization problem and offered a genetic algorithm and a pattern search method to solve this problem, effectively controlling data distortion. However, the method [43] is strictly limited by the high requirements for data distribution, and it is challenging to define the optimization problem for categorical attributes with discrete and fixed value ranges. Furthermore, it necessitates the use of primary keys for the partitioning algorithm.

**Reversible Database Watermarking.** In recent years, many studies have focused on reversible watermarking schemes, where the watermark can be extracted from the watermarked data and the original data completely restored. Watermarking techniques based on histogram shifting are a promising solution. [23] introduced a reversible watermark algorithm based on histogram shifting in groups. This method also utilized a message authentication code (MAC) calculated with a private key and a primary key for grouping, defined a statistical quantity “prediction error” for plotting histograms, and shifted the histogram to embed watermark bits in each group. Similar histogram-based methods include [31], which proposed a robust reversible watermarking mechanism based on consecutive columns in histograms. Methods like those in [23, 31] are limited to integer attributes and are not suitable for floating-point attributes which are widely present in tabular datasets for regression tasks. Thus, their application scenarios are constrained. Due to the poor robustness of LSB methods, Li et al. [30] proposed a reversible watermarking method by abandoning the LSB approach and embedding watermark bits into the decimal digits. However, it still has the same issues as the LSB methods.

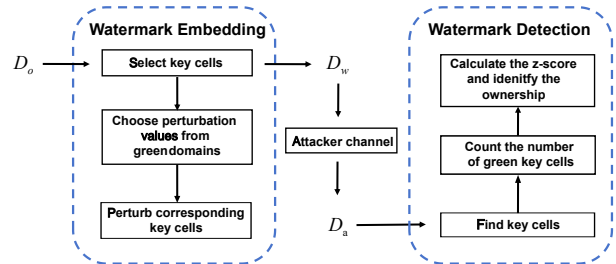
**Table 1: The summary of frequently used notations**

Notation	Description
$n$	Number of cells
$n_w$	Number of key cells
$n_g$	Number of green key cells
$p$	Positive axis boundary of the perturbation range
$k$	Number of unit domains
$\alpha$	Significance level of the test for detecting a watermark
$\beta$	Proportion affected by a certain attack relative to the overall dataset
$\gamma$	Ratio of the lengths of green domains to red domains
$D_o$	Original dataset
$D_w$	Watermarked dataset
$D_s$	Suspicious dataset

These schemes are mainly designed for database management systems where the original database is constantly updated. In contrast, we focus on ML applications where a dataset is not frequently modified, such as static data repositories in UCI Machine Learning Repository [6]. TabularMark is designed based on this setting and solves the above limitations. In addition to the proposed limitations, existing relational database watermarking schemes ignore the effect on the performance of downstream ML models trained on the datasets but only consider changes in the mean or standard deviation of certain attributes for traditional tasks such as query answering. The target of this paper is to reduce the ML utility cost of data owners by limiting the impact of schemes on models trained on the watermarked dataset and to increase the ML utility cost of attackers by aggravating degradation in model performance when removing the watermark.

### 3 ALGORITHMS

In this section, we present the watermark embedding and detection algorithms in detail. The process of TabularMark is illustrated in Figure 1 and involves two stages: watermark embedding and watermark detection. Table 1 summarizes the frequently used notations.



**Figure 1: Flowchart of TabularMark, where  $D_o$  is the original dataset,  $D_w$  is the watermarked dataset,  $D_a$  is the watermarked dataset after suffered attacks.**

#### 3.1 Threat Model

To delimit the scope of this paper, we describe the abilities of attackers as follows.

Considering the characteristics of TabularMark, we assume that the attackers 1) have full access to the watermarked dataset; 2)

are aware of the general process of the watermark embedding and detection algorithms but do not know secret information, such as the distribution of key cells or the specific attributes used for choosing key cells; 3) cannot access the original dataset; and 4) take into consideration not to excessively impact ML utility of the watermarked dataset. For example, in alteration attacks, the range of alteration to the cells would not exceed the perturbation range of the watermark embedding algorithm by a large margin.

### 3.2 Watermark Embedding

In this section, we present the details of the watermark embedding algorithm.

The watermark embedding approach draws inspiration from the random deviations that exist between specific cells within a suspicious dataset and those within the original dataset. This deviation may arise from biases in the measuring instruments used during data collection, or it could stem from differences in the data processing procedures. For example, several institutions independently survey the housing prices in different areas of a city. Due to factors such as the scope of the survey, their data will not be exactly the same and will have certain deviations relative to each other.

If the deviation range is randomly partitioned into two divisions, the probability that the deviation of one cell of suspicious datasets relative to the original dataset falls into one of these divisions is 0.5. If the deviations of multiple cells simultaneously fall into one division, the probability will exponentially decrease. Conversely, the characteristics of the deviation formed by artificially adding perturbation from a particular division can serve as a watermarking scheme. We assume that data owners can access the original dataset to calculate the deviation when detecting the watermark. As we focus on ML applications where a dataset is not frequently modified, it is feasible to access the original dataset.

It is reasonable to assume that the deviation will not be excessively large. We approximate the deviation range using a custom range  $[-p, p]$  and utilize it as the perturbation range. The perturbation range will be partitioned into two divisions, and values will be selected from one division as noise to perturb key cells as a watermark. Figure 2 shows an example of a domain partition. We divide the range  $[-p, p]$  into  $k$  unit domains which are further divided into green domains and red domains in a 1:1 ratio. As illustrated in Figure 2, the range  $[-p, p]$  is partitioned into various red and green domains of different sizes but with an overall equal length. Based on the probability that the deviation of all the detected key cells that fall into the specific type will decrease exponentially, the watermark embedding algorithm only needs to perturb a small number of cells. The characteristics of the deviation distribution can be tested with one proportion z-test, which will be explained in detail in Section 3.3.




Figure 2: An example of domain partition

For brevity, this paper focuses on the widely used tasks of classification and regression within the realm of ML. Tabular datasets employed by these tasks chiefly comprise two types of attributes: numerical attributes and categorical attributes, which can be used by TabularMark to embed watermarks. Besides, we do not consider the character traits of categorical attributes, but only the encoded categorical attributes such as 0-6. Hence, the categorical attributes can be regarded as special numerical attributes.

	C1	C2	C3	C4	...	C496	C497	C498	C499	C500
Private	24.86	-3.45	11.72	37.21	...	-8.14	44.68	18.33	28.07	2.15
Private	24.86	-3.45	11.72	37.21	...	-8.14	44.68	18.33	28.07	2.15
			167890	e3f2a1				bcdelf0		
			-4.2	+3.7				-3.1		
Published	24.86	-7.65	11.72	40.91	...	-8.14	44.68	15.23	28.07	2.15

Figure 3: An example of watermark embedding

A tabular dataset contains attributes  $A_0, A_1, \dots, A_{v-1}$  and we assume ML models using  $A_{v-1}$  as the prediction target. For simplicity, we select key cells from one attribute. In Section 4, we verify the effectiveness of TabularMark by embedding the watermark into the attribute  $A_{v-1}$ . The algorithm can easily extend when the watermarking scheme selects key cells from multiple attributes. Algorithm 1 gives detailed information on the watermark embedding algorithm when we select key cells from numerical attributes. We first choose an attribute  $A_i$  to which the watermark will be embedded (Line 1). Then, the number of key cells  $n_w$  and the boundary of the perturbation range  $p$  are determined based on the strength of the watermark (Line 2). On the original dataset  $D_o$ , we pick  $n_w$  cells as key cells from the chosen attribute  $A_i$  (Line 3). For each key cell, we divide the perturbation range  $[-p, p]$  into  $k$  domains equally, then randomly divided into  $0.5k$  green domains and  $0.5k$  red domains, and the random number seed is kept as a secret. In practice, we tend to choose a relatively large value for  $k$ , such that within a small part of  $[-p, p]$ , the overall lengths of green domains and red domains approach equality. Perturb the key cell by randomly choosing a number from the green domains (Lines 4-9). After perturbing all key cells, we finish the watermark embedding process. The algorithm can be converted to categorical attributes with simple modifications. As the value range of encoded categorical attributes is discrete and relatively small, it does not necessitate a predefined perturbation range. We choose to randomly divide the possible categories into two domains, green and red, and then select a category from the green domain to replace the original category attribute.

*Example 3.1.* Figure 3 shows an example of the watermark embedding algorithm for numerical attributes. Suppose the attribute we embed a watermark has 500 cells. ① We select 50 key cells from these cells, and 3 key cells out of them are indicated by checkmarks. We set the perturbation range to  $[-5, 5]$ . ② For each key cell, we use a random number to seed a random number generator. The random seeds corresponding to each key cell are represented below in hexadecimal format. Utilizing these seeds, we partition the range  $[-5, 5]$  as illustrated in Figure 2. Since each key cell has a unique seed, the resulting green and red domains are also different for

each cell. Then, for each key cell, we select a random number as data noise from its corresponding green domains (represented by the green numbers under the random seeds) to perturb it. ③ After perturbing all key cells in this manner, we obtain the watermarked dataset, which is then ready for publishing.

