

# The Fast Maximum Distance to Average Vector (F-MDAV): An Algorithm for $k$ -Anonymous Microaggregation in Big Data

Ana Rodríguez-Hoyos<sup>a,b</sup>, José Estrada-Jiménez<sup>a,b</sup>, David Rebollo-Monedero<sup>b</sup>, Ahmad Mohamad Mezher<sup>b</sup>,  
Javier Parra-Arnau<sup>c</sup>, Jordi Forné<sup>b,\*</sup>

<sup>a</sup>*Departamento de Electrónica, Telecomunicaciones y Redes de Información, Escuela Politécnica Nacional (EPN), Ladrón de Guevara, E11-253 Quito, Ecuador*

<sup>b</sup>*Department of Telematics Engineering, Universitat Politècnica de Catalunya (UPC), C. Jordi Girona 1-3, E-08034 Barcelona, Spain*

<sup>c</sup>*Department of Computer Engineering and Mathematics, Universitat Rovira i Virgili, Av. Països Catalans 26, E-43007 Tarragona, Catalonia.*

---

## Abstract

The massive exploitation of tons of data is currently guiding critical decisions in domains such as economics or health. But serious privacy risks arise since personal data is commonly involved.  $k$ -Anonymous microaggregation is a well-known method that guarantees individuals' privacy while preserving much of data utility. Unfortunately, methods like this are computationally expensive in big data settings, whereas the application domain of data might require an immediate response to make "life or death" decisions.

Accordingly, this paper proposes five strategies to simplify the internal operations (such as distance calculations and element sorting) of the maximum distance to average vector method, the de facto microaggregation standard. For the sake of its usability in large-scale databases, they, e.g., reduce the number of operations necessary to compute distances from  $3m$  to  $2m$ , where  $m$  is the number of attributes of the data set. Also, the complexity of sorting operations gets reduced from  $\mathcal{O}(n \log n)$  to  $\mathcal{O}(n)$  where  $n$  is the number of records. Through extensive experimentation over multiple data sets, we show that the new algorithm gets significantly faster. Interestingly, the speedup factor by each technique is not greater than 2, but the multiplicative effect of combining them all turns the algorithm four times faster than the original microaggregation mechanism. This remarkable speedup factor is achieved, literally, with no additional cost in terms of data utility, i.e., it does not incur greater information loss.

**Keywords:** MDAV,  $k$ -anonymous microaggregation, speed-up, data privacy, big data

---

## 1. Introduction

Big data is bringing new, unprecedented business opportunities to companies around the world. Currently, it is possible to collect and process vast amounts of information from which more, new, better and varied customer knowledge is mined. As a result, better decisions can be made in sectors like health care, banking,  
5 marketing and transportation [1, 2, 3].

Despite these benefits, within such an abundance of data, it is common to find personal sensitive information, which poses serious privacy risks. First, in the name of this data revolution, information is more prone to be openly published or shared with untrusted third parties. Also, although identifiers are typically suppressed, other demographic attributes, when combined, can be used to reidentify individuals [4, 5, 6]. Thus,

---

\*Corresponding author

Email addresses: ana.rodriguez@epn.edu.ec (Ana Rodríguez-Hoyos), jose.estrada@epn.edu.ec (José Estrada-Jiménez), david.rebollo@upc.edu (David Rebollo-Monedero), ahmad.mohamad@upc.edu (Ahmad Mohamad Mezher), javier.parra@urv.cat (Javier Parra-Arnau), jforne@upc.edu (Jordi Forné)

10 sensitive attributes might be easily linked to the subjects to whom the disclosed information corresponds, which might lead to discrimination, retaliation, and blackmail [7].

*Statistical disclosure control* (SDC) aims at addressing these privacy issues in the special case of microdata files. Microdata are database tables whose records carry data concerning individual subjects. The typical scenario in microdata SDC is a data curator holding the original data set and perturbing the so-called *quasi-identifier* attributes (i.e., attributes that, in combination, may be linked with external information to reidentify individuals in the data set). The goal is to keep disclosure risk as low as possible, while ensuring that only useful statistics or trends are learned by the recipients of data. One of the most common strategies to keep this risk under control is the “privacy first” approach. Here, the data curator enforces a privacy model, which usually depends on a privacy parameter, to ensure an upper bound on the re-identification risk. Some of the best-known privacy models comprise  $k$ -anonymity [8, 4] and  $\epsilon$ -differential privacy [9].

