

Achieving Privacy and Tracing Unauthorised Data Usage: Anonymisation-based Fingerprinting for Private Data

Tanja Šarčević

SBA Research, Vienna, Austria
tsarcevic@sba-research.org

Rudolf Mayer

SBA Research, Vienna, Austria
rmayer@sba-research.org

Philipp Adler

Vienna University of Technology, Austria

Abstract—Since many types of data nowadays contain personally identifiable information about individuals, it is important to apply privacy protection techniques to mitigate disclosure risks. One approach is k -anonymity, where hiding the identity within a group of k similar entities reduces the risk of re-identification. Another risk when distributing data is the loss of control over their further re-distribution and sharing. This risk is frequently addressed by fingerprinting, a method that allows to identify the recipient of a specific copy of the data, by embedding a generally invisible mark.

In this paper, we specifically implement and adapt an intrinsic fingerprint scheme that makes use of k -anonymity and the fact that multiple, differently perturbed versions of a dataset can be found that all fulfil a certain k -anonymity, and share a rather similar level of data precision. Thus, these different datasets can be seen each as a fingerprinted version of the original.

One research question we address in this paper is the evaluation of the most common generalisation algorithms according to their generalisation strategy and their influence on data utility and the number of resulting release candidates, i.e. fingerprints and execution time. In addition, we investigate the properties and robustness of these fingerprints against intentional (adversarial) manipulation through attack simulations. We further provide recommendations and guidelines on how fingerprinting can be best achieved based on the results of our evaluation.

Index Terms—anonymisation, fingerprinting, robustness

I. INTRODUCTION

Data publishing is required in many settings, e.g. for collaborative data analysis, or if the analysis shall be outsourced to a third party, or if data in general be made available to the general public for research. However, data publishing comes with several pitfalls. On the one hand, disclosure risks threaten the privacy of the individuals in the dataset, e.g. via (re-)identification attacks. Anonymisation strategies have been proposed to deal with these risks, among them k -anonymity and its extensions such as l -diversity. These reduce the disclosure risks, normally at the cost of data utility, as they modify statistical properties, as well as the utility for e.g. predictive machine learning models trained thereupon.

On the other hand, if the data is only made available to a selected number of users, further re-distribution of the data is generally not desired by the original data controller. Thus, illicit re-distribution shall be deterred from and made detectable. Fingerprinting is a reactive technique that can identify the original recipient from data that is re-distributed,

by embedding a recipient-specific mark into the data, e.g. by altering attribute values in relational data. This approach thus inevitably lowers the utility of the fingerprinted data. A scheme [1] for fingerprinting relational data proposes to combine k -anonymisation of a dataset with fingerprinting, by leveraging the fact that there are generally multiple solutions to achieve the same level of k -anonymity, namely by generalising different attributes to a different degree. The level to which each attribute then gets generalised represents an intrinsic and recognisable property of that one specific instance of a k -anonymised dataset of a specific original dataset. The scheme thus proposes to generate multiple different solutions and treat each of them as a uniquely fingerprinted version of the original dataset – which can thus be released to the respective recipients. This relies on the assumption that there are (enough) different solutions to provide a sufficient number of possible recipient copies that, however, should not differ (too much) in data utility. With this assumption, fingerprinting thus comes at no additional “cost” in terms of data utility, then what is already introduced by applying k -anonymity, and becomes an interesting option in case the dataset will anyhow be anonymised before releasing it – likely a reasonable assumption in many scenarios, considering data protection regulations

While the proposed scheme poses some intriguing properties, it has neither been implemented nor empirically validated until now. In this paper, we provide an implementation and evaluation of the key properties of the scheme. Our contributions are thus:

- An implementation of the k -anonymity based fingerprinting scheme
- An evaluation of the assumption that there is a sufficiently large number of roughly equivalent k -anonymous solutions
- An evaluation of the robustness of the scheme against attacks, i.e. intentional, malicious modifications to the released data with the goal of reducing the effectiveness of the fingerprint
- Evaluation of the success of the collusion attack, where the recipients collaborate to remain untraceable, and the discussion on the efficiency of the scheme against this

attack

The remainder of this paper is structured as follows. Section II discusses related work, before Section III details the approach for k -anonymous fingerprinting. We evaluate the utility and robustness of the scheme in Section IV, and provide conclusions in Section V.

II. RELATED WORK

A. k -anonymity

k -anonymity was first introduced by Samarati and Sweeney [2] to obfuscate sensitive datasets in order to be able to share them with other parties. In a dataset, we can generally distinguish different types of attributes. On the one hand, (*directly identifying attributes*) directly reveal the identity of a data record and should be removed from the dataset before publishing. *Quasi-identifiers* (QIs) do not directly identify a person but may become uniquely identifying when used in combination with other QIs. QIs often hold significant information, which is required in analysis processes. In medical analysis, for example, it is often important to differentiate between age groups, the type of job, or information on the location of the residence of patients. Thus, this information cannot simply be completely omitted. *Sensitive data* is contained in attributes that for example hold information about a certain type of illness or the salary of an individual. These are generally the main target in statistical analysis, and can therefore not be omitted or obfuscated.

As each level of generalisation invokes an increasing loss of specificity, the goal of k -anonymisation is to minimise the overall information loss, making it an NP-hard problem [3]. Several different, mostly heuristic, approaches have been proposed for optimising the level of suppression and generalisation for achieving a specific level of k -anonymity. Samarati [4] introduces the concept of *minimal generalisation* that captures the property of the release process not to distort the data more than needed to achieve k -anonymity. There are a few proposed globally optimal anonymisation algorithms, such as Flash [5], a greedy depth-first manner algorithm and OLA (Optimal Lattice Anonymisation) [6].

1) *Data Precision / Utility Metrics*: Protecting the privacy of data inevitably leads to some loss of information. Since the usefulness of data often depends on the use case, which from the publisher's point of view is typically unknown, choosing a fitting metric is challenging. A wide variety of general-purpose data precision, utility or quality metrics have been proposed; according to [7], these metrics are divided into the *cell-* and *column-oriented* and *row-oriented* quality models, and we will briefly introduce representative metrics. A column-oriented metric *height* [4] estimates the information loss as the sum of the generalisation level steps applied to the sets of microdata. However, since the total number of possible generalisation levels is not taken into account in the measurement, it can happen that complete suppression leads to the same result as a simple generalisation. Consequently, these circumstances might have a huge impact on the final result. As a cell-oriented metric, *precision* quantifies the data quality based

on normalised generalisation levels. Accordingly, the amount of distortion of attribute values is measured as the generalisation level of an attribute value relative to the height of the attribute's generalisation hierarchy. Row-oriented metrics measure the information loss according to modifications on the sample level. For instance, a *modified discernibility metric* (DM^*) [8] measures the size of the equivalence classes and penalises suppressed data records.