---

**Algorithm 1:** Watermark Embedding Algorithm.
 

---

**input** : original dataset  $D_o$ , selected attribute  $A_i$   
**output** : watermarked dataset  $D_w$

- 1 choose an attribute  $A_i$  to which the watermark is to be added;
- 2 determine the values of  $p$ ,  $n_w$ , and  $k$ ;
- 3 select  $n_w$  cells as key cells from the selected attribute;
- 4 **foreach** key cell **do**
  - 5 use a random number to seed a random number generator;
  - 6 partition the range  $[-p, p]$  into  $k$  unit domains;
  - 7 with this seed, randomly divide these domains into 0.5k green domains and 0.5k red domains;
  - 8 from the green domains, select a random number;
  - 9 perturb the key cell with this random number;

---

### 3.3 Watermark Detection

In this section, we provide the details of the watermark detection algorithm. Specifically, the watermark detection algorithm employs a rigorous statistical measure to examine suspicious datasets, ensuring successful detection even in the face of substantial attacks. We assume that  $D_o$  can be accessed during the watermark detection process. Next, we introduce this measure in detail.

**3.3.1 One Proportion Z-test.** The one proportion z-test [18] is the statistical tool utilized in the watermark detection algorithm. One proportion z-test is a statistical test that determines whether a single sample rate (e.g., the success rate) is significantly different from a known or hypothesized population rate. We define the null hypothesis  $H_0$  to detect the watermark as follows,

$H_0$  : The tabular dataset is not generated by attacking the watermarked dataset.

Then we will test a statistic for a one proportion z-test as

$$z = \frac{2 \times (n_g - 0.5n_w)}{\sqrt{n_w}},$$

where  $n_w$  is the total number of key cells and  $n_g$  is the number of green cells counted in the detection phase. The meaning of this equation is to subtract the expected value from the observed value, and then divide by the standard deviation of the observed value. In the watermark detection algorithm,  $n_g$  is the main observed value. Given that in the watermark embedding algorithm, the length ratio of green domains to red domains is one-to-one, the probability of the deviations of the detected key cells relative to the key cells in the original dataset falling within the green domains can be considered

	C1	C2	C3	C4	...	C496	C497	C498	C499	C500
Private	24.86	-3.45	11.72	37.21	...	-8.14	44.68	18.33	28.07	2.15
Suspicious	24.86	-6.65	11.72	39.91	...	-8.14	44.68	15.03	28.07	2.15
		-3.20		+2.70				-3.30		
		f67890		e3f2a1				bcdef0		
		-3.20		+2.70				-3.30		

①  
②  
③  
④

**Figure 4:** An example for watermark detection

as 0.5. Therefore, the expectation of  $n_g$  is  $0.5n_w$  and its standard deviation is  $\sqrt{0.25n_w}$ .

A significance level  $\alpha$  needs to be defined according to the strength of the scheme. Commonly,  $\alpha = 0.05$  is used, which implies a 5% risk of incorrectly rejecting the null hypothesis ( $H_0$ ). Subsequently, we need to calculate the  $p$ -value. The  $p$ -value is the probability of observing a test statistic as extreme as, or more extreme than, the value observed, under the assumption that the null hypothesis is true. If the  $p$ -value is less than or equal to  $\alpha$ , we reject the null hypothesis. This means that there is statistical evidence to suggest that the sample proportion is significantly different from the expected proportion. Conversely, if the  $p$ -value is greater, we fail to reject the null hypothesis, suggesting that any observed difference could reasonably occur by random chance, implying that no watermark is detected in the suspicious dataset  $D_s$ . The  $p$ -value can directly measure the magnitude of the false positive rate because a larger  $p$ -value indicates a higher probability of observing the sample under the null hypothesis condition, which means a higher false positive rate when we reject the null hypothesis.

**3.3.2 Detection.** In the watermark detection algorithm,  $\alpha$  will be set directly to the value related to the z-score. As z-score and  $p$ -value are negatively correlated, the larger the z-score, the smaller the  $p$ -value. Users can customize the threshold to adjust the strength of the watermark detection algorithm. If the threshold of the  $p$ -value is set to 0.05, the corresponding z-score threshold is 1.96. In Section 4, we default the z-score threshold to 1.96.

Algorithm 2 details the steps of the watermark detection algorithm for numerical attributes. For each key cell of  $D_s$ , the partition of green domains and red domains is recovered using the corresponding random seed (Lines 3-4). Calculate the difference of this key cell between  $D_s$  and  $D_o$ , and determine whether the difference belongs to the divided green domains. If so, we mark this key cell as a green cell (Lines 5-7). The process concludes with the calculation of the z-score, which is compared to a predetermined threshold  $\alpha$ . If the computed z-score surpasses  $\alpha$ , it indicates the presence of the watermark in the dataset  $D_s$ . In contrast, if the z-score falls below  $\alpha$ , it suggests that the watermark is not detected in  $D_s$  (Lines 8-12). The algorithm can be converted to categorical attributes by directly determining whether a key cell in  $D_s$  belongs to the green domain to count the green cells.

**Example 3.2.** An example of the watermark detection algorithm is shown in Figure 4. ① First, we calculate the difference of all 50 key cells between the private data and the suspicious data. ② Next, we use the random seeds previously employed in watermark embedding to restore the green domains and red domains corresponding to each key cell. ③ If the difference of a key cell in the suspicious

**Algorithm 2: Watermark Detection Algorithm.**


---

**input** : original dataset  $D_o$ , suspicious dataset  $D_s$ , selected attribute  $A_i$ , perturbation range  $[-p, p]$ , number of unit domains  $k$ , exact set of key cells and random seeds

**output** : true or false

- 1  $n_g \leftarrow 0$ ;
- 2 **foreach** key cell **do**
- 3     use the corresponding random number to seed a random number generator;
- 4     restore the green domains and red domains of  $[-p, p]$ ;
- 5     calculate the difference between a key cell of  $D_s$  and that of  $D_o$ ;
- 6     **if** the difference  $\in$  green domains **then**
- 7          $n_g \leftarrow n_g + 1$ ;
- 8      $z \leftarrow \frac{2(n_g - 0.5n_w)}{\sqrt{n_w}}$ ;
- 9     **if**  $z \geq \alpha$  **then**
- 10         **return** true;
- 11 **else**
- 12     **return** false;

---

data falls within its corresponding green domains, the key cell is marked as a green cell. Otherwise, it is marked as a red cell. We find that 30 out of the total 50 key cells are green, this indicates that 40% of the green cells have been flipped by the attacker into red cells.

④ Given that the total number of key cells is 50, the corresponding z-score is calculated to be 1.414. By determining whether the z-score is greater than  $\alpha$ , we can determine whether the suspicious data is watermarked.

**3.3.3 Matching.** In many scenarios, the relative positions of the tuples requiring examination in the detected dataset are inconsistent with those in the watermarked dataset. This inconsistency can be caused by attacks, such as insertion or deletion attacks, which may change the relative positions of tuples. Therefore, it is necessary to locate these perturbed tuples. Considering that watermark schemes reliant on primary keys are susceptible to primary key replacement attacks, we need a matching algorithm to address this issue. Inspired by the basic fact that for a tabular dataset, the probability that the MSBs of multiple attributes are simultaneously equal at the same time is low, and it's difficult for attackers to substitute them all, we choose MSBs of several attributes to form a special primary key and check whether the primary key of a tuple in  $D_s$  is equal to the primary key of a tuple containing a key cell in  $D_o$ .

Algorithm 3 details the steps to locate tuples with key cells in  $D_s$ . We first select MSBs of  $k$  attributes as a special primary key (Line 1). For each tuple containing key cells in  $D_o$ , we use the corresponding primary key to compare with the primary key of each tuple in  $D_d$ . We find the corresponding key cell when the tuple in  $D_d$  is matched (Lines 3-8). Our practice shows that the effect of the watermark detection algorithm is already effective when two attributes are chosen, and the selected attributes are almost equivalent to a true primary key when we choose three attributes.

**Algorithm 3: Matching Algorithm.**


---

**input** : original dataset  $D_o$ , detected dataset  $D_d$

**output** : key cells in  $D_o$

- 1 select MSBs of  $k$  attributes as a primary key;
- 2  $k_{ds} \leftarrow \emptyset$ ;
- 3 **foreach** tuple containing key cell in  $D_o$  **do**
- 4     **foreach** tuple in  $D_d$  **do**
- 5         **if** primary key of  $t_d =$  primary key of  $t_o$  **then**
- 6              $k_d \leftarrow$  key cell in  $t_d$ ;
- 7              $k_{ds} \leftarrow k_{ds} \cup \{k_d\}$ ;
- 8     **return**  $k_{ds}$ ;

---




Figure 5: An example of matching tuples

To combine Algorithm 2 with Algorithm 3, we simply run Algorithm 3 at the beginning of Algorithm 2 to match tuples, which is designed to identify the perturbed tuples containing key cells, especially in scenarios where the order of tuples is disrupted due to some attacks such as insertion or deletion attacks.

*Example 3.3.* Figure 5 shows an example of using Algorithm 3 to find perturbed tuples. The left side of the figure shows the detected dataset and the right side is the original dataset. It can be observed that the order of the detected dataset is scrambled compared to the original dataset, making it difficult to find perturbed tuples directly. We assume that the first tuple of the original dataset is the watermarked tuple containing a key cell. We then select MSBs of the first two attributes,  $A_0$  and  $A_1$ , as a special primary key. Using this special primary key, we compare tuples in the detected dataset. When we compare to the third tuple of the detected dataset, both values of MSBs of  $A_0$  and  $A_1$  equal the values in the special primary key of the tuple. Finally, we successfully identify the perturbed tuple in the detected dataset.

### 3.4 Analysis on Watermark Removal

In this section, to demonstrate the robustness of TabularMark, we model the potential strategies of alteration attacks and mathematically prove the lower bound of the number of cells that need to be altered by the attacker to achieve a high confidence level. According to our partition strategy, when a green cell is perturbed out of the range of the unit green domain, there is a 0.5 chance that it will be flipped to a red cell. According to the threat model, the attackers cannot discern the key cells. Therefore, to flip green key cells into red, attackers can only randomly select  $n_h$  cells and introduce noise. We simulate alteration attacks using three representative methods of adding noise, including uniform distribution noise, Gaussian distribution noise, and Laplace distribution noise, all with a mean of 0 and a standard deviation of  $\sigma$ .