20 Privacy models rely on a variety of anonymization mechanisms, the common denominator binding all them is *data perturbation*. Essentially, all such mechanisms modify the original data set to guarantee the chosen privacy model, inevitably at the cost of some loss in data utility [10]. Evidently, a balance between privacy and utility should be found so that the protected data are useful in real practice, that is, they approximate well the original data. However, in a big data domain, privacy protection also requires mechanisms to execute in a reasonable amount of time, despite the size of the data.

Examples of anonymization mechanisms include microaggregation, suppression, generalization and noise addition. Among them,  $k$ -anonymous microaggregation is a high-utility approach that obfuscates demographic attributes. By carefully aggregating these attributes, a minimum level of distortion must be applied to data. Unfortunately, current microaggregation algorithms entail a very high computational cost when anonymizing big data [11]. Thus, since utility extraction from big data is a priority, and already time consuming, privacy protection might be easily neglected. That is why some works are starting to propose strategies to reduce the running time of privacy enhancing mechanisms while preserving the utility of data.

### 1.1. Contributions of this work

35 In this work, we propose and evaluate various improvements to a widely-known microaggregation algorithm for  $k$ -anonymity protection, the maximum distance to average vector (MDAV). The fundamental aim of such improvements is to effectively reduce the computational complexity of MDAV. Without yielding any additional loss in data utility, we obtain remarkable reductions in running time by diminishing the number of operations necessary to aggregate data, while protecting the privacy of the individuals involved. The proposed computational improvements aim at adapting high-utility SDC to large-scale data sets with numerical demographic attributes. This approach is interesting since the reduction of the computational cost of privacy protection algorithms may encourage its implementation, especially when computation usually entails important economic costs for companies exploiting big data. More concretely, our contributions are the following:

- 45 • We propose five strategies for efficient computation that significantly accelerate the running time of  $k$ -anonymous microaggregation for multivariate data, which is ideal for big data scenarios.
- Our improvement strategies are based on algebraic and algorithmic operations. First, we rely on algebraically optimized computation of Euclidean distances. Then, algorithmically, we resort to the precomputation and reuse of distances. Also, we use partial selection of a few  $(k - 1)$  candidates among a list of nearest and furthest data points without requiring total sorting. When calculating centroids, we eliminate redundant operations of sum, thus considerably reducing the resulting complexity. Finally, we have numerical data represented with single-precision floating numbers in lieu of double-precision for faster computation, with utterly negligible impact in accuracy for the application at hand. Said strategies do not alter the functionality of MDAV but the resulting improvements of execution time do make this algorithm more usable in big-data contexts.
- 55 • We conduct an extensive, thorough performance evaluation of our modified MDAV. Through systematic experimentation on three standardized data sets, we find the magnitude of the computational improvements, quantified as the resulting speed-up.

- The overall speed of our privacy algorithm is increased by a factor of 400 percent in relation to its original version. In terms of operativity, this would have a significant impact on the economic costs of a company dedicated to exploit big data privately.
- The proposed enhancements do not incur additional information loss. Namely, there is no increase in data distortion. We modify the most complex parts of MDAV so that less computations are needed to obtain the same microcells. As a result, our algorithm generates the same anonymized data that the original MDAV would produce.

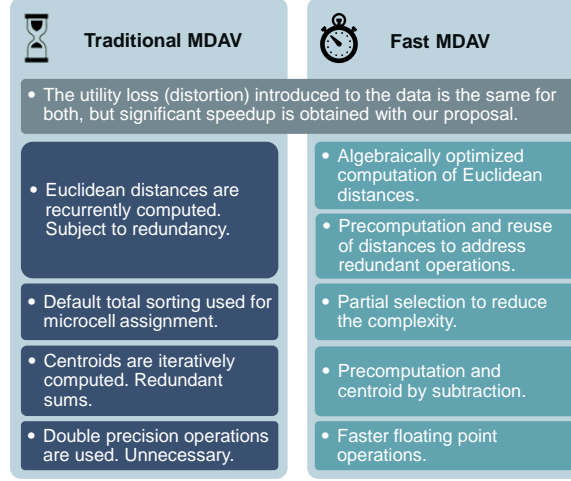


Figure 1: Summary of the five computational improvements of F-MDAV with respect to traditional MDAV.