B. Watermarking and Fingerprinting

Watermarking and fingerprinting are forms of steganography, the information-hiding techniques that embed information about the owner (watermarking) and the recipient (fingerprinting) into digital data. Watermarking allows identifying the *ownership* of digital objects by embedding secret owner-specific information into the dataset. Fingerprinting can be seen as an extension of watermarking, which in addition encodes, for each copy, an identifier of the particular recipient. Fingerprint combines thus secret owner-specific and recipient-specific information embedded in a specific release of a digital object. Fingerprinting and watermarking of digital data first appeared for multimedia data and have been extensively studied over the last two decades [9], [10]. For tabular data, which is the focus of this work, most of the current state-of-the-art fingerprinting methods extend the watermarking technique proposed by Agrawal [11], such as [12], [13], [14]. The techniques in principle contain two algorithms: fingerprint insertion (embedding) and fingerprint detection (extraction). In the embedding phase, the fingerprint is created as a bit-string uniquely representing the recipient and additionally encoding the owner's secret and is embedded into the least significant bits (LSBs) of the data values. The extraction is the reverse process where the encoded values are extracted from the LSBs and the fingerprint can be assigned to the specific recipient.

1) *k -anonymity based fingerprinting*: Schrittwieser et al. [15] propose a scheme that is based on k -anonymity. They note that as multiple solutions for achieving the same level of k -anonymity exist, and that these datasets then differ in the granularity of the attributes, this granularity can serve as a fingerprint identifying the exact copy that was released. The approach is based on the assumption that the different solutions are also comparable to each other in data quality. This scheme uses a slightly modified version of the OLA algorithm [6] for revealing all possible k -anonymous datasets. The algorithm incorporates the following steps (description from [15]):

- 1) Define a minimum k for the k -anonymity criterion, the minimum and maximum levels of data loss l_{min} and l_{max} and the data precision metric to be used
- 2) Define the generalisation strategies for each quasi-identifier
- 3) Calculate the lattice diagram derived from all possible generalisations
- 4) Choose a node at middle height and decide whether it is at least k -anonymous

- a) In case it is not, rule out all nodes below in the lattice diagram.
- b) In case it is, mark all nodes above the chosen one as possible solutions.
- 5) Start with step four for the remaining sub-graph, similar to the original algorithm
- 6) In case no sub-graph is left, start by choosing another initial node at middle height and proceed with step four until all nodes are evaluated
- 7) For each at least k -anonymous solution, calculate data precision and the actual k . Remove all solutions with data precision outside the bounds of l_{min} and l_{max} .
- 8) Classify and cluster the solutions by their data precision
- 9) Create "similar" sets of microdata based on results in one cluster and distribute them to the recipients

The watermark detection relies on a pattern list that stores the generalisations performed to achieve a certain k -anonymised dataset and the recipient of this dataset. From the dataset in question, the generalisation patterns are observed (by checking the granularity of the values for each attribute), and the match is retrieved from the pattern list. Note that this detection method assumes that an attacker does not voluntarily increase the generalisation of some of the attributes in the dataset they obtained. In further work [16], the authors propose a collusion filter applied before data distribution and an approach to resolve the collusion attack (cf. Section II-B2), i.e. identify the recipients participating in the malicious collaboration. The design of the collusion filter also implies a theoretical bound for the number of collusion-free fingerprints. One major drawback of both the base and extended approach is the lack of evaluation and a lack of study on the success of the collusion attack.

There are other approaches that simultaneously achieve data privacy and ownership protection or tracing; Bertino et al. use binning [17], Gambis et al. use (α, β) -sanitisation [18], Ji et al. [19] differential privacy, [20] watermark data by firstly applying the classical watermark, then k -anonymity.

2) *Robustness of Fingerprinting Schemes*: The robustness of a fingerprinting scheme is measured as the resilience of the scheme against modifications of the dataset, which can happen as a result of benign updates or malicious attacks. The resilience manifests as the success of the detection scheme to extract the fingerprint from a data copy and associate it to its correct recipient. In literature, frequently mentioned attacks against fingerprinting schemes for relational datasets are different manipulations of LSBs (e.g. flipping attack where the attacker flips a portion of LSBs of data values), attribute-oriented alterations (deletion, transformation) or row-oriented alterations (addition, deletion) or [21]. Another common attack is *additive attack*, where the attacker produces and embeds their own fingerprint on top of the existing one, to try and claim false ownership of the data. An attack specific to fingerprinting is a *collusion attack* where multiple malicious recipients collaborate to obfuscate or hide the received fingerprints.

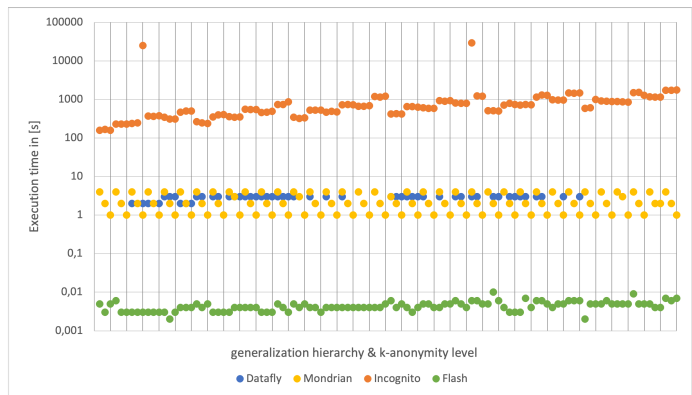


Fig. 1: k -anonymity algorithm performance on KDD; the x-axis represents the increasing privacy level k combined with increasing complexity of generalisation hierarchies

III. FINGERPRINTING VIA k -ANONYMITY

A. Approach

The starting point of our approach is the algorithm of Schrittwieser et. al. [15] (cf. Section II-B1). In the following, we further detail the steps we adapted.

a) *Step 1: Definition of input parameters (number of possible data users, clustering method and data utility metrics)*: The step encompasses the definition of all input parameters, including the dataset, privacy level k , suppression limit $MaxSup$, data utility metric (*DUM*) for measuring the applied information loss during generalisation, clustering method and the number of data users u to whom the unique fingerprints will be distributed. The specification of the number of recipients is necessary to guarantee a minimised intra-cluster solution with at least the required number of fingerprints. To provide roughly the same level of precision to all data users, the choice of an appropriate utility metric is an important aspect. In Section IV, we demonstrate a few metrics that are particularly suitable for our goal: the generation of a cluster with high data utility and low variance. The choice of the clustering algorithm also impacts the final cluster result. Furthermore, we could observe from experimentation that defining l_{min} and l_{max} beforehand is not beneficial for the algorithm, as it usually requires post-processing steps (in Step 7), where these attributes need to be relaxed to allow a larger range of candidate solutions. For this, we adapt the baseline approach by avoiding the pre-definition of these thresholds and consequently remove step 7 from the baseline algorithm.