**PROPOSITION 3.4.** Denote by  $n_h$  the amount of cells an attacker tamper with. When each cell is added with an i.i.d noise from uniform distribution  $\epsilon \sim U[-6\sigma, 6\sigma]$ . The expectation of  $n_h$  for the watermark removing is

$$E[n_h] = \frac{n(n_w - n_\alpha)}{n_w p_\sigma},$$

where  $n_\alpha = \alpha\sqrt{\frac{n_w}{4}} + \frac{n_w}{2}$  is the least amount of green cells to make the z-score achieve  $\alpha$  and  $p_\sigma = 0.5 - \frac{p}{12k\sigma}$ . For example, when  $n = 10000$ ,  $n_w = 400$  and  $n_\alpha = 250$ , we have  $E[n_h] > 7500$  since  $p_\sigma < 0.5$ . Moreover, to achieve a 95% confidence interval, we have

$$n_h > n_w + \left(\frac{(n - n_w)!}{0.05n!}\right) \sum_{n_x=0}^{n_w - n_\alpha} \sum_{n_y=n_x}^{n_w} \binom{n_y}{n_x} \binom{n_w}{n_y} p_\sigma^{n_x} (1 - p_\sigma)^{n_y - n_x} \frac{-1}{n_w}.$$

**PROOF.** We state the case where  $\gamma = 0.5$  and  $3\sigma \geq p/k$ . Consider a green cell  $x \sim U[0, 2p/k]$ , the probability  $x + \epsilon$  is not in the same interval as  $x$  is  $Pr(x + \epsilon > 2p/k) + Pr(x + \epsilon < 0)$ , which is the  $Pr(\epsilon > 2p/k - x) + Pr(\epsilon < -x) = 1 - \frac{p}{6k\sigma}$ . Therefore, the probability that  $x + \epsilon$  is red is  $p_\sigma = 0.5 - \frac{p}{12k\sigma}$ .

To achieve the watermark removal, we need to add noise on at least  $n_w - n_\alpha$  cells with the watermark. Therefore, the expectation of  $n_h$  is  $\frac{n(n_w - n_\alpha)}{n_w} (0.5 - \frac{p}{12k\sigma})^{-1}$ .

Consider the case of a 95% confidence interval. Denote by  $N_x$  the amount of key cells changing from green to red. Then,  $Pr(N_x = n_x) = \sum_{n_y=n_x}^{n_w} \frac{\binom{n_y}{n_x} \binom{n_w}{n_y} \binom{n - n_w}{n_h - n_y}}{\binom{n}{n_h}} p_\sigma^{n_x} (1 - p_\sigma)^{n_y - n_x} > \frac{\binom{n - n_w}{n_h - n_w}}{\binom{n}{n_h}} \sum_{n_y=n_x}^{n_w} \binom{n_y}{n_x} \binom{n_w}{n_y} p_\sigma^{n_x} (1 - p_\sigma)^{n_y - n_x}$  when we consider  $n_h$  is not very large, which means  $\sum_{n_x=0}^{n_w - n_\alpha} Pr(N_x = n_x) > (n_h - n_w) n_w \frac{(n - n_w)!}{n!} \sum_{n_x=0}^{n_w - n_\alpha} \sum_{n_y=n_x}^{n_w} \binom{n_y}{n_x} \binom{n_w}{n_y} p_\sigma^{n_x} (1 - p_\sigma)^{n_y - n_x}$ . Then, we can derive the aforementioned inequality.  $\square$

**PROPOSITION 3.5.** Denote by  $n_h$  the amount of cells an attacker tamper with. When each cell is added an i.i.d noise from Gaussian distribution  $\epsilon \sim \mathcal{N}(0, \sigma^2)$ . The expectation of  $n_h$  for the watermark removing is

$$E[n_h] \geq \frac{n(n_w - n_\alpha)}{n_w p_\sigma},$$

where  $n_\alpha = \alpha\sqrt{\frac{n_w}{4}} + \frac{n_w}{2}$  is the least amount of green cells to make the z-score achieve  $\alpha$ ,  $p_\sigma = \frac{1}{4} + \frac{1}{2} \int_{\frac{2p}{k\sigma}}^{\infty} \psi(x) dx$ , and  $\psi(\cdot)$  denotes probability density function (PDF) of the standard normal distribution. Moreover, to achieve a 95% confidence interval, we have

$$n_h > n_w + \left(\frac{(n - n_w)!}{0.05n!}\right) \sum_{n_x=0}^{n_w - n_\alpha} \sum_{n_y=n_x}^{n_w} \binom{n_y}{n_x} \binom{n_w}{n_y} p_\sigma^{n_x} (1 - p_\sigma)^{n_y - n_x} \frac{-1}{n_w}.$$

**PROOF.** Similarly, we only need to compute the maximum probability that a green cell is changed to red and substitute it into  $p_\sigma$  of Proposition 3.4, with  $Pr(\epsilon > 2p/k - x) + Pr(\epsilon < -x) \leq Pr(\epsilon > 2p/k) + Pr(\epsilon < 0) = \frac{1}{2} + \int_{\frac{2p}{k\sigma}}^{\infty} \psi(x) dx$ .  $\square$

**PROPOSITION 3.6.** Denote by  $n_h$  the amount of cells an attacker tamper with. When each cell is added an i.i.d noise from Laplace distribution  $\epsilon \sim \text{Lap}(0, \sigma/\sqrt{2})$ . The expectation of  $n_h$  for the watermark

removing is

$$E[n_h] \geq \frac{n(n_w - n_\alpha)}{n_w p_\sigma},$$

where  $n_\alpha = \alpha\sqrt{\frac{n_w}{4}} + \frac{n_w}{2}$  is the least amount of green cells to make the z-score achieve  $\alpha$  and  $p_\sigma = \frac{1}{4} (1 + \exp\{-\sqrt{2}p/k\sigma\})$ . Moreover, to achieve a 95% confidence interval, we have

$$n_h > n_w + \left(\frac{(n - n_w)!}{0.05n!}\right) \sum_{n_x=0}^{n_w - n_\alpha} \sum_{n_y=n_x}^{n_w} \binom{n_y}{n_x} \binom{n_w}{n_y} p_\sigma^{n_x} (1 - p_\sigma)^{n_y - n_x} \frac{-1}{n_w}.$$

**PROOF.** Similarly, we only need to compute the maximum probability that a green cell is changed to red, with  $Pr(\epsilon > 2p/k - x) + Pr(\epsilon < -x) \leq Pr(\epsilon > 2p/k) + Pr(\epsilon < 0) = \frac{1}{2} + \int_{\frac{\sqrt{2}p}{k\sigma}}^{\infty} \frac{1}{2} e^{-x} dx = \frac{1}{2} (1 + \exp\{-\sqrt{2}p/k\sigma\})$ .  $\square$

**Analysis and Conclusions.**  $n_h$  required to erase the watermark with a 95% confidence is obviously greater than the expected  $n_h$  when erasing the watermark, so we analyze the robustness of the watermark from  $E[n_h]$ . The upper bound of  $p_\sigma$  is strictly less than 0.5, since there is still a fifty percent chance that the changed interval remains green. Moreover, the upper bound of  $p_\sigma$  is greater in the case of Gaussian noise or Laplace noise when the variance of the noise is small, which reflects the potential of the latter two strategies to reduce  $E[n_h]$  in such cases.

For  $E[n_h]$ , even if we assume that every perturbation can flip green cells, we have  $E[n_h] \geq \frac{n(n_w - n_\alpha)}{n_w p_\sigma} = \frac{n(n_w - n_\alpha)}{n_w} = n \left(1 - \frac{n_\alpha}{n_w}\right)$  when added noises are from uniform distribution, Gaussian distribution noise, or Laplace distribution. Considering that  $n_\alpha$  approaches half of  $n_w$ ,  $E[n_h]$  is also close to half of  $n$ , much greater than  $n_w$ , indicating that attackers need to pay a much higher cost than the data owner to erase the watermark.

## 4 EXPERIMENTS AND ANALYSIS

In this section, we organize experiments by answering five research questions (RQs). Following the three desiderata mentioned in Section 1, we design four related RQs to demonstrate the superiority of TabularMark. To further explore the usage of TabularMark, we complement one additional RQ on hyper-parameter selection. The code is available at <https://github.com/yihzheng258/TabularMark>.

- RQ1)** Is the TabularMark detectable for various tabular datasets? (Section 4.1)
- RQ2)** Can TabularMark achieve non-intrusiveness across various tabular datasets used for ML? (Section 4.2)
- RQ3)** Can TabularMark be robust against common malicious attacks? (Section 4.3)
- RQ4)** How does TabularMark perform compared to the related state-of-the-art schemes? (Section 4.4)
- RQ5)** What are the trade-offs among the hyper-parameters of TabularMark? (Section 4.5)

**Experimental Setup.** Experiments are conducted on a server comprising two Intel(R) Xeon(R) Platinum 8383C CPU@2.70GHz, running Ubuntu 18.04 LTS 64-bit with 256GB memory. We use NumPy to randomly generate two-dimensional normal distribution data as synthetic datasets, with a mean of  $\mu = 0$ , a variance of  $\sigma = 20$ , and a

total of 2000 tuples to run. Forest Cover Type [8], HOG feature based on the digits dataset [4], and Boston Housing Prices [22] are employed as real-world datasets. Notably, the HOG feature dataset is generated with the histogram of oriented gradients (HOG) features extracted from the digits dataset, combined with their categories. For ML models, XGBoost [13], Random Forest [10], and Linear Regression [33] are used to cover both classification and regression tasks. Besides, the z-score threshold defaults to 1.96 to maintain the theoretical constraint of a 5% false positive rate. To keep the content coherent and focused, some less commonly used dataset settings will be explained when introducing specific tasks.