Our contributions are summarized in Fig. 1.

### 1.2. Related work and impact of this approach

The novelty of our contribution lies mainly in our effort to analyze the microaggregation algorithm, i.e., finding the components subject to be accelerated, and devise the mechanisms and algebraic properties that could implement such improvements. Some of said operations could have different applications and part of our contribution is, in fact, tailoring them to the specific case of  $k$ -anonymous microaggregation.

Recent works have shown to follow similar approaches, harnessing known mechanisms from different domains (e.g., machine learning), aiming to increase the efficiency of privacy protection algorithms, not only in terms of runtime [12, 13, 14, 15], but also in terms of resulting data utility. For instance, [16] developed an efficient clustering mechanism to deal with large databases while preserving the data utility through a partitioning method of a modified minimum spanning tree. In the same line, [17] designed an efficient and effective microaggregation mechanism based on calculating the distance among records as the mutual information (entropy) among them. Finally, [18] introduced fast data-oriented microaggregation (FDM), a method capable of getting multiple protected versions of a large data set (for different values of  $k$ ) in a single load.

This may suggest that this is a topic of interest. Although utility in this context might have received more attention from researchers, the computational time of privacy algorithms has to do with its *usability* and it is, from our perspective, at least as important.

Although data is the new oil in the big data era, applications of big data are currently possible just because the algorithms that process it can be executed much faster than in the past. However, the execution time is still a bottleneck for some highly demanding applications, which does not favor the implementation of further processing privacy routines. Consequently, we could say that accelerating the execution of privacy protection mechanisms is a fundamental approach to encourage its adoption in the era of big data.

Certainly, besides the runtime acceleration by our method, the null degradation of data utility is part of the added value of our approach. On the other hand, the works previously cited offer reduction of runtime at a cost in data distortion. Ours does not imply any distortion to microaggregated data.

Another not negligible detail about our approach is that it has a multiplicative effect that can be combined with that of other proposals. Namely, if ours (speedup factor,  $4\times$ ) is applied along with another strategy that multiplies the speed of MDAV by two, in practice, the resulting speedup would be  $8\times$ . Naturally, said effect enables an interesting opportunity for capitalizing on several of the efficiency approaches proposed for microaggregation and, particularly for MDAV. Interestingly, these computational improvements can also be combined with the functional (data utility) improvements introduced by other works where microaggregation is involved [19, 20, 21]

Another rationale regarding the novelty of our proposal follows. Due to the massification of Internet access, the vast amounts of data containing personal information may grow and change very dynamically, commonly feeding online services. Then, microaggregation, in this context, is likely to be implemented as an ongoing process, running as fast as possible, rather than as a static one-time job. In fact, if data is sufficiently vast, microaggregating it once could be unfeasible in practice for some, e.g., real-time, applications due to the quadratic complexity of MDAV, so optimizing its running time in the big data era seems mandatory.

Inevitably, all these details derived from our contribution have important economic implications. Despite the obstacles behind implementing privacy, some big tech companies are turning their privacy stance into a huge competitive advantage. Thus, the companies that best adapt their operation to privacy requirements (preserving data utility and algorithm usability) will be in better position to exploit such advantage. Interestingly, in this context, increasing the efficiency of privacy protection mechanisms (e.g., reducing their runtime) could become a powerful value generator.

The remainder of this paper is organized as follows. §2 reviews the background on  $k$ -anonymous microaggregation, particularly focusing on MDAV, and reviews the state of the art in microaggregation algorithms for SDC. Next, §3 describes the adaptations we apply to MDAV in order to reduce its running time. §4 shows and analyzes the experimental results of our improvements for a variety of data sets. Lastly, conclusions are drawn in §5.