b) *Step 2: DGH definition*: The definition of the domain generalisation hierarchy for each quasi-identifier serves as an indicator for the theoretical upper limit of the number of solutions, $\prod_{i=1}^N n_i$, where N denotes the number of QIs and n_i the number of applicable generalisation steps of the corresponding QI. Although this formula does not apply in practice, since not all anonymised datasets meet the specified privacy criterion, a high number of possible DGH_{D_i} is recommended to increase the chance of finding further solution candidates.

c) *Steps 3-6: k-anonymisation:* Achieving k -anonymity can be treated as an independent modular part of the approach, therefore we combine steps 3 – 6, originally describing the OLA algorithm, into one step that can be achieved as well with any other k -anonymity algorithm. Due to a large amount of computation (to achieve multiple k -anonymous datasets), we opted for the most time-efficient algorithm, Flash, according to our comparative study on the KDD datasets (Figure 1).

d) *Step 8: Clustering the solution space:* For selecting release candidates it is important to evaluate the clusters according to several criteria. Since multiple recipients get mutually different anonymised copies of the same dataset, it is important to keep the copies on a similar quality level to ensure fairness; we also need to be able to provide a sufficient number of anonymised copies. Therefore, we model our criteria as follows:

- The *average* utility of the fingerprints (anonymised datasets) within the cluster is given by the cluster *centroid*. Clusters with a high data utility are more useful for the recipients.
- The utility *variance* is an indicator of similarity between the copies of one cluster. Hence, clusters with a small variance will ensure fairness in the distribution of similarly useful anonymised datasets.
- The *number of nodes*, or *cluster cardinality* shall be at least as large as the number of expected data recipients, and rather be maximised in case this number is not fully known upfront.

An optimal cluster for distribution should satisfy all three of the above criteria simultaneously. We thus use the following weighted average sum of the three quality criteria to measure the quality of a cluster C_i .

$$Q_{C_i} = \frac{f_i}{f_{max}} \times \frac{1}{3} + \frac{c_i}{c_{max}} \times \frac{1}{3} + (1 - \frac{v_i}{v_{max}}) \times \frac{1}{3} \quad (1)$$

, where v is variance, c the centroid and f the number of fingerprints.

e) *Step 9: Distribution of the datasets from the optimal cluster:* The datasets from the resulting cluster undergo further considerations regarding collusion possibilities discussed in Section III-B. The *collusion filter* is applied to eliminate datasets with a high risk of collaboration and the final set of k -anonymous solutions is distributed to the recipients.

B. Collusion filter

In [16], the authors set the theoretical upper bound of the number of recipients to the number of QIs in the dataset. According to their discussion, the "signature" of recipient A in collusion is one specific QI, where recipient A has got the largest granularity of that QI (i.e. the lowest generalisation level) out of all other recipients. Since this needs to be true for all recipients to guarantee the collision resolution, the number of recipients can only be as large as the number of QIs. While this statement in theory indeed limits the number of recipients, it is based on the assumption that the recipients can successfully collaborate by matching the other attributes

in the dataset (not affected by the k -anonymity). In practice, this is not easily achievable, since the non-sensitive attributes have to satisfy some specific properties. We demonstrate and evaluate the collusion attack in Section IV-D.

C. Robustness of fingerprints

Fingerprinting schemes should be robust against different attacks that aim to prevent the correct detection of the fingerprint. We analyse the robustness using measures proposed in [22]:

- *Misattribution false hit*, the probability of detecting an incorrect but valid fingerprint that belongs to a different recipient. This measure describes the success of the attacker's **obfuscation**, the ability to disguise their own identity into someone else's.
- *False negative*, the probability of not detecting the valid fingerprint from fingerprinted data i.e. the detected fingerprint finds no direct match to the released ones. This measure describes **hiding**, the likelihood of concealing behind a group of potential fingerprints such that no recipient can be suspected. The higher the group, the higher the success rate for the attacker.

To summarise, both attacker characteristics attempt to modify the anonymised dataset in such a way that the underlying resulting fingerprint can not be directly assigned to the original data recipient. The attacks evaluated in Section IV-C follow this procedure:

- S1: *Identification of manipulable QIs:* Any value that is not fully suppressed can potentially be manipulated.
- S2: *Attribute analysis:* Using the unique attribute values, an attacker can determine which data type an attribute belongs to. For instance, the continuous attributes will have values in integer or decimal formats, however, values represented in ranges potentially also represent generalised continuous attributes.
- S3: *Attribute selection:* The attacker chooses any number of QIs from the manipulable set, which he manipulates to subvert the embedded fingerprint pattern.
- S4: *Attribute manipulation:* Any selected QI is generalised. The final result of the fingerprint alteration process ultimately depends on the existing background knowledge and the associated data type. The straightforward way to generalise QIs is to fully suppress them (i.e. replace the values with the most general value, *). Whether the attacker has the ability to apply other transformations depends on their access to the original hierarchies of the QIs, or their ability to guess the correct generalisations. Step S4 requires distinction of the attackers based on the knowledge, skills, and resources available to them.
 - An *uninformed attacker* has no information about the generalisation patterns or any background knowledge that could be used to generalise the attribute to anything more granular than the full suppression (*).
 - *Attacker with some background knowledge:* an attacker who does not have an access to the generalisation hierarchies can still transform QIs using assumptions based on

common practices in generalisation patterns. For instance, demographic attributes such as zip code, place of birth etc. are often part of the datasets that are anonymised. Hence, some generalisation patterns can resemble each other in similar domains, or might even be standardised. E.g., it is a common practice to generalise zip codes such that the last digits are gradually suppressed on each generalisation level (90011→9001*→900**→90***→9****). Similarly, numerical attributes, such as *age* are usually generalised withing ranges, and these ranges might double in each hierarchy level (32→[30-34]→[30-39]→...), etc.

- An *Informed attacker* is familiar with the generalisation hierarchies used for each QI (either provided originally by the data owner, or obtained through malicious actions). Manipulations in this case exactly follow the generalisation patterns of the original anonymisation process. The probability of transforming the fingerprint to another valid one (the one that belongs to a different recipient) is hence higher compared to less informed attackers.

IV. EVALUATION

To evaluate the fingerprinting scheme, we conduct a series of experiments with two real-world datasets that have already been used in several previous studies [6], [23] in the area of *k*-anonymity and also fingerprinting.

A. Datasets

The first dataset is the Adult dataset¹, which is an excerpt from the 1994 US census database; it is through 15 attributes in 32,561 samples. As pre-processing steps, we removed all records with missing values², leaving us with 30,162 samples. The second dataset was initially introduced in the KDD Cup 1998³ (KDD for brevity), and represents the donations of the Paralysed Veterans Association, veterans with spinal cord injuries or diseases with 95,412 observations. After removing instances with missing values, 63,441 of the 95,412 tuples remained. The dataset consists of eight QIs, six of which belong to the continuous data type.