**Reproducibility.** For reproducibility, we set a specific random seed to ensure that the data distortion and model training effects are repeatable on the same hardware and software stack. To eliminate the randomness of the results, the experimental outcomes are generally averaged over multiple runs.

#### 4.1 Detectability

In this section, we investigate the z-scores on the original, watermarked, perturbed datasets to answer **RQ1**. One synthetic and three real-world datasets are utilized to verify the detectability of TabularMark.

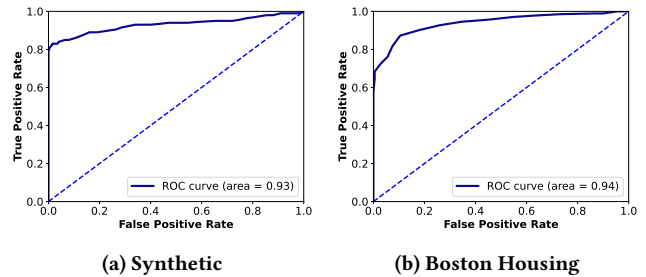
In the experiment, synthetic datasets are watermarked with  $n_w = 300$ ,  $p = 2\sigma$ , and  $k = 500$  on the first dimension. Forest Cover Type dataset is watermarked with  $n_w = 400$  on the attribute *Cover\_Type*. HOG feature dataset is watermarked with  $n_w = 150$  on the attribute *category*. Boston Housing Prices dataset is watermarked with  $p = 25$ ,  $k = 500$ , and  $n_w = 50$  on the attribute *MEDV*. Subsequently, watermark detection is performed on the original dataset  $D_o$ , the watermarked dataset  $D_w$ , and the dataset  $D_p$  generated after randomly perturbing  $D_o$ . Results of detected z-scores are shown in Table 2.

**Table 2: Z-scores**

Dataset	$D_o$	$D_w$	$D_p$
Synthetic	-0.0531	17.3	-0.0393
Forest	-0.0159	18.6	-0.0467
HOG	-0.327	12.3	0.128
Boston Housing	-0.113	6.91	-0.142

In practice, there are many potential suspicious datasets that may cause false positive problems. To evaluate the false positive problem of TabularMark, we construct neighboring suspicious datasets for synthetic datasets and the real Boston Housing Price dataset and then conduct watermark detection. Specifically, for the synthetic datasets, 100 independently generated datasets, 100 perturbed original datasets, and 200 perturbed watermarked datasets are used to construct a set of 400 neighboring datasets. For the Boston Housing Price dataset, we randomly divide the original dataset into two equal parts, A and B. We use 100 perturbed Bs, 100 perturbed As, and 200 perturbed watermarked As to construct a set of 400 neighboring datasets. The detection results of TabularMark are expressed using ROC curves in Figure 6.

The ROC curves and the AUC results (both greater than 0.9) show that TabularMark is highly effective in detecting the watermark with a low false positive rate.



**Figure 6: ROC curves**

#### 4.2 Non-Intrusiveness

In this section, we investigate ML-related metrics on the original and watermarked datasets to answer **RQ2**. Three real-world datasets, including the Forest Cover Type dataset, the HOG feature dataset, and the Boston Housing Prices dataset, are utilized to verify the non-intrusiveness of TabularMark.

We train XGBoost classification models separately on  $D_o$  and  $D_w$ . Then we evaluate their classification performance. The  $F_1$ -scores for categories 2, 4, and 6 are presented in Table 3. It can be observed that the ML utility of  $D_w$ , compared to  $D_o$ , does not significantly decrease, which implies that TabularMark hardly impairs the ML utility of the watermarked dataset. This indicates that TabularMark is non-intrusive.

**Table 3:  $F_1$ -scores**

Dataset	Category 2	Category 4	Category 6
$D_o$	0.888	0.940	0.848
$D_w$	0.887	0.937	0.845

**More Datasets.** We continue to verify the non-intrusiveness of TabularMark on the HOG feature dataset and Boston Housing Prices dataset. Random Forest and linear regression models are employed to assess their ML utility. Experimental results are demonstrated in Table 4.

**Table 4: ML utility on more datasets**

Dataset	(a) Accuracy on HOG		(b) MSEs on Boston Housing	
	$D_o$	$D_w$	Dataset	$D_o$ $D_w$
Accuracy	0.942	0.940	MSE	24.8 25.6

Additionally, we test the non-intrusiveness of TabularMark on a small dataset Zoo [19] (with 100 tuples), which is watermarked with  $n_w = 15$  on attribute *type*. Random Forest is used to measure the ML utility of datasets. Experimental results are demonstrated in Table 5.

**Table 5: Z-scores and accuracy on Zoo**

Dataset	$D_o$	$D_w$
Z-score	0.366	3.87
Accuracy	0.900	0.870



The experimental results show that the ML-related metrics, including MSE and accuracy, are slightly reduced after watermarking, demonstrating the non-intrusiveness of TabularMark. For small datasets, TabularMark is still non-intrusive. However, since the watermark strength is proportional to the absolute number of key cells, the distortion of the watermark dataset is more significant as the proportion of tuples perturbed is relatively larger in small datasets.

### 4.3 Robustness

In this section, we investigate the z-scores and ML-related metrics on the attacked datasets to answer **RQ3**. Following the proposed schemes [2, 23, 31], we mainly verify the robustness of TabularMark on the real-world dataset Forest Cover Type.

**Alteration Attack.** An alteration attack on a dataset refers to a type of attack where an attacker intentionally modifies the attributes of the dataset in a random or seemingly random manner. The randomness of the changes may make it difficult to detect watermarks. In this experiment, we assume that the attacker is aware of the attribute to which the watermark has been embedded (in actual scenarios, the attribute used for watermarking is generally kept secret). We simulate the attacker randomly perturbing 20%, 40%, 60%, 80% and 100% of the attribute Cover Type and examine the z-scores under various alteration attacks along with the ML utility impact on the XGBoost model. Table 6 and Table 7 respectively show the results of the z-scores and the classification performance for the dataset under alteration attacks.

**Table 6: Z-scores on Forest Cover Type Dataset**

Alteration Attack (%)	20	40	60	80	100
Z-score	13.8	8.67	3.68	-1.11	-6.05

**Table 7:  $F_1$ -scores on Forest Cover Type Dataset**

Alteration Attack (%)	Category 2	Category 4	Category 6
20	0.879	0.853	0.813
40	0.867	0.791	0.761
60	0.847	0.661	0.645
80	0.679	0.0631	0.115
100	0.00203	0.00180	0.000773

We can observe that, given  $\alpha = 1.96$ , when the proportion of the alteration attack reaches 80%, the watermark detection algorithm can no longer detect the watermark in the attacked dataset. However, at this point, the  $F_1$ -scores for these three categories decrease significantly, indicating that the ML utility of the attacked dataset for the model has been greatly reduced. Although the attacker successfully erases the watermark, the trained model becomes inoperable, which implies that TabularMark is robust under alteration attacks.

**Insertion Attack.** An insertion attack on a dataset refers to a type of attack where tuples are created randomly and inserted into the dataset to destroy the watermark. In this experiment, we simulate the attacker randomly inserting 20%, 40%, 60%, 80%, 100% tuples into

the watermarked dataset and examine the z-scores under various insertion attacks. Due to the insertion attack disrupting the order of tuples in the dataset, we employ the first two attributes and the first three attributes of the Forest Cover Type as primary keys respectively to match tuples. Table 8 and Table 9 display the z-score detection results for these two matching methods, respectively.

**Table 8: Z-scores (two-attribute matching)**

Insertion Attack (%)	20	40	60	80	100
Z-score	8.23	7.43	6.87	6.24	5.73

**Table 9: Z-scores (three-attribute matching)**

Insertion Attack (%)	20	40	60	80	100
Z-score	18.3	18.1	18.1	17.9	17.8

The results from the two tables indicate that, under the assumption of a threshold  $\alpha = 1.96$ , watermarks can still be effectively detected even when up to 100% of tuples are inserted relative to the size of the original dataset, which signifies that TabularMark is robust under insertion attacks.

**Deletion Attack.** A deletion attack on a dataset refers to a type of attack where tuples are deleted randomly to destroy the watermark. In this experiment, we simulate the attacker randomly deleting 20%, 40%, 60%, 80%, 100% tuples from the watermarked dataset and examine the z-scores under various deletion attacks. Due to the same reason with the insertion attack, we also employ the first two attributes and the first three attributes of the Forest Cover Type as primary keys respectively to match tuples. Tables 10 and 11 display the z-score detection results for these two matching methods, respectively.

**Table 10: Z-scores (two-attribute matching)**

Deletion Attack (%)	20	40	60	80	100
Z-score	8.42	7.45	6.13	4.59	\

**Table 11: Z-scores (three-attribute matching)**

Deletion Attack (%)	20	40	60	80	100
Z-score	16.7	14.4	11.6	8.21	\

The results from the two aforementioned tables indicate that, under the assumption of a threshold  $\alpha = 1.96$ , the watermark can still be effectively detected even in the event of a substantial deletion attack, which suggests that TabularMark is robust under deletion attacks.