Identifiers		Quasi-Identifiers		Confidential Attributes		Microaggregated Quasi-Identifiers			Confidential Attributes	
Name	Age	Marital Status	ZIP Code	Annual Salary	Sexual Orientation	Age	Marital Status	ZIP Code	Annual Salary	Sexual Orientation
Alice Adams	32	1	94024	45 K	♀♂	33	0.33	94***	45 K	♀♂
Bob Brown	34	0	94305	35 K	♂♂	33	0.33	94***	35 K	♂♂
Chloe Carter	33	0	94024	15 K	♀♀	33	0.33	94***	15 K	♀♀
Dave Diaz	43	0	90210	55 K	♀♀♂	45	0.67	9021*	55 K	♀♀♂
Eve Ellis	47	1	90210	70 K	♀♂	45	0.67	9021*	70 K	♀♂
Frank Fisher	45	1	90213	60 K	♀♀♂	45	0.67	9021*	60 K	♀♀♂

$k$ -Anonymized Records

Figure 2: Example of  $k$ -anonymous microaggregation of published data with  $k = 3$ , showing probably sensitive information (annual salary and sexual orientation) as confidential attributes, in relation to demographic variables (age, marital status and ZIP code) as quasi-identifiers. The original table codifies the marital-status attribute with 1 and 0 to refer to “married” and “single” respectively. When microaggregation is applied, the marital-status attribute is anonymized by replacing the values in a cell by the proportion of married people in that cell, as MDAV does. For example, in the first cell (first three instances) of the microaggregated table, the marital-status attribute corresponds to  $1/3$ , since, in the original table, one of the three people in that cell is married.

## 2. Background and state of the art

Currently, data is more prone to be shared with external parties or even openly, e.g., for research purposes, in order to better exploit its utility. Thus, malicious observers may illegitimately take advantage, since

sensitive information might still be encoded within released data, although previously some information had been suppressed. Unfortunately, conventional privacy services against unintended observers and based on cryptography, fail to address the practical dilemma when the intended recipient of the information is not fully trusted.

As a first approach to protect the anonymity of individuals, it is common to just eliminate their identifiers (e.g., full names or social security numbers) before a data set is published or shared. However, this practice was proved to be insufficient in [4], where it was shown that 87% of the population in the United States could be unequivocally identified solely on the basis of the triple consisting of their date of birth, gender and 5-digit ZIP code, according to 1990 census data. Due to the discriminative potential of a few combined demographic attributes, more sophisticated approaches have been proposed to obscure the identity of the respondents appearing in the released data set.

### 2.1. Background on $k$ -Anonymous Microaggregation

Privacy protection techniques usually focus on databases carrying information concerning individual respondents (from a survey or a census). Said databases, also known as *microdata sets*, contain a set of attributes that may be classified into identifiers, quasi-identifiers and confidential attributes. Firstly, *identifiers*, such as full names or medical record numbers, can single out individuals from a data set and are commonly removed to preserve the anonymity of respondents. Secondly, *quasi-identifiers* or key attributes may include age, gender, address, or physical features, which combined and linked with other external information can be used to reidentify respondents. Finally, a data set may contain *confidential attributes* with sensitive information on the respondents, such as salary, health condition, and religion.

In Fig. 2, we illustrate how a perturbed, and thus more private, version of a data set is obtained to be published instead of the original one. In the figure, the original data set combines attributes common in census and medical surveys. It has three quasi-identifiers; age, marital status and ZIP code; and two confidential attributes: annual salary and sexual orientation. The figure at hand shows how, in order to preserve the privacy of respondents, perturbation is applied to quasi-identifiers. This technique, called microaggregation, is applied to enforce  $k$ -anonymity [4], a privacy model that guarantees that each tuple of key-attribute values is identically shared by at least  $k$  records in a data set. Rather than making the original table available, a perturbed version is published where aggregated records of quasi-identifying values are replaced by a common representative tuple. The result is a microaggregated data set that may prevent reidentification attacks.