The generalisation hierarchies for both datasets range between two and six levels. With the number of hierarchy steps, we could theoretically achieve 20,736 and 1,128,960 different combinations of generalisations, and thus fingerprints, for Adult and KDD, respectively. However, not all of these generalisations will fulfil the desired level of *k*. In fact, only considering combinations that achieve at least 2-anonymity, the number of transformations in the solution space is 96 for KDD, and 12,096 for Adult.

B. Influence of utility metrics

In this section, we show the effects different utility metrics have on the resulting cluster. The quality of the clusters of

fingerprints is evaluated using Equation (1), where the centroid and variance are calculated using a predefined utility metric.; we evaluate:

- *Height* [4], i.e. the generalisation level of an attribute value
- *non-uniform (NU) entropy* [6] measures the differences in the distributions of attribute values induced by data transformations.
- *Precision* [24] is the generalisation level of an attribute value relative to the height of the attribute's generalisation hierarchy
- *Loss* [25] measures the granularity of data by determining the fraction of an attribute's domain that is covered by the transformed values.
- *Average Equivalence Class size (AEC)* [26] measures the average size of classes of indistinguishable records
- *Modified discernibility (DM*)* [6] is the sum of squared equivalence class sizes
- *Ambiguity Metric (AM)* [27] measures the degree of uncertainty in the anonymised data. The uncertainty of an anonymised record measures the number of tuples in the data domain that could have been generalised to it.

The row-oriented data utility metrics *AEC*, *AM* and *DM** have generally large utility scores and hence produce the clusters with high-utility centroids, which results in a high impact on the weighted average used to determine the general quality of the cluster, as shown in Tables I and III. This in result ranks some clusters with higher cardinality lower according to the weighted average as shown in Table II.

The metrics describing the absolute (*height*) or relative (*precision*) generalisation level are subject to the lower granularity of the obtained utility values, hence a large number of obtained anonymised datasets appear the same in terms of quality. This results in large clusters with low variance, especially evident in Table I.

Furthermore, using different utility metrics, the resulting clusters will differ in size, which in our case denotes the number of fingerprints. This effect is of great relevance when the data holder wants to maximise the number of resulting fingerprints to be able to distribute to many recipients. Hence, a metric that maximises the cluster sizes should be preferred over the others. For instance, the greatest cardinality in overall best-quality clusters is achieved via *height*, which is especially evident from the experiments on Adult data Figures 2 and 3. Nonetheless, a data publisher should be aware of the consequences of clusters with a high number of fingerprints when distributing them: while it might be advantageous to have clusters with many elements when the set of recipients is uncertain, *if* the number of data users is known in advance, clusters with more elements might lead to the disadvantage of having a lower centroid or higher variance.

C. Robustness

In this section, we evaluate the robustness of fingerprinting against tampering and removal towards the two attack types, *obfuscation*, and *hiding*. These attacks are carried by a single

¹<http://archive.ics.uci.edu/ml/datasets/adult>

²Dealing with missing data in the process of anonymisation is left outside of the scope of this manuscript.

³<https://archive.ics.uci.edu/dataset/129/kdd+cup+1998+data>

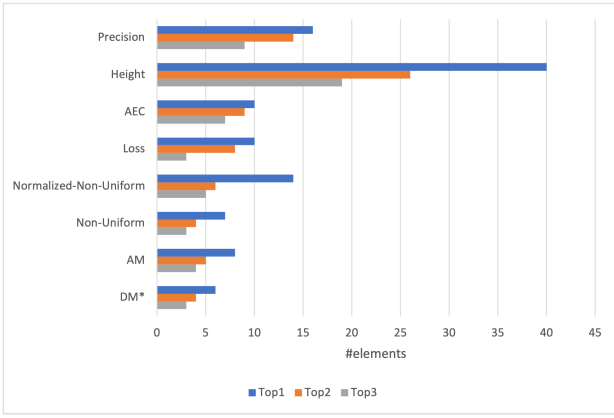


Fig. 2: k-means cluster size results for k:2, gen:[1,1,1,2,3,2,1,2,1] and min 3 recipients

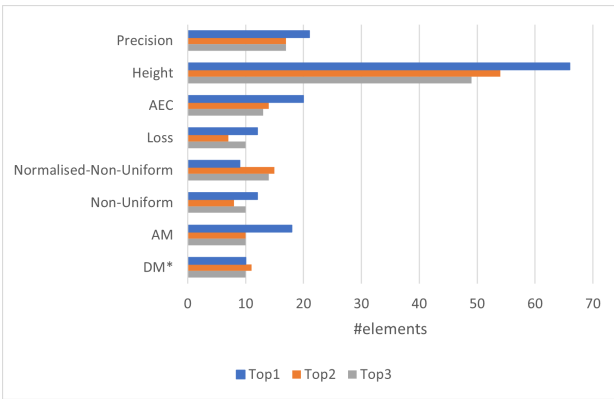


Fig. 3: k-means cluster size results for k=2, gen:[1,3,1,3,4,3,2,3,2] and min 3 recipients

recipient alone, in contrast to the collusion attacks discussed in Section IV-D. Therefore, we inspect the properties of each released dataset in relation to the characteristics of the cluster from which the fingerprints originate. In particular, we investigate the necessary trade-off in connection with the applied attack settings and the success rate according to our two defined manipulation goals. The insights gained from this analysis advise the data publishers on which properties are influential in the generation and distribution of fingerprints and how clusters are designed in a way to remain robust against manipulation.

Data publishers who grant third parties access to the desired dataset have to create a pattern list, which links the recipient and the fingerprint pattern and thus allows tracking of the source in the event of an improper release. As a result of the preceding clustering step, the allocation shown in Tables IV and V was made.

1) *Experiment design*: Based on the assumption that every data recipient can potentially take on the role of an attacker and re-distribute its corresponding dataset without permission to third parties, we designed our experiments in such a way that we allow each recipient any number of modifications, thus

TABLE I: Top 3 optimal clusters per utility metric for the Adult dataset, gen=[1,3,1,2,3,2,1,2,1]

#	metric	#elements	Clustering results		
			centroid	variance	w.avg.
1	AM	18	98,043	0,122	0,753
2	AM	10	96,930	0,031	0,710
3	AM	10	96,922	0,123	0,710
1	AEC	20	89,070	1,383	0,733
2	AEC	14	96,970	0,251	0,730
3	AEC	13	96,685	0,001	0,724
1	DM*	10	93,571	0,221	0,698
2	DM*	11	91,531	0,219	0,696
3	DM*	10	90,593	0,293	0,688
1	Height	66	18,750	0	0,729
2	Height	54	18,750	0	0,669
3	Height	49	25,000	0	0,664
1	Precision	21	22,220	0	0,514
2	Precision	17	22,220	0	0,494
3	Precision	17	22,220	0	0,494
1	Loss	12	35,120	0,176	0,512
2	Loss	7	42,543	3,993	0,509
3	Loss	10	33,580	0,021	0,496
1	NU	12	20,829	0,083	0,464
2	NU	8	26,506	1,659	0,461
3	NU	10	21,963	0,118	0,457
1	Norm.NU	9	33,767	0,191	0,492
2	Norm.NU	15	22,265	0,004	0,484
3	Norm.NU	14	22,217	0,022	0,478