**Data Cleaning Attacks.** Considering that relational data often contains functional dependencies and integrity constraints, attackers may use some data cleaning algorithms such as HoloClean [41] to reverse perturbed key cells back. However, considering attackers have less background knowledge than the data owner, the correlation among attributes that can be exploited is relatively limited.

Besides, it is unlikely that there are explicit functional dependencies involving the target attribute. Otherwise, there would be no need to predict the target attribute through a trained ML model. Therefore, if the data owner chooses to perturb the target attribute, it will unlikely be affected by the data cleaning algorithms. We utilize the example datasets (Adult and Hospital) in the HoloClean repo<sup>1</sup> for experiments. Adult is watermarked with  $n_w = 300$  on the attribute *Income*. Hospital is watermarked with  $n_w = 300$  on the attribute *Condition*. Constraints among these datasets are defined by [41]. Table 12 shows the results of the z-scores for the watermarked dataset  $D_w$  and the attacked dataset  $D_a$  generated by adopting the HoloClean for data cleaning attacks.

**Table 12: Z-scores of Adult and Hospital under data cleaning**

Dataset	$D_w$	$D_a$
Adult	17.3	17.3
Hospital	17.3	10.7

The experimental results demonstrate the robustness of TabularMark against data cleaning attacks.

**Data Anonymization and Synthesization Attacks.** Attackers may also consider stronger and utility-preserving data anonymization and synthesization attacks, such as  $k$ -anonymization [7, 11], synthetic data generation [9, 37, 52], and differential privacy based anonymizations [1, 5, 51] to erase watermarks. We examine the robustness of TabularMark against these attacks using dataset Iris which is watermarked with  $n_w = 20$  on the attribute *class*. Iris datasets are used to train a Random Forest classifier on features of flowers to determine the class of flower from among 3 classes, and the accuracy of the trained model is the metric to measure ML utility.

For  $k$ -anonymization [7, 11], we employ the mean to represent each group for anonymization and set  $k = 3$ . We use different numbers of columns for anonymization (['petal\_width'], ['petal\_length', 'petal\_width'], and ['sepal\_width', 'petal\_length', 'petal\_width']) and use the remaining columns to perform Algorithm 3 for watermark detection. Tables 13 and 14 show that although  $k$ -anonymization can impact the watermark, it also causes a significant decrease in data utility.

**Table 13: Accuracy of  $D_o$  and  $D_w$  on Iris**

Dataset	$D_o$	$D_w$
Accuracy	1.00	0.933

**Table 14: Results of  $k$ -anonymization on different #columns**

#columns	1	2	3
Z-score	5.85	5.38	1.10
Accuracy	0.911	0.873	0.786

For synthetic data generation, many works [9, 37, 52] have been proposed to generate synthetic tabular data following the distribution of original data. With this approach, attackers may completely

<sup>1</sup><https://github.com/HoloClean/holoclean>

erase the embedded watermark since there is no correspondence between the synthetic records and the original records. We adopt [9] as an attack method and keep other settings the same as in the previous experiments.  $D_o$  and  $D_w$  are used to train generative models and generate  $D'_o$  and  $D'_w$ , respectively. Table 15 shows that the data distortion caused by watermarks significantly reduces the utility of the downstream synthetic data  $D'_w$  compared to the synthetic data  $D'_o$ . This suggests that while the watermark may be erased by synthetic data, the significant utility drop still helps protect against the use of the data.

**Table 15: Accuracy of generation on Iris**

Dataset	$D_o$	$D'_o$	$D_w$	$D'_w$
Accuracy	1.00	1.00	0.933	0.778

Attackers can also use differential privacy mechanisms [1, 5, 51] to generate a perturbed statistical summary or synthetic data to prevent the correspondence to  $D_o$ . However, these mechanisms require adding noise in the generative model to ensure differential privacy [47], which further reduces the data utility. Therefore, attackers pay an even greater cost than synthetic data generation attacks to erase the watermark.

In summary, the above experimental analysis and results show that TabularMark can still safeguard the data against these stronger attacks by significantly degrading the utility of the downstream use of the attacked data.

**More Datasets.** We continue to validate the robustness of TabularMark on the HOG feature dataset and Boston Housing Prices dataset. From Section 4.3, due to TabularMark being difficult to remove by other types of attacks, we only simulate alteration attacks. Experimental results are demonstrated in Tables 16 and 17. Besides, we also test the robustness of TabularMark on the small dataset Zoo [19] under alteration attacks. Experimental results are shown in Table 18.

**Table 16: Z-scores and accuracy on HOG**

Alteration Attack (%)	20	40	60	80	100
Z-score	9.71	7.13	4.78	2.55	-0.0816
Accuracy	0.933	0.910	0.825	0.530	0.100

**Table 17: Z-scores and MSEs on Boston Housing**

Alteration Attack (%)	20	40	60	80	100
Z-score	5.30	3.76	2.17	0.571	-1.10
MSE	26.1	27.6	34.0	35.4	42.3

**Table 18: Z-scores and accuracy on Zoo**

Alteration Attack (%)	20	40	60	80	100
Z-score	3.35	2.32	1.29	-0.774	-0.774
Accuracy	0.740	0.550	0.550	0.290	0.190

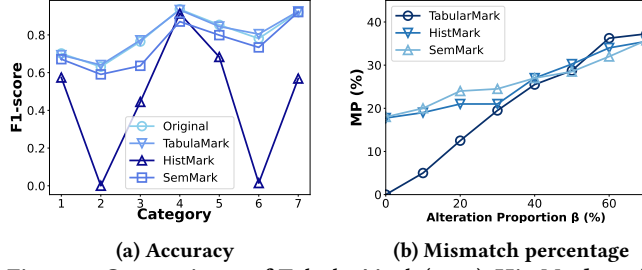


Figure 7: Comparisons of TabularMark (ours), HistMark, and SemMark on (a) non-intrusiveness and (b) robustness.

The results show that attackers would incur substantial costs to make the watermark undetectable, the cost is so high that it renders the trained model nearly unusable. This means that TabularMark can effectively safeguard the ownership of these datasets regardless of their size.

#### 4.4 Comparison

In this section, we compare related watermarking schemes with TabularMark to answer **RQ4**. The real-world dataset Forest Cover Type is utilized to verify TabularMark outperforms related work.

Due to the inconsistency of the detection metrics of TabularMark with related relational database watermarking works, direct comparison with related work is not feasible. Most other works employ the bit error rate (BER) to measure the integrity of the extracted watermark information. In contrast, we use the z-score of the one proportion z-test as our detection metric. As we do not embed watermark plaintext, we use the proportion of key cells flipped from green domains to red domains as a metric similar to the BER. These metrics are uniformly represented as Mismatch Percentage (MP). Although this proportion does not directly indicate the strength of the watermark, it is the closest method available for comparison with metrics used in related work.

We consider two recent relational database watermarking schemes for comparison with TabularMark: 1) HistMark: Hu et al. [23] utilize histogram shifting to embed watermark information. 2) SemMark: Li et al. [30] alter the random least significant position of the decimal numerical attributes to embed watermark information. Both methods are considered state-of-the-art reversible watermarking schemes and are capable of effectively controlling distortion. We will implement TabularMark, HistMark, and SemMark on the Forest Cover Type dataset and compare them based on the three goals of a watermarking scheme. We use a subset of the Forest Cover Type dataset to avoid the adverse impact of having a much larger number of instances in certain categories on watermarking schemes. For the experiments, we randomly select 2000 tuples from each category. The number of key cells for TabularMark and the length of the watermark bit sequence is set to 400, ensuring that the quantity of 0s and 1s in the watermark bit sequence is equal. The watermark is embedded in the attribute Cover\_Type. The classification model XGBoost is utilized to test the ML utility. Figure 7a and Figure 7b exhibit the results of the comparative experiments.

Compared to the original dataset, the watermarked dataset of TabularMark exhibits a negligible decline in the  $F_1$ -scores across

various categories, with the lines almost overlapping. In contrast, the line for SemMark presents a noticeable gap relative to TabularMark. This discrepancy is due to the requirement of SemMark to ensure successful watermark embedding and the use of a majority voting mechanism for robustness, which introduces a higher degree of data distortion. The outcomes for HistMark are even less satisfactory, which is due to histogram shifting that alters a large amount of data in every group. Even changes as minimal as  $\pm 1$  can introduce excessive data distortion for categorical attributes.

When the attack proportion is set to 0, the MP for all three watermarking schemes tends to be low, indicating that detectability is generally assured. However, the mechanisms employed by SemMark and HistMark may result in some watermark bits not being successfully embedded into the dataset. For HistMark, if the peak bin of the histogram of each group matches the maximum or minimum value of the corresponding attribute, it may not be possible to shift it, leading to the failure of watermark embedding for that group. In the case of SemMark, the location where a watermark bit is embedded within a tuple is randomly selected, which may cause some positions to be unselected. TabularMark, on the other hand, does not rely on the embedding of specific watermark bits and is not susceptible to this issue.

As shown in Figure 7b, when the alteration attack proportion increases, aside from an initially slightly lower MP for TabularMark, the plots of all three watermarking schemes converge beyond a 30% proportion, suggesting that there is no significant disparity in the MP metric among the schemes. Although TabularMark employs a threshold-based z-score for watermark detection rather than MP, this implies that the robustness of TabularMark is comparable to that of state-of-the-art methods to a certain extent.

#### 4.5 Design Trade-Offs

In this section, we conduct a qualitative analysis of the crucial parameters of TabularMark and analyze their trade-offs to answer **RQ5**. Synthetic datasets are utilized to explore trade-offs to exclude external factors' interference.