As illustrated in Fig. 3, if tuples of key attributes in a data set could be represented as points in the Euclidean space,  $k$ -anonymous microaggregation would consist in partitioning these points in cells of size  $k$ , and quantizing each cell and its elements with a representative point. Perturbed key attributes would be characterized by the set of representative points. The de facto standard for numerical microaggregation is the MDAV algorithm. It was proposed in [22] as a practical evolution of a multivariate fixed-size microaggregation method and conceived in [23]. MDAV is one of the best-known and most widely used fixed-size microaggregation algorithms for numerical data due to its high performance both in terms of running time and the resulting (even empirical) utility of anonymized data [24, 25].

We provide, in Algorithm 1, a simplified version of that given in [26] and termed “MDAV generic”.

From the pseudocode of Algorithm 1, we can infer that the main operations of MDAV involve the calculation of centroids (averages), the calculation of distances, and numerical sorting. Mainly due to sorting operations, MDAV shows a quadratic computational complexity; more specifically, its running time grows as  $\frac{n^2}{k}$ , where  $n$  is the number of records of the data set. Thus, in general, if the number of records is very large, the microaggregation process would take too long and might not be feasible.

### 2.2. State of the Art of $k$ -Anonymous Microaggregation and Sorting Operations

With microaggregation, as with generalization and suppression, distortion is applied to key attributes to satisfy the  $k$ -anonymity privacy model [8, 4]. This model guarantees that each individual’s information contained in a released data set cannot be distinguished from that of at least other  $k - 1$  individuals whose

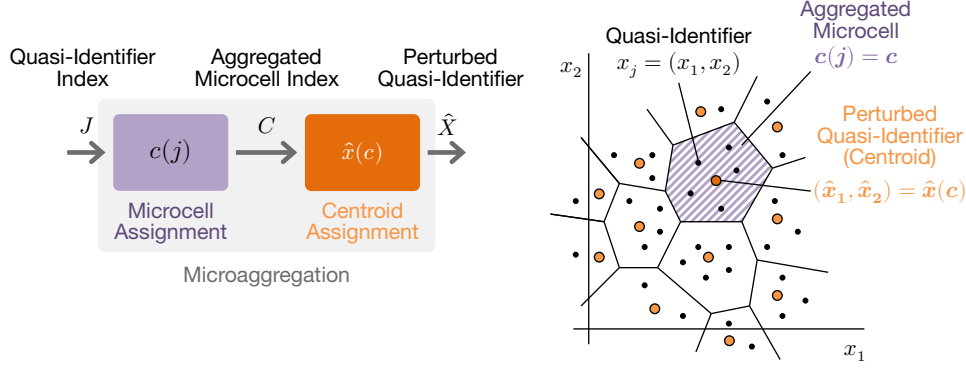


Figure 3:  $k$ -Anonymous microaggregation as a minimum-distortion quantizer design problem with a constraint on the size of the quantizer cells [27].

**Algorithm 1** MDAV “generic”, functionally equivalent to Algorithm 5.1 in [26].

---

```

function MDAV
input  $k, (x_j)_{j=1}^n$             $\triangleright$  Anonymity parameter  $k$ , quasi-ID portion  $x_1, \dots, x_n \in \mathbb{R}^m$ 
                                 $\triangleright$  of a dataset of  $n$  records
                                 $\triangleright$  Assignment function from records to microcells  $j \mapsto q(j)$ 
output  $q$ 
1: while  $2k$  points or more in the dataset remain to be assigned to microcells do
2:   find the centroid (average)  $C$  of those remaining points
3:   find the furthest point  $P$  from the centroid  $C$ , and the furthest point  $Q$  from  $P$ 
4:   select and group the  $k-1$  nearest points to  $P$ , along with  $P$  itself, into a microcell, and
   do the same with the  $k-1$  nearest points to  $Q$ 
5:   remove the two microcells just formed from the dataset
6: if there are  $k$  to  $2k-1$  points left then
7:   form a microcell with those and finish
8: else                                $\triangleright$  At most  $k-1$  points left, not enough for a new microcell
9:   adjoin any remaining points to the last microcell            $\triangleright$  Typically nearest microcell

```

---

information also appears in the data set. The original formulation of  $k$ -anonymity as a privacy criterion was modified into the microaggregation-based approach in [28, 23, 26, 29].