TABLE II: Ranking of the clusters for the Adult dataset, gen=[1,3,1,2,3,2,1,2,1], min k=2, metric=AM

#	#elements	Clustering results		
		centroid	variance	weighted avg.
1	18	98,043	0,122	0,753
...
6	11	93,478	0,057	0,703
7	10	92,669	0,084	0,695
8	10	92,657	0,077	0,695
9	7	96,409	0,004	0,693
10	11	90,394	0,274	0,693
...
16	13	84,040	0,314	0,688
...
19	13	85,437	1,107	0,685
...
25	12	85,719	0,354	0,682

covering the full range of possible manipulation scenarios for evaluation. Since we can not know in advance what background knowledge an attacker possesses, we must consider all attacker skill sets described in Section III-C. We thus perform the following steps; from Step 3 on, they are repeated 100 times to cover many combinations:

- S1: *Cluster selection for distribution*. First of all, a data publisher has to define which cluster he wants to release. This is followed by assigning each recipient a fingerprint.
- S2: *Specification of the attack properties*. Definition of how an attacker approaches the manipulation in terms of the background knowledge of the attacker; either an *informed*

TABLE III: Top 3 optimal clusters per utility metric for the KDD dataset, gen=[1,2,4,3,1,1,1,1]

#	metric	#elements	Clustering results		w.avg.
			centroid	variance	
1	AM	13	91,178	16,798	0,971
2	AM	11	90,222	13,903	0,917
3	AM	11	84,943	58,396	0,892
1	AEC	13	92,298	15,903	0,975
2	AEC	11	91,364	13,123	0,921
3	AEC	11	86,023	54,703	0,896
1	DM*	13	88,588	26,932	0,961
2	DM*	11	87,399	22,575	0,906
3	DM*	9	94,512	4,675	0,882
1	Height	12	24,405	12,391	0,721
2	Height	10	25,000	22,941	0,670
3	Height	9	17,463	12,588	0,620
1	Precision	9	8,798	3,167	0,593
2	Precision	8	15,754	3,921	0,590
3	Precision	8	14,061	4,609	0,585
1	Loss	10	22,822	7,602	0,665
2	Loss	8	22,044	5,660	0,611
3	Loss	8	21,693	3,084	0,611
1	NU	8	9,263	1,207	0,569
2	NU	8	7,239	2,005	0,562
3	NU	8	5,469	1,216	0,556
1	Norm.NU	8	12,766	3,844	0,580
2	Norm.NU	8	3,934	5,380	0,550
3	Norm.NU	7	9,024	1,395	0,543

TABLE IV: Pattern list for chosen clusters and manipulation opportunities per type of attacker - Adult dataset

recipient	utility score	fingerprint pattern	A1	A2	A3
A	96,67	[0,2,0,2,2,2,1,2,1]	7	7	7
B	96,67	[0,2,0,2,3,2,1,1,1]	7	7	7
C	96,67	[0,2,1,2,1,2,1,1,1]	7	7	11
D	96,67	[0,2,1,2,3,1,1,1,1]	7	7	7
E	96,67	[1,1,1,2,1,2,1,1,1]	7	7	11
F	96,67	[1,1,1,2,3,1,1,1,1]	7	7	7
G	96,67	[1,2,0,1,2,2,1,2,1]	7	7	7
H	96,67	[1,2,0,1,3,2,1,1,1]	7	7	7
I	96,67	[1,2,0,2,3,2,1,1,0]	7	7	7
J	96,67	[1,2,1,1,1,2,1,1,1]	7	7	11
K	96,67	[1,2,1,1,3,1,1,1,1]	7	7	7
L	96,67	[1,2,1,2,1,2,1,1,0]	7	7	11

attacker or attacker with some background knowledge.

- S3: *Random selection of an attacker.* From the set of data recipients, a random one is selected who takes on the role of the attacker.
- S4: *Data manipulation.* The attacker tries to manipulate the dataset based on his capabilities in a way that he remains undetected by the detection mechanism. First, the attacker identifies the QIs that can still be manipulated. Depending on the specified number of attributes an attacker is supposed to influence, he selects a random one from the identified candidates and checks whether this can be manipulated according to the corresponding specified criteria, and the number of generalisation steps. If the prerequisites of *Step 2* are met, the attribute values will

TABLE V: Pattern list for chosen clusters and manipulation opportunities per type of attacker - KDD dataset

recipient	utility score	fingerprint pattern	A1	A2	A3
A	98,18	[1,1,1,2,1,1,1,1]	7	11	11
B	96,42	[1,1,1,3,1,1,1,1]	3	5	5
C	96,47	[1,1,2,2,1,1,1,1]	7	7	7
D	93,06	[1,1,2,3,1,1,1,1]	3	3	3
E	98,03	[1,1,3,0,1,1,1,1]	3	3	7
F	96,28	[1,1,3,1,1,1,1,1]	3	5	5
G	93,00	[1,1,3,2,1,1,1,1]	3	3	3
H	96,28	[1,2,1,0,1,1,1,1]	3	5	11
I	92,95	[1,2,1,1,1,1,1,1]	3	7	8
J	92,80	[1,2,2,0,1,1,1,1]	3	3	7

be replaced with the new generalised values. If not, we move on to the next candidate attribute. We repeat this procedure until either the requirements are met, or the list of candidates is empty. In the second case, we discard the release and select a new attacker.

- S5: *Fingerprint detection.* In the last step, the detection phase, the distinct values of the attributes in the dataset under evaluation are compared to the generalisation hierarchy, and the resulting pattern is compared to the pattern list. If no fingerprint exactly matches the detected one, potential candidates are determined based on their distance.

Relevant metrics for our analysis are (i) the attack success rate with respect to two attack targets: *obfuscation* and *hiding* and (ii) the required manipulation, particularly the induced data loss to achieve this objective, which represents the cost of the attack.

Note that the obfuscation success rate is \leq than the hiding success rate, as a successful identity obfuscation implies successful identity hiding, while the latter can be achieved without the former.

We summarised in Tables IV and V the number of manipulations that a recipient can do depending on the attacker role they take (A_1 corresponds to an uninformed attacker, A_2 to some background knowledge and A_3 to an informed attacker).

- 2) *Analysis:* Based on the possible manipulation strategies, we perform the analysis and examine the average success rate as well as the incurred data loss.