To generate labels for classification experiments, data across two dimensions is combined with weights in the range of  $[-1, 1]$ , and the weighted sum is passed through a logistic function to compute the probability of an event occurring. This probability is then used to conduct a binomial distribution experiment to generate a target array. The numerical attribute Dimension 1 is utilized to investigate the trade-offs of the parameters. A logistic regression model is trained to evaluate data utility. We assume that the attacker employs an alteration attack, aiming to modify the dataset subtly in an attempt to remove the watermark without making overt changes that would degrade the ML utility of the dataset. We employ a logistic regression model as the classifier for experiments.

**4.5.1 The Perturbation Range  $[-p, p]$ .** The parameter  $p$  is related not only to the strength of the watermark but also to the ML utility of the dataset. Here, we treat  $p$  as a variable, with values set at  $0.5\sigma$ ,  $1\sigma$ ,  $1.5\sigma$ ,  $2\sigma$ , and  $2.5\sigma$ . Other parameters are set as  $k = 500$  and  $n_w = 300$ . We fix the alteration range at  $2\sigma$ , meaning that they perturb the dataset within a range of  $[-2\sigma, 2\sigma]$ . The attack proportions are set at 20%, 40%, 60%, 80%, and 100%.

In the following, we denote these proportions using the symbol  $\beta$ . Table 19 illustrates the impact of watermarking on the ML utility of the dataset, measured by the classification accuracy of the trained logistic regression model. Figure 8a and Figure 8b respectively illustrate the variations in z-scores and classification accuracy of the dataset under attacks as the parameter  $p$  changes. Table 20 presents the specific values of z-scores and classification accuracy for the attacked datasets.

**Table 19: Accuracy**

Dataset	$0.5\sigma$	$1.0\sigma$	$1.5\sigma$	$2.0\sigma$	$2.5\sigma$
$D_o$	0.977	0.977	0.977	0.977	0.977
$D_w$	0.976	0.973	0.965	0.956	0.952



**Figure 8: Accuracy and z-scores on  $\beta$ s**

With the increase in  $p$ , the ML utility of the watermarked dataset degrades. In alteration attacks, it can be observed that as  $\beta$  increases, the z-score decreases rapidly. Meanwhile, the decrease in the utility of attacked datasets is more pronounced. With the increase in  $p$ , the ML utility of the attacked dataset degrades. This is because as  $p$  increases, the perturbed data are more likely to fall within the green domains, forcing attackers to increase the rate of alteration, resulting in making the watermark unusable. The larger  $\beta$ , the poorer the performance of the trained model. It is important to note that  $p$  cannot be increased indefinitely, as key cells may be identified as outliers.

**4.5.2 The Number of Key Cells  $n_w$ .** The number of key cells, denoted by  $n_w$ , is also a crucial parameter. Here, we treat  $n_w$  as a variable, with values set at 100, 200, 300, 400, and 500. Other parameters are set as  $p = 2.0\sigma$  and  $k = 500$ . We also fix the alteration range of perturbation at  $2\sigma$ . The attack proportions are set at 20%, 40%, 60%, 80%, and 100%. Table 21 illustrates the impact of watermarking on the ML utility of the dataset, measured by the classification accuracy of the trained logistic regression model. Figure 9a and Figure 9b respectively illustrate the variations in z-scores and classification accuracy of the dataset under attack as the parameter  $n_w$  changes. Table 22 presents the specific values of z-scores and classification accuracy for the attacked datasets.

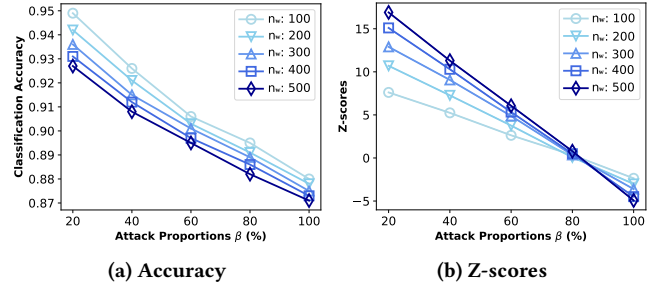
It can be observed that as  $n_w$  increases, the ML utility of the watermarked dataset declines due to the increase of perturbed data, and there is a slight decrease in the classification performance of the trained model. Simultaneously, the cost of attack rises, as a larger  $n_w$  results in a higher z-score under the same attack intensity, making the watermark less susceptible to being erased.

**Table 20: Z-scores and accuracy on  $p$ s**

$p$	Z-score					Accuracy				
	20	40	60	80	100	20	40	60	80	100
$0.5\sigma$	11.4	5.23	-0.903	-6.97	-13.0	0.953	0.929	0.908	0.898	0.883
$1.0\sigma$	12.2	7.01	1.84	-3.70	-8.68	0.950	0.927	0.906	0.895	0.881
$1.5\sigma$	12.8	8.23	3.41	-1.19	-5.68	0.945	0.922	0.904	0.892	0.879
$2.0\sigma$	13.1	9.06	4.91	0.457	-3.69	0.936	0.916	0.901	0.888	0.876
$2.5\sigma$	13.3	9.39	5.24	1.37	-2.43	0.928	0.908	0.896	0.882	0.872

**Table 21: Accuracy**

Dataset	100	200	300	400	500
$D_o$	0.977	0.977	0.977	0.977	0.977
$D_w$	0.972	0.963	0.957	0.953	0.949




**Figure 9: Accuracy and z-scores on  $\beta$ s**

**Table 22: Z-scores and accuracy on  $n_w$ s**

$n_w$	Z-score					Accuracy				
	20	40	60	80	100	20	40	60	80	100
100	7.61	5.24	2.63	0.344	-2.36	0.949	0.926	0.906	0.895	0.880
200	10.7	7.27	3.76	0.0735	-3.01	0.942	0.921	0.903	0.891	0.878
300	12.9	9.06	4.88	0.351	-3.61	0.936	0.915	0.901	0.889	0.875
400	15.1	10.3	5.33	0.496	-4.49	0.931	0.912	0.897	0.886	0.873
500	16.9	11.3	6.06	0.785	-4.93	0.927	0.908	0.895	0.882	0.871

**4.5.3 The Proportion of Green Domains  $\gamma$ .** This represents a unique trade-off, as in the watermark embedding and detection algorithms, the ratio of the length of green domains to the total length of the perturbation range, represented as  $\gamma$ , is set at 0.50. This decision is based on the assumption that the deviation of cells in suspicious datasets from those in the original dataset falls within a certain range randomly. However, the impact of variations in  $\gamma$  on TabularMark is worth exploring. Here, we treat  $\gamma$  as a variable, with values set at 0.25, 0.33, 0.5, 0.67, 0.75,  $n_w = 300$ . We also fix  $p = 1.5\sigma$  and  $k = 500$ . We fix the alteration range at  $1.5\sigma$ . Table 23 illustrates the impact of watermarking on the ML utility of the dataset, measured by the classification accuracy of the trained logistic regression model. Figures 10a and 10b respectively illustrate the variations in z-scores and classification accuracy of the dataset under attack as the parameter  $\gamma$  changes. It can be observed that, as  $\gamma$  increases, the ML utility of the watermarked dataset remains nearly unchanged, indicating that variations in  $\gamma$  hardly alter the degree of data distortion. When attackers increase the attack proportion, the classification accuracy for different  $\gamma$  values tends to converge, while smaller  $\gamma$  values result in larger z-scores. This is



(a) Accuracy (b) Z-scores  
Figure 10: Accuracy and z-scores on  $\beta$ s

dictated by the calculation formula of the z-score, indicating that smaller  $\gamma$  values can enhance the robustness of TabularMark.

Table 23: Accuracy

Dataset	0.25	0.33	0.50	0.67	0.75
$D_o$	0.958	0.958	0.958	0.958	0.958
$D_w$	0.958	0.958	0.956	0.957	0.957

To further investigate the trade-off afforded by  $\gamma$ , we tested the z-scores of different watermarking schemes on the original dataset. The average results are presented in Table 24. It can be observed that when  $\gamma$  is either very small or very large, the z-scores tend to be larger. This suggests that when the threshold is low, there is a potential risk of false positives, where an unwatermarked original dataset might be misidentified as watermarked. This occurs because if the green cells of the original dataset deviate slightly from the expected values when  $\gamma$  is small, it can result in a disproportionately large z-score. Conversely, when  $\gamma$  is large, the key cells of the original dataset naturally fall within the green domains with a higher probability, leading to larger z-scores as well. Therefore, selecting  $\gamma = 0.50$  may represent a more balanced choice.

**4.5.4 Conclusion.** Increasing the size of  $p$  and  $n_w$  can enhance the robustness of the watermark, but it may compromise non-intrusiveness. We suggest that data owners adjust these two parameters based on their actual needs. Additionally, keeping  $\gamma$  at 0.50 helps to reduce the potential risk of false positives so we suggest data owners use it.

Table 24: Z-scores

Dataset	0.25	0.33	0.50	0.67	0.75
$D_o$	1.47	0.489	0.346	0.122	0.933

## 5 DISCUSSION

In this section, we discuss some optimization strategies to improve TabularMark and extend its applicability.

### 5.1 Noise Selection

In this section, we propose a new perturbation noise selection method to enhance the non-intrusiveness of TabularMark.