Although  $k$ -anonymity is a very popular privacy criterion, it is not flawless. Since the criterion strictly operates with key attributes, the statistical properties of confidential attributes are neglected. In general,  $k$ -anonymity overlooks the knowledge a potential attacker may already have or obtain about the data set, giving rise to similarity, skewness or background-knowledge attacks [30, 31, 32]. Stricter criteria addressing similarity attacks are  $p$ -sensitive [33, 34] and  $l$ -diversity. Other criteria such as  $t$ -closeness [35],  $\delta$ -disclosure [36], and average privacy risk [37, 31] deal with similarity and skewness attacks by imposing further requirements on the semantics and distribution of confidential attributes.

Evidently, the  $k$ -anonymous microaggregation process entails distortion or information loss on the data since the original data is modified. The usual criterion to quantify the distortion of microaggregated data is the mean-squared error (MSE) for numerical key attributes representable as points in the Euclidean space. MSE can be computed as

$$\text{MSE} = \sum_{j=1}^n \|x_j - \hat{x}_j\|^2,$$

where  $n$  is the number of records of the data set,  $m$  is the number of attributes of each record,  $x_j \in \mathbb{R}^m$  is the  $j^{\text{th}}$  record, and  $\hat{x}_j$  is the tuple representative of the  $j^{\text{th}}$  record. While MSE is the general metric of the degradation of data utility after microaggregation, other more empirical utility metrics exist (e.g., accuracy of learning models) but their application may be limited to certain types of data sets. In general, the implementations of microaggregation have been oriented to significantly reduce the inherent distortion due to perturbation [38, 39, 40]; as a result, this privacy mechanism can be considered part of the high-utility



SDC spectrum. However, such benefit in utility commonly derives in more sophisticated and significantly costlier implementations in terms of computational time [11].

At a time like the present in which data is generated and processed at a large scale, these new issues in computational complexity must be considered in order to guarantee that privacy protection algorithms are usable in real practice. Notwithstanding, this new dimension of usability has not been addressed in the realm of SDC. To the best of our knowledge, the only approach in this respect is presented in [41], and here we offer a substantial extension of such work.

From the pseudo code in Algorithm 1, it is straightforward to see that the running time  $t$  of the algorithm grows as  $t(n) = \frac{n^2}{k}$  (in time units relative to the case  $\frac{n^2}{k}$ ) for  $n \gg k$ , i.e., it has a quadratic complexity, which implies a superadditive running time. On the one hand, said complexity could make MDAV difficult to apply on big data contexts since its running time would grow significantly as the square of a large number of records. On the other hand, superadditivity is a very useful characteristic of high-utility microaggregation that opens the door to further strategies to reduce the running time of MDAV, e.g., the “divide and conquer” algorithmic approach. This superadditivity can be particularly exploited in sorting operations.

*Sorting* is repetitively used by MDAV to build  $k$ -anonymous groups, based on the “distance” of a reference tuple to the rest of tuples. Consequently, sorting may represent an important part of the execution time of the microaggregation algorithm.

Interestingly, since MDAV only requires finding the  $k - 1$  nearest tuples (see 1), strict sorting (total sorting) is not necessary, but a more relaxed, less complex, operation called partial sorting that finds closest or furthest points faster.

Sorting problems are implemented through two main efficient algorithms: quicksort [42], for total sorting, and quickselect [43], for partial sorting. While quicksort has, in average, a complexity  $\mathcal{O}(n \log n)$  when sorting  $n$  items, quickselect presents a complexity  $\mathcal{O}(n)$ . Thus, choosing quickselect or any of its improved versions could significantly reduce the running time of MDAV. Other versions of quicksort and quickselect [44, 45] have been introduced to improve their performance both on the average and worst cases, giving rise to very efficient sorting applications.