Table VI shows the difference in the number of modification strategies, the resulting data loss, and the success rate based on the level of knowledge an attacker possesses for the Adult dataset. Even though an attacker with additional knowledge can manipulate the dataset with less information loss and has 16 more attempts to disguise his identity due to the *QI education*, the average success rate is nearly the same as an uninformed attacker. However, in many settings, the attack comes at a relatively high cost, indicated by the relative utility loss. While an acceptable loss depends on the usage scenario, even the attack with the lowest loss incurs an almost 7% relative loss penalty, while only avoiding detection in a bit more than a third of the cases. Attacks that would succeed with around 70% chance incur a utility loss of around 30%, which seems prohibitively costly in most settings. It can also

be seen that obfuscation is very difficult to achieve with the generalisation hierarchies present in this dataset.

TABLE VI: Attack and detection phase outcomes; Adult

Input parameter				Modification results			
no of attrib.	levels	mixed	knowledge	options	avg. obfuscation success	avg. hiding success	avg. rel. utility loss
1	1	⊥	⊥	36	0,00%	44,44%	8,05%
1	1	⊥	⊥	36	0,00%	37,96%	6,77%
1	2	⊥	⊥	4	0,00%	58,33%	13,79%
2	1	⊥	⊥	36	0,00%	75,00%	31,70%
2	1	⊥	⊥	36	5,56%	71,94%	25,80%
2	2	⊥	⊥	8	0,00%	76,25%	39,66%
3	1	⊥	⊥	12	0,00%	91,67%	100,00%
3	1	⊥	⊥	12	33,33%	94,44%	77,01%
3	2	⊥	⊥	4	0,00%	91,67%	100,00%
			⊥	84	0,00%	64,29%	31,32%
			⊥	100	6,00%	63,00%	28,69%

For the KDD dataset, the continuous attributes *age* and *income* are the only potential candidate for a fine-grained generalisation. All the others, including *zip* that belong in each release to at least a domain of level 1, do not give an informed attacker any additional advantage because their hierarchy is based on a simple suppression approach. When comparing the same input parameter sets of Table VII, it again shows that an attacker with internal knowledge has a clear advantage. He suffers from less information loss, achieves a higher success rate, and has more manipulation options available. In general, however, on this dataset, the incurred utility losses are much smaller compared to the Adult dataset. In several settings, it is possible to achieve a hiding success rate of more than 70%, incurring relative utility loss below 15%. On this dataset, also relatively high success rates for obfuscation can be achieved, with around 50% average success rate. The relatively low utility loss can be attributed to often using the continuous attribute *income*, which can be generalised three times according to its underlying domain tree.

TABLE VII: Attack and detection phase outcomes; KDD

Input parameter				Modification results			
no of attrib.	levels	mixed	knowledge	options	avg. obfuscation success	avg. hiding success	avg. rel. utility loss
1	1	⊥	⊥	22	31,82%	72,27%	14,90%
1	1	⊥	⊥	22	40,91%	74,70%	11,74%
1	2	⊥	⊥	6	16,67%	73,89%	13,20%
1	2	⊥	⊥	9	22,22%	73,33%	12,69%
1	3	⊥	⊥	3	0,00%	65,56%	27,27%
2	1	⊥	⊥	14	71,43%	96,43%	52,07%
2	1	⊥	⊥	14	85,71%	98,21%	38,41%
2	2	⊥	⊥	1	0,00%	87,50%	100,00%
2	2	⊥	⊥	2	50,00%	93,75%	74,56%
2	2	⊥	⊥	6	33,33%	91,67%	77,04%
2	2	⊥	⊥	10	80,00%	97,50%	50,06%
2	3	⊥	⊥	4	50,00%	93,75%	68,19%
3	1	⊥	⊥	2	100,00%	100,00%	74,75%
3	1	⊥	⊥	2	100,00%	100,00%	74,75%
3	2	⊥	⊥	1	100,00%	100,00%	100,00%
3	2	⊥	⊥	1	100,00%	100,00%	100,00%
			⊥	52	44,23%	83,90%	37,46%
			⊥	67	55,22%	85,26%	32,30%

In summary, manipulations resulting in a higher information

loss have a better success rate, as the number of patterns in question increases. The underlying attribute data types have a significant influence on the number of possible attack strategies and their success rate.

3) *Observations and recommendations:* Regarding the *dataset characteristics*, certain properties of the dataset have a major impact on the attacker’s success. A higher number of attributes in conjunction with a large distance between the represented and the maximum possible generalisation level favours an attacker, granting more attack opportunities. Generalisation hierarchies of continuous data types allow for the generation of a high number of fingerprints, with the disadvantage of increasing the risk of a successful attack. Categorical attributes provide the benefit of limiting the number of attacks, as well as increasing data loss induced by the attack. Since their domain structure can be more complex and harder to predict compared to continuous, attackers without background knowledge often have no other option than to apply total suppression.

Data publishers can influence the robustness of fingerprints with the *choice of appropriate fingerprint generation parameters*. The lower the data utility, implying a high degree of generalisation up to suppression, the lower the number of opportunities for manipulation. On the one side, data publishers achieve thereby a limitation of the attack and thus more robust fingerprints, on the other side, the data is less useful for further analysis. Therefore, while choosing the cluster of fingerprints with the high centroid (utility), the data owner must take into account the trade-off with fingerprint robustness.

Finally, the *number of fingerprints* affects their robustness. The smaller the number of potential recipients, the smaller the eligible set for an attacker to hide behind becomes, thus making the attack harder. Of course, there are certain limitations in restricting the number of recipients, and if there are commercial incentives (like selling the data), then there is a desire to maximise the number of recipients, which is then in direct contradiction to the aim of limiting the number of recipients for attack resilience purposes.

D. Collusion attack

In this part, we conduct two case studies on the Adult Census dataset to analyse the collaboration between data recipients. Firstly, we analyse the success rate of the collusion attack for all fingerprints in the cluster, before applying the collusion filter. Secondly, we apply the collusion filter as described in Section III-B and demonstrate the collusion resolution. Lastly, we challenge the strong definition of the upper limit of the number of recipients (equalling the number of QIs in the dataset as discussed in Section III-C).

a) *Collusion analysis:* Multiple data recipients may combine their datasets by matching the data values they have in common, typically the non-QI attributes that remain unchanged in the anonymisation/fingerprinting process, but also on the QIs that may share exact same values due to the same anonymisation levels. We first show how successful the attackers get regardless of the collusion filter. Given that their goal is to reduce their received generalised values (assignment

shown in Table VIII), the success rate denotes the percentage of records recovered to minimal generalisation levels available from collaboration.

TABLE VIII: Recipients and their gen. configurations

Recipient	A	B	C	D	E
Gen. config.	[4,1,2,1,1,1,0,1,0]	[3,1,3,1,1,1,0,1,0]	[3,1,2,2,1,1,0,1,0]	[3,1,2,1,1,1,0,1,1]	[4,1,3,0,1,1,0,1,0]

TABLE IX: Collusion attack.