In Line 6 of Algorithm 1, we randomly select a number from the green domains for perturbation. However, this uniform random

selection strategy may pick up a fair amount of substantial noise to perturb the original dataset, which could potentially have a noticeable impact on the ML utility of the data. To address this issue, we improve the original strategy by adopting the normal distribution probability. This means that numbers closer to the original value have a higher probability of being selected, which is intuitively reasonable as it implies that larger deviations are less likely to occur. We characterize this transition with the corresponding PDF. In the watermark embedding phase, we continuously sample from the probability density function on  $[-p, p]$  until the sample falls within the green domains.

The PDF for uniform random selection is given by:

$$f(x) = \begin{cases} \frac{1}{2p}, & \text{if } -p \leq x \leq p, \\ 0, & \text{otherwise.} \end{cases}$$

The PDF for selection based on a normal distribution is given as:

$$f(x) = \begin{cases} \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right) \Big/ (\Phi(p) - \Phi(-p)), & \text{if } -p \leq x \leq p, \\ 0, & \text{otherwise.} \end{cases}$$

We continue experiments on the Boston Housing Prices dataset. Apart from the strategy for noise selection, all other experimental settings remain the same as those described in Section 4.3. We investigate whether the ML utility of the new watermarked dataset, with the improved strategy for noise selection, shows improvement compared to the original watermark dataset. Additionally, we examine the robustness of the new watermark dataset against alteration attacks, with the attack strategy being the same as before. Table 25 displays the results for the original dataset, the watermarked dataset with uniform random perturbation, and the watermarked dataset with random perturbation based on a normal distribution probability. Table 26 displays the results for the new watermark dataset under alteration attacks.

Table 25: MSEs on  $D_o$  and  $D_w$ s

Dataset	$D_o$	Completely Random	Normal Distribution
MSE	24.8	25.6	25.0

Table 26: Z-scores and MSEs

Alteration Attack (%)	20	40	60	80	100
Z-score	5.46	3.88	2.18	0.794	-0.690
MSE	25.9	27.1	34.9	33.8	41.8

The experimental results indicate that perturbing data based on sampling from a normal distribution probability effectively reduces the impact of watermarking on the ML utility of the dataset while still maintaining robustness.

### 5.2 Key Cell Selection

For the watermark embedding algorithm (Algorithm 1), the key cell selection strategy is random. It means that data points crucial to ML model performance may be perturbed, thus significantly reducing

ML utility. Data valuation methods like Shapley value [20] can be utilized to select key cells with less contribution to ML utility. During the watermark embedding process, we calculate the Shapley value of each tuple in advance and select  $n_w$  tuples with the smallest Shapley values as key cells. The ML utility of watermark datasets under two different strategies is shown in Table 27.


**Table 27: Results of MSEs**

Dataset	Completely Random	Shapley Strategy
MSE	24.1	23.3

The results show that our new key cell selection strategy with Shapley value can reduce data distortion during watermark embedding and further improve the non-intrusiveness of TabularMark.

### 5.3 Matching

The matching algorithm (Algorithm 3) uses a row-by-row search strategy to locate key cells, which is very inefficient, especially when the dataset is large. We explore a more efficient algorithm by adding a sorting step and a binary search step, which can theoretically reduce the time complexity from  $O(n^2)$  to  $O(n \log n)$ . We use the Boston Housing Prices dataset to verify the improvement in time efficiency. Figure 11a shows the z-score results of the two strategies under insertion attacks. Figure 11b shows the time consumption of the two strategies.



**Figure 11: Z-scores and time on  $\beta$ s**

The results show that our new matching strategy can greatly reduce time consumption without affecting the effect of watermark detection.

### 5.4 Multi-Relational Database

Currently TabularMark considers that only one relation needs to be watermarked. Multiple relations are also common, especially in multi-relational databases. In the context of ML, multi-relational databases consist of multiple tables, each potentially containing different attributes related to a central theme or dataset. Two possible approaches can be utilized to watermark multi-relational databases. First, when preparing such databases for ML tasks, it is a common practice to merge these various tables into a single, comprehensive table. This process involves joining the tables based on the foreign keys so that the resultant table includes the target attribute (the main variable of interest, such as a label in supervised learning) and all other attributes that serve as features for model training. Once this consolidated table is created, TabularMark can be effectively

utilized. Second, each individual relation (or table) within the multi-relational database can be watermarked separately before joining. This approach has a significant advantage: it supports detection efforts even when only a subset of the relations is combined.

## 6 CONCLUSION

In this paper, we propose a simple yet effective watermarking scheme for tabular datasets in ML, TabularMark, showing superior detectability, non-intrusiveness, and robustness. TabularMark is inspired by the randomness in the deviation between a suspicious dataset and the corresponding original dataset. When embedding, TabularMark artificially divides the data with deviations from the original data into two divisions and selects from one division to perturb certain cells in the original dataset. One proportion z-test is adopted in detection to identify these intentional perturbations and determine the ownership of the dataset. To detect the watermark in datasets where some tuple sequences are inconsistent with the original dataset, we propose a matching algorithm to locate tuples containing key cells. Experimental results on real and synthetic datasets validate the effectiveness of TabularMark and demonstrate its applicability to various ML datasets and models.

The main limitation of TabularMark is that it uses a non-blind approach to detect watermarks, which requires the knowledge of the original dataset during the detection process. It is a valuable improvement direction to study a blind version of TabularMark so that it can be extended to watermarking a continuously updated database without accessing the original dataset. Beyond the blind property, several other related and practical directions are worthy of future research as well. While TabularMark provides a general minimal effect in the performance of various ML models trained on tabular datasets, minimizing the performance effect for a specific ML model remains a pivotal challenge. Meanwhile, considering more possible attack methods, supporting more types of datasets, and enlarging the scope of the threat model can be further studied.

## ACKNOWLEDGMENT

The authors would like to thank the anonymous reviewers for their helpful comments. This work was supported in part by the National Key RD Program of China (2021YFB3101100, 2022YFB3103401), NSFC (62102352, U23A20306), The Zhejiang Province Pioneer Plan (2024C01074), JSPS KAKENHI (JP23K24851), JST PRESTO (JPMJPR23P5) and JST CREST (JPMJCR21M2).

## REFERENCES

- [1] Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. 2016. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*. 308–318.
- [2] Rakesh Agrawal and Jerry Kiernan. 2002. Watermarking relational databases. In *Vldb '02: Proceedings of the 28th International Conference on Very Large Databases*. Elsevier, 155–166.
- [3] Sajjad Bagheri Baba Ahmadi, Gongxuan Zhang, Mahdi Rabbani, Lynda Boukela, and Hamed Jelodar. 2021. An intelligent and blind dual color image watermarking for authentication and copyright protection. *Applied Intelligence* 51 (2021), 1701–1732.
- [4] E. Alpaydin and Fevzi. Alimoglu. 1998. Pen-Based Recognition of Handwritten Digits. UCI Machine Learning Repository. DOI: <https://doi.org/10.24432/C5MG6K>.
- [5] Sergul Aydore, William Brown, Michael Kearns, Krishnaram Kenthapadi, Luca Melis, Aaron Roth, and Ankit A Siva. 2021. Differentially private query release