The use of these algorithms and their adapted versions is pretty wide in computer science. For instance, to implement the C++ `sort` and `nth_element` functions, the C standard library uses, respectively, a refinement of quicksort called introsort [44] and a refinement of quickselect called introselect [44]. The algorithm introsort involves executing quicksort for the average case and switching to heapsort when the worst case is detected. The same combination is done for introselect where, for its worst case, quickselect switches to the Floyd-Rivest algorithm [46, 47, 48].

Although not documented since it is a proprietary application, the `sort` command in Matlab R2017b surely derives from the quicksort algorithm or any of its variants such as introsort. Similarly, it is reasonable to think that the Matlab `mink` command, which selects the ordered  $k$  smallest values from a list, is internally executing introselect. In §4 we give more details on the impact of a more efficient sorting algorithm on the running time of MDAV.

As a final note, the strategies developed in this study could be synergically combined with the proposals presented in [19] and [20] where both functional and even further computational improvements are introduced.

### 3. Computational improvements on MDAV

As described in §2, MDAV (see Algorithm 1) creates partitions or microcells from a data set by aggregating neighboring records. Since MDAV operates with numerical attributes, each record is seen as an  $m$ -dimensional point ( $x_s \in \mathbb{R}^m$ ) in the Euclidean space, being  $m$  the number of attributes of the data set. Note that the microaggregation process iteratively extracts pairs of cells while  $2k$  points or more in the data set remain to be assigned. First, a centroid  $C$  is calculated as the average of the remaining points. Then, from  $C$ , two points  $P$  and  $Q$  are found from the data set, which serve as references to obtain the neighboring points of each of the two new microcells: one formed by  $P$  with its  $k - 1$  nearest points and another by  $Q$  with its  $k - 1$  nearest points.  $P$  is obtained as the furthest point from  $C$  (the maximum distance to average vector) and  $Q$  as the furthest point from  $P$ .

The previous description reveals a set of mathematical operations over the records of the data set. These operations mainly involve centroids calculation, distances calculation, and sorting. Since these operations are used repetitively and executed over a vast number of tuples, there is an interesting chance for improvement in the overall performance of MDAV. Next, we describe in detail the improvements we propose on these operations; Table 1 summarizes the MDAV tasks improved in this work, the respective strategy followed, and gives a brief description of each one. In §4, we show the benefits of these strategies through extensive experimentation.

### 3.1. Algebraic improvement

From line 3 of Algorithm 1, we can devise that much of the MDAV runtime is intended to calculate distances in three moments: when finding the furthest point P from centroid C, when finding the furthest point Q from P, and when obtaining the  $k - 1$  nearest points from P and from Q to build two microcells.

The fact that distances between the (in some cases the same) points of a data set are continuously calculated turns each iteration very redundant. The computation complexity of MDAV, then, derives from such redundancy, which we tackle through this improvement.

Since MDAV considers each record of the data set as an  $m$ -dimensional point in the Euclidean space, a distance  $D_j$  between a reference point  $x_0$  (which is C, P and Q, depending on the moment) and a collection of points  $x_j$  for  $j \in \{1, \dots, n\}$  is calculated as a quadratic Euclidean distance, i.e.,

$$D_j = \|x_j - x_0\|^2.$$

Then, to get  $D_j$ , for each  $m$ -dimensional point  $x_j$ , an element by element subtraction ( $m$  operations) and a square norm ( $2m - 1$  operations) must be calculated. That is, a total of  $3m - 1$  operations for each point of the data set.

To reduce the resulting runtime, we consider finding an analogous expression to calculate  $D_j$  such that less operations are performed. In this case, we harness the polarization identity of the inner product to put the expression of  $D_j$  in terms of the inner product of  $x_j$  and  $x_0$ . So we expand the last expression such that

$$\|x_j - x_0\|^2 = \|x_j\|^2 + \|x_0\|^2 - 2\langle x_j, x_0 \rangle.$$

If both sides of the last equation are subtracted  $\|x_0\|^2$  and multiplied by  $1/2$ , we obtain

$$\frac{1}{2} \left( \|x_j - x_0\|^2 - \|x_0\|^2 \right) = \frac{1}{2} \|x_j\|^2 - \langle x_j, x_0 \rangle.$$

Although the expressions on both sides of the equation no longer represent the real value of  $D_j$ , they are a metric still useful to compare distances since they were summed and multiplied by a constant. Thus, when the calculation of distances in MDAV is used to determine the furthest point from a fixed point  $x_0$ , we can safely use the right part of the last expression for comparison issues.