*including records that have 2 potential matches among collaborators' datasets

**including records that have 3 potential matches among collaborators

Collaborators	Success rate	Success rate*	Success rate**
A & B	1478 (4,90%)	2150 (7,12%)	2585 (8,57%)
A & C	1782 (5,91%)	2626 (8,71%)	3190 (10,58%)
A & D	1752 (5,81%)	2556 (8,47%)	3099 (10,27%)
A & E	1640 (5,44%)	2398 (7,95%)	2935 (9,73%)
B & C	2060 (6,83%)	3036 (10,07%)	3762 (12,47%)
B & D	1982 (6,57%)	2930 (9,71%)	3533 (11,71%)
B & E	1378 (4,57%)	1996 (6,62%)	2404 (7,97%)
C & D	2437 (8,08%)	3571 (11,84%)	4402 (14,59%)
C & E	1378 (4,57%)	1996 (6,62%)	2404 (7,97%)
D & E	1256 (4,16%)	1842 (6,11%)	2193 (7,27%)

The ability of the collaborators to successfully match their datasets depends majorly on the amount and distribution of non-QIs. As confirmed by our analysis presented in Table IX, only in average about 5,69% of records are successfully transformed by two recipients colluding. The remainder of the datasets leaves uncertainty for the matches across the two datasets, we show results for up to 2 and up to 3 potential matches per record.

After applying the collusion filter to the potential cluster of recipients from Table VIII as described in Section III-B, the resulting group of anonymised datasets for distribution is [A,B,E] and the collusion resolution is possible by using a single record. This is demonstrated in Table X. Exemplified is the case of A and E in collusion, where we can single out one of 1640 of the samples (cf. Table IX) transformed into the least generalised values available to the union of the colluders, i.e. [4,1,2,0,1,1,0,1,0]. Since only the recipient A has the level of detail for the "education" attribute (generalisation level = 2) and only E has the level of detail for the "marital-status" (generalisation level = 0) among all collusion-free recipients, we infer that indeed the recipients A and E had to be in the collaboration to obtain the presented data sample. Without the loss of generality, this applies to any combination or number of collaborators from [A,B,E].

On the other hand, we can demonstrate the unsuccessful collusion resolution on the result of the collaboration of recipients C and D, those excluded by the collusion filter step, i.e. if we used all five datasets from Table VIII. This collaboration results in 2437 records (cf. Table IX) that have minimal attribute generalisations available to the union of the recipients; i.e. the configuration [3,1,2,1,1,1,0,1,0]. From this configuration, it is not possible to detect the correct colluders since there is no unique level of detail that only one recipient received. For instance, the generalisation level = 1 of "marital-status" (4th QI) could have originated from recipients A, B or

C. This applies to any collusion which contains at least one recipient outside of [A,B,E].

b) *Relaxing the upper limit on the number of recipients:*

A rather low percentage of the recovered records indicates that the collaborators do not necessarily benefit greatly from joining the datasets. The threat can even be diminished by suppressing selected records. This method can be found useful in scenarios where a larger number of recipients is a critical requirement. For our example from above we found that similar sets of individual records contribute to success rates for all the collaborations in Table IX. By suppressing these records, the resulting collusion attack success rates considerably decrease. We find that suppressing $\approx 10\%$ of data records leads to cutting the success rates by more than half, suppressing $\approx 15\%$ decreases the success below 1% for most of the collusions, and finally, the perfect collusion can be evaded by suppressing $\approx 23\%$ of the records, per results shown in Table XI. Note that the collusion cannot be resolved in cases where non-collusion-free recipients collaborate, however, their chances of finding the matching records are so small that one could rightfully claim that collusion is no longer a threat.

c) *Discussion:* The colluding attackers might have dual intentions; (i) reducing the generalisation levels, hence improving data utility, or (ii) re-identification of the data records. The first one cannot be prevented, but we discuss how the involved parties can be detected from such collaborations. The re-identification goal is prevented intrinsically by the collusion filter such that any reduction of generalisation levels still results in a dataset that satisfies k -anonymity property. It is important to note that this cannot be guaranteed when relaxing the collusion-free requirement. To make sure that the re-identification is not possible for the collaborations in the relaxed scenario, one either needs to ensure that the combination of all minimal provided attribute generalisations still satisfies k -anonymity, or that the attacks success rates are decreased to 0, hence preventing the collusion altogether.

V. CONCLUSIONS

In this paper we analysed the approach of achieving anonymity and unauthorised usage verification of the datasets via a unified process, combining thus two common requirements for sharing private and sensitive datasets. The approach uses the k -anonymity generalisation patterns as the fingerprint information, according to which the owner can trace the data copy to its original recipient. The analysis of our implemented approach shows that the data owner needs to solve certain trade-offs in order to obtain robust fingerprints and ensure high data quality simultaneously for all the recipients. The utility of the fingerprints is preserved better in the settings where attributes have high granularity of potential generalisations (such as numeric attributes), however, the robustness is achieved better in settings including attributes with a rather low granularity of potential generalisations, thus limiting any potential attacks. Additionally, limiting the amount of distributed datasets contributes to better robustness. This is especially the case for collusion attacks where the attackers combine

TABLE X: Demonstration of collusion resolution with and without collusion filter. One record of the Adult Census dataset obtain from the collusion of recipients A and B represents the unsuccessful collusion resolution, and C and D and the corresponding generalisation configurations.

	age	workclass	education	marital-status	occupation	relationship	race	sex	capital-gain	capital-loss	hours-per-week	native-country	salary-class
A&E	[57, 97]	*	higher edu.	Widowed	*	not-in-family	*	Male	0	0	6	*	> 50k
	4	1	2	0	1	-	1	0	-	-	-	1	0
C&D	[77, 97]	*	higher edu.	other	*	not-in-family	*	Male	0	0	50	*	> 50k
	3	1	2	1	1	-	1	0	-	-	-	1	0

TABLE XI: Collusion attack after suppressing selected records.

Collaborators	suppressed records					
	0%	1,19%	1,33%	10,08%	15,02%	22,73%
A & C	5,91%	4,98%	4,97%	2,10%	0,39%	0
A & D	5,81%	4,91%	4,91%	2,04%	0,30%	0
B & C	6,83%	5,88%	5,88%	3,00%	0,47%	0
B & D	6,57%	5,68%	5,68%	2,78%	0,31%	0
C & D	8,08%	7,69%	7,55%	4,14%	2,93%	0
C & E	4,57%	4,22%	4,22%	0,91%	0,16%	0
D & E	4,16%	3,84%	3,70%	3,68%	3,26%	0

information from their datasets to either improve the utility of their dataset or attempt the re-identification of the records. We show that the method is resilient against those attempts. We further show that even for a larger number of recipients, the collusion threat can be diminished by suppressing records that allow collaboration.