- through adaptive projection. In *International Conference on Machine Learning*. PMLR, 457–467.
- [6] K. Bache and M. Lichman. 2013. UCI Machine Learning Repository. <http://archive.ics.uci.edu/ml>
  - [7] Roberto J Bayardo and Rakesh Agrawal. 2005. Data privacy through optimal k-anonymization. In *21st International conference on data engineering (ICDE'05)*. IEEE, 217–228.
  - [8] Jock Blackard. 1998. Covertype. UCI Machine Learning Repository. DOI: <https://doi.org/10.24432/C50K5N>.
  - [9] Vadim Borisov, Kathrin Seßler, Tobias Leemann, Martin Pawelczyk, and Gjergji Kasneci. 2023. Language Models are Realistic Tabular Data Generators. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net. <https://openreview.net/pdf?id=cEygmnQNOel>
  - [10] Leo Breiman. 2001. Random forests. *Machine learning* 45 (2001), 5–32.
  - [11] Ji-Won Byun, Ashish Kamra, Elisa Bertino, and Ninghui Li. 2007. Efficient k-anonymization using clustering techniques. In *International conference on database systems for advanced applications*. Springer, 188–200.
  - [12] I-Fei Chen and Chi-Jie Lu. 2017. Sales forecasting by combining clustering and machine-learning techniques for computer retailing. *Neural Computing and Applications* 28 (2017), 2633–2647.
  - [13] Tianqi Chen and Carlos Guestrin. 2016. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*. 785–794.
  - [14] Miranda Christ, Sam Gunn, and Or Zamir. 2024. Undetectable watermarks for language models. In *The Thirty Seventh Annual Conference on Learning Theory*. PMLR, 1125–1139.
  - [15] Paulo Cortez, A. Cerdeira, F. Almeida, T. Matos, and J. Reis. 2009. Wine Quality. UCI Machine Learning Repository. DOI: <https://doi.org/10.24432/C56S3T>.
  - [16] Matthew F Dixon, Igor Halperin, and Paul Bilokon. 2020. *Machine learning in finance*. Vol. 1170. Springer.
  - [17] Linkang Du, Zheng Zhu, Min Chen, Shouling Ji, Peng Cheng, Jiming Chen, and Zhikun Zhang. 2024. WIP: Auditing Artist Style Pirate in Text-to-image Generation Models. In *Proceedings of the Network and Distributed System Security Symposium, NDSS 2024*. <https://www.ndss-symposium.org/ndss-paper/autodraft-534/>
  - [18] Joseph L Fleiss, Bruce Levin, and Myunghee Cho Paik. 2013. *Statistical methods for rates and proportions*. John Wiley & Sons.
  - [19] Richard Forsyth. 1990. Zoo. UCI Machine Learning Repository. DOI: <https://doi.org/10.24432/C5R59V>.
  - [20] Amirata Ghorbani and James Zou. 2019. Data shapley: Equitable valuation of data for machine learning. In *International conference on machine learning*. PMLR, 2242–2251.
  - [21] Ali Hamadou, Xingming Sun, Saeed Shah, and Lingyun Gao. 2011. A weight-based semi-fragile watermarking scheme for integrity verification of relational data. *International Journal of Digital Content Technology and its Applications* 5, 8 (2011), 148–157.
  - [22] David Harrison and Daniel L Rubinfeld. 1978. Hedonic housing prices and the demand for clean air. *Journal of Environmental Economics and Management* 5, 1 (1978), 81–102. [https://doi.org/10.1016/0095-0696\(78\)90006-2](https://doi.org/10.1016/0095-0696(78)90006-2)
  - [23] Donghui Hu, Dan Zhao, and Shuli Zheng. 2018. A new robust approach for reversible database watermarking with distortion control. *IEEE Transactions on Knowledge and Data Engineering* 31, 6 (2018), 1024–1037.
  - [24] Jing-Zhi Huang and Zhan Shi. 2023. Machine-Learning-Based Return Predictors and the Spanning Controversy in Macro-Finance. *Manag. Sci.* 69, 3 (2023), 1780–1804. <https://doi.org/10.1287/MNSC.2022.4386>
  - [25] Min-Shiang Hwang, Ming-Ru Xie, and Chia-Chun Wu. 2020. A reversible hiding technique using LSB matching for relational databases. *Informatica* 31, 3 (2020), 481–497.
  - [26] Fei Jiang, Yong Jiang, Hui Zhi, Yi Dong, Hao Li, Sufeng Ma, Yilong Wang, Qiang Dong, Haipeng Shen, and Yongjun Wang. 2017. Artificial intelligence in healthcare: past, present and future. *Stroke and vascular neurology* 2, 4 (2017).
  - [27] Muhammad Kamran and Muddassar Farooq. 2018. A Comprehensive Survey of Watermarking Relational Databases Research. *CoRR abs/1801.08271* (2018). arXiv:1801.08271 <http://arxiv.org/abs/1801.08271>
  - [28] Muhammad Kamran, Sabah Suhail, and Muddassar Farooq. 2013. A robust, distortion minimizing technique for watermarking relational databases using once-for-all usability constraints. *IEEE Transactions on Knowledge and Data Engineering* 25, 12 (2013), 2694–2707.
  - [29] John Kirchenbauer, Jonas Geiping, Yuxin Wen, Jonathan Katz, Ian Miers, and Tom Goldstein. 2023. A watermark for large language models. In *International Conference on Machine Learning*. PMLR, 17061–17084.
  - [30] Wenling Li, Ning Li, Jianen Yan, Zhaoxin Zhang, Ping Yu, and Gang Long. 2022. Secure and high-quality watermarking algorithms for relational database based on semantic. *IEEE Transactions on Knowledge and Data Engineering* (2022).
  - [31] Yan Li, Junwei Wang, and Hongyong Jia. 2020. A robust and reversible watermarking algorithm for a relational database based on continuous columns in histogram. *Mathematics* 8, 11 (2020), 1994.
  - [32] Chia-Chen Lin, Thai-Son Nguyen, and Chin-Chen Chang. 2021. LRW-CRDB: Lossless robust watermarking scheme for categorical relational databases. *Symmetry* 13, 11 (2021), 2191.
  - [33] Douglas C Montgomery, Elizabeth A Peck, and G Geoffrey Vining. 2021. *Introduction to linear regression analysis*. John Wiley & Sons.
  - [34] Chejarla Venkat Narayana, Chinta Lakshmi Likhitha, Syed Bademiya, and Karre Kusumanjali. 2021. Machine learning techniques to predict the price of used cars: predictive analytics in retail business. In *2021 second international conference on electronics and sustainable communication systems (ICESC)*. IEEE, 1680–1687.
  - [35] Warwick Nash, Tracy Sellers, Simon Talbot, Andrew Cawthorn, and Wes Ford. 1995. Abalone. UCI Machine Learning Repository. DOI: <https://doi.org/10.24432/C55C7W>.
  - [36] Ahmet Murat Ozbayoglu, Mehmet Ugur Gudelek, and Omer Berat Sezer. 2020. Deep learning for financial applications: A survey. *Applied Soft Computing* 93 (2020), 106384.
  - [37] Neha Patki, Roy Wedge, and Kalyan Veeramachaneni. 2016. The synthetic data vault. In *2016 IEEE international conference on data science and advanced analytics (DSAA)*. IEEE, 399–410.
  - [38] Zeyu Qin, Liuyi Yao, Daoyuan Chen, Yaliang Li, Bolin Ding, and Minhao Cheng. 2023. Revisiting Personalized Federated Learning: Robustness Against Backdoor Attacks. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, KDD 2023, Long Beach, CA, USA, August 6-10, 2023*, Ambuj K. Singh, Yizhou Sun, Leman Akoglu, Dimitrios Gunopoulos, Xifeng Yan, Ravi Kumar, Fatma Ozcan, and Jieping Ye (Eds.). ACM, 4743–4755. <https://doi.org/10.1145/3580305.3599898>
  - [39] R. Quinlan. 1993. Auto MPG. UCI Machine Learning Repository. DOI: <https://doi.org/10.24432/C5859H>.
  - [40] Khansa Rasheed, Adnan Qayyum, Mohamed Ghaly, Ala I. Al-Fuqaha, Adeel Razi, and Junaid Qadir. 2022. Explainable, trustworthy, and ethical machine learning for healthcare: A survey. *Comput. Biol. Medicine* 149 (2022), 106043. <https://doi.org/10.1016/J.COMPBIOMED.2022.106043>
  - [41] Theodoros Rekatsinas, Xu Chu, Ihab F. Ilyas, and Christopher Ré. 2017. HoloClean: Holistic Data Repairs with Probabilistic Inference. *Proc. VLDB Endow.* 10, 11 (2017), 1190–1201. <https://doi.org/10.14778/3137628.3137631>
  - [42] David Rhind. 1992. Data access, charging and copyright and their implications for geographical information systems. *International Journal of Geographical Information Systems* 6, 1 (1992), 13–30.
  - [43] Mohamed Shehab, Elisa Bertino, and Arif Ghafoor. 2007. Watermarking relational databases using optimization-based techniques. *IEEE transactions on Knowledge and Data Engineering* 20, 1 (2007), 116–129.
  - [44] Radu Sion, Mikhail J. Atallah, and Sunil Prabhakar. 2004. Rights Protection for Relational Data. *IEEE Trans. Knowl. Data Eng.* 16, 12 (2004), 1509–1525. <https://doi.org/10.1109/TKDE.2004.94>
  - [45] Devendra Swami, Alay Dilipbhai Shah, and Subhrajit K. B. Ray. 2020. Predicting Future Sales of Retail Products using Machine Learning. *CoRR abs/2008.07779* (2020). arXiv:2008.07779 <https://arxiv.org/abs/2008.07779>
  - [46] Mingtian Tan, Tianhao Wang, and Suresh Jha. 2023. A Somewhat Robust Image Watermark against Diffusion-based Editing Models. *CoRR abs/2311.13713* (2023). <https://doi.org/10.48550/ARXIV.2311.13713> arXiv:2311.13713
  - [47] Toan V Tran and Li Xiong. 2024. Differentially Private Tabular Data Synthesis using Large Language Models. *arXiv preprint arXiv:2406.01457* (2024).
  - [48] Haiming Wang, Zhikun Zhang, Min Chen, and Shibo He. 2023. Making Watermark Survive Model Extraction Attacks in Graph Neural Networks. In *IEEE International Conference on Communications, ICC 2023, Rome, Italy, May 28 - June 1, 2023*. IEEE, 57–62. <https://doi.org/10.1109/ICC45041.2023.10278974>
  - [49] Jenna Wiens and Erica S Shenoy. 2018. Machine learning for healthcare: on the verge of a major shift in healthcare epidemiology. *Clinical Infectious Diseases* 66, 1 (2018), 149–153.
  - [50] Xiangrong Xiao, Xingming Sun, and Minggang Chen. 2007. Second-LSB-dependent robust watermarking for relational database. In *Third international symposium on information assurance and security*. IEEE, 292–300.
  - [51] Liyang Xie, Kaixiang Lin, Shu Wang, Fei Wang, and Jiayu Zhou. 2018. Differentially private generative adversarial network. *arXiv preprint arXiv:1802.06739* (2018).
  - [52] Lei Xu, Maria Skoularidou, Alfredo Cuesta-Infante, and Kalyan Veeramachaneni. 2019. Modeling tabular data using conditional gan. *Advances in neural information processing systems* 32 (2019).
  - [53] Mohamed Yamni, Hicham Karmouni, Mhamed Sayyouri, and Hassan Qjidaa. 2022. Efficient watermarking algorithm for digital audio/speech signal. *Digital Signal Processing* 120 (2022), 103251.
  - [54] Xiaorui Zhang, Xun Sun, Xingming Sun, Wei Sun, and Sunil Kumar Jha. 2022. Robust Reversible Audio Watermarking Scheme for Telemedicine and Privacy Protection. *Computers, Materials & Continua* 71, 2 (2022).
  - [55] Xin Zhong, Pei-Chi Huang, Spyridon Mastorakis, and Frank Y Shih. 2020. An automated and robust image watermarking scheme based on deep neural networks. *IEEE Transactions on Multimedia* 23 (2020), 1951–1961.