The future work includes comparing different methods that ensure privacy and unauthorised usage tracing including those that incorporate some stronger privacy definitions and those applying these requirements sequentially, i.e. classical fingerprint on top of the k -anonymous data.

ACKNOWLEDGEMENTS

c This work received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No 826078 (FeatureCloud). This publication reflects only the authors’ view and the European Commission is not responsible for any use that may be made of the information it contains. SBA Research (SBA-K1) is a COMET Center within the COMET - Competence Centers for Excellent Technologies Programme and funded by BMK, BMAW, and the federal state of Vienna. The COMET Programme is managed by FFG.

REFERENCES

- [1] P. Kieseberg, S. Schrittwieser, M. Mulazzani, I. Echizen, and E. Weippl, “An algorithm for collusion-resistant anonymization and fingerprinting of sensitive microdata,” *Electronic Markets*, vol. 24, pp. 113–124, June 2014.
- [2] P. Samarati and L. Sweeney, “Protecting privacy when disclosing information: k -anonymity and its enforcement through generalization and suppression,” tech. rep., 1998. Carnegie Mellon University.
- [3] A. Meyerson and R. Williams, “On the complexity of optimal K -anonymity,” in *Proceedings of the twenty-third ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems - PODS ’04*, (Paris, France), p. 223, ACM Press, 2004.
- [4] P. Samarati, “Protecting respondents identities in microdata release,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 13, pp. 1010–1027, Dec. 2001.
- [5] F. Kohlmayer, F. Prasser, C. Eckert, A. Kemper, and K. A. Kuhn, “Flash: Efficient, Stable and Optimal K -Anonymity,” in *International Conference on Privacy, Security, Risk and Trust and 2012 International Conference on Social Computing, PASSAT*, (Amsterdam, Netherlands), pp. 708–717, IEEE, Sept. 2012.
- [6] K. E. Emam, F. K. Dankar, R. Issa, E. Jonker, D. Amyot, E. Cogo, J.-P. Corriveau, M. Walker, S. Chowdhury, R. Vaillancourt, T. Roffey, and J. Bottomley, “A globally optimal k -anonymity method for the de-identification of health data,” vol. 16, pp. 670–682, 2009.
- [7] F. Prasser, J. Eicher, H. Spengler, R. Bild, and K. A. Kuhn, “Flexible data anonymization using ARX—Current status and challenges ahead,” *Software: Practice and Experience*, vol. 50, pp. 1277–1304, July 2020.
- [8] K. E. Emam and B. Malin, “Appendix b: Concepts and methods for de-identifying clinical trial data,” 2015.
- [9] F. A. Petitcolas and S. Katzenbeisser, *Information Hiding Techniques for Steganography and Digital Watermarking (Artech House Computer Security Series)*. Artech House, 2000.
- [10] I. J. Cox, M. L. Miller, J. A. Bloom, and C. Honsinger, *Digital watermarking*, vol. 53. Springer, 2002.
- [11] R. Agrawal, P. J. Haas, and J. Kiernan, “Watermarking relational data: framework, algorithms and analysis,” *The VLDB Journal—The International Journal on Very Large Data Bases*, vol. 12, no. 2, pp. 157–169, 2003.
- [12] Y. Li, V. Swarup, and S. Jajodia, “Fingerprinting relational databases: Schemes and specialties,” *IEEE Transactions on Dependable and Secure Computing*, vol. 2, no. 1, pp. 34–45, 2005.
- [13] J. Lafaye, D. Gross-Amblard, C. Constantin, and M. Guerrouani, “Watermill: An optimized fingerprinting system for databases under constraints,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 20, no. 4, pp. 532–546, 2008.
- [14] T. Sarcevic and R. Mayer, “A correlation-preserving fingerprinting technique for categorical data in relational databases,” in *ICT Systems Security and Privacy Protection: 35th IFIP TC 11 International Conference, SEC 2020, Maribor, Slovenia, September 21–23, 2020, Proceedings 35*, pp. 401–415, Springer, 2020.
- [15] S. Schrittwieser, P. Kieseberg, I. Echizen, S. Wohlgemuth, N. Sonehara, and E. Weippl, “An Algorithm for k -Anonymity-Based Fingerprinting,” in *Digital Forensics and Watermarking*, pp. 439–452, Berlin, Heidelberg: Springer Berlin Heidelberg, 2012.
- [16] P. Kieseberg, S. Schrittwieser, M. Mulazzani, I. Echizen, and E. Weippl, “An algorithm for collusion-resistant anonymization and fingerprinting of sensitive microdata,” *Electronic Markets*, vol. 24, pp. 113–124, Jun 2014.
- [17] E. Bertino, Beng Chin Ooi, Yanjiang Yang, and R. Deng, “Privacy and Ownership Preserving of Outsourced Medical Data,” in *International Conference on Data Engineering, ICDE*, (Tokyo, Japan), pp. 521–532, IEEE, 2005.
- [18] S. Gams, J. Lohve, and J.-M. Robert, “Entwining Sanitization and Personalization on Databases,” in *Asia Conference on Computer and Communications Security, ASIA CCS*, (Incheon Republic of Korea), pp. 207–219, ACM, May 2018.
- [19] T. Ji, E. Ayday, E. Yilmaz, and P. Li, “Privacy-Preserving Database Fingerprinting,” Mar. 2022. arXiv:2109.02768 [cs].
- [20] J. Yu, S. Yuan, Y. Yuan, Y. Li, and C. Chen, “A k -Anonymity-Based Robust Watermarking Scheme for Relational Database,” in *Science of Cyber Security*, vol. 13580, pp. 557–573, Cham: Springer International Publishing, 2022. Series Title: Lecture Notes in Computer Science.
- [21] M. Kamran and M. Farooq, “A Comprehensive Survey of Watermarking Relational Databases Research,” Jan. 2018. Number: arXiv:1801.08271 arXiv:1801.08271 [cs].
- [22] Y. Li, V. Swarup, and S. Jajodia, “Fingerprinting relational databases: schemes and specialties,” vol. 2, pp. 34–45, 2005.

- [23] F. Kohlmayer, F. Prasser, C. Eckert, A. Kemper, and K. A. Kuhn, "Flash: Efficient, stable and optimal k-anonymity," 2012.
- [24] L. Sweeney, "Achieving K-anonymity Privacy Protection Using Generalization and Suppression," *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, vol. 10, pp. 571–588, Oct. 2002.
- [25] V. S. Iyengar, "Transforming data to satisfy privacy constraints," in *Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 279–288, 2002.
- [26] K. LeFevre, D. DeWitt, and R. Ramakrishnan, "Mondrian Multidimensional K-Anonymity," in *International Conference on Data Engineering, ICDE*, (Atlanta, GA, USA), pp. 25–25, IEEE, 2006.
- [27] M. E. Nergiz and C. Clifton, "Thoughts on k-anonymization," *Data & Knowledge Engineering*, vol. 63, no. 3, pp. 622–645, 2007.