Conveniently, for each compared point  $x_s$ , the value of  $\frac{1}{2}\|x_j\|^2$  can be precomputed once before MDAV is initiated, out of the redundant process, and avoiding significant recalculation in every iteration of MDAV. Thus, in this case, the distance comparison is reduced to the calculation of the inner product  $\langle x_j, x_0 \rangle$ , which consists in an element by element multiplication, i.e.,  $m$  operations, and a sum of the resulting  $m$  terms, i.e.,  $m - 1$  operations, for each point  $x_j$ . For this representation of distances, we have a grand total of  $2m - 1$  operations.

By analytically operating on an expression, we get less operations than the original expression of distance  $D_s$ . More precisely, the number of operations is reduced from  $3m - 1$  to  $2m - 1$  for each distance calculation where  $m$  is the number of quasi-identifiers of each record in the data set.

Not only the number of operations is reduced, but they are algebraic in nature, and that is something for which much of the current code is optimized (e.g., in Matlab or the C standard library). Consequently, it is reasonable that Matlab uses vectorized code or more efficient CPU instructions, that is, advanced vector instructions (AVX) through the Intel math kernel library (MKL). In fact, there are instructions that compute an accumulated sum and a product, designed for the efficient computation of vector and matrix products,



MDAV	Fast MDAV	
$\underbrace{\ x_j - x_0\ ^2}_{3m - 1 \text{ operations}}$	$\underbrace{\frac{1}{2}\ x_j\ ^2}_{\text{precomputed before iterating}}$	$\underbrace{\langle x_j, x_0 \rangle}_{2m - 1 \text{ operations}}$
		$m \text{ if FMA is employed}$

Figure 4: Representation of the distance calculation performed for each  $m$ -dimensional point  $x_j$  when  $k$ -anonymous microcells are built. We can see that our approach Fast MDAV is able to reduce the number of operations from  $3m - 1$  to  $2m - 1$  for each of these  $n$  records. Furthermore, since the inner product  $\langle x_j, x_0 \rangle$  is subject to optimization if FMA is used, the number of operations could be even reduced to  $m$ .

called multiply-accumulate operations and fused multiply-add (FMA), which are included in certain Intel processors (e.g., Haskell). Interestingly, if FMA were implemented, our proposal would lead to reduce the number of operations here analyzed to  $m$ .

In Fig. 4, we summarize the analysis carried out in the last paragraphs. Finally, this improvement lends itself to the implementation of MDAV in graphics processing units (GPU).

### 3.2. Distance reuse

The high utility offered by MDAV comes from smartly grouping the closest points into microcells. As already pointed out, this process certainly involves several steps where distances need to be calculated. In this section we concentrate on steps 3 and 4 of Algorithm 1 with the aim of reducing the runtime of MDAV.

We can see that, when aggregating a microcell, given a reference point P, two distance-calculation operations are performed. First, to find the  $k - 1$  nearest points to P necessary to form a microcell, the distances from all points to P have to be calculated previously. Afterwards, the furthest point Q from P is needed to serve as a reference for building a new microcell; this also entails calculating distances to P.

Evidently, for both steps, the calculations of distances to P can be calculated once and reused. Here, our proposal is using the distances from every point  $x_s$  to P both to find the points nearest to P and to find the furthest point from P.

### 3.3. Partial sorting

Sorting is another time-consuming operation within the original version of MDAV. It is recurrently implemented, e.g., to find the points closest to a given centroid in order to establish the most appropriate members for each microaggregated cell, as posed in line 4 of Algorithm 1.

Finding the points closest to a centroid C implies sorting all the distances from those points to C upwards (total sorting) and then getting the first ones, i.e., the shortest ones. Interestingly, in the case of MDAV, such total sorting is not necessary because only the  $k - 1$  nearest points are required. In fact, their corresponding distances to the centroid do not even need to be ordered. Finding the  $k$  smallest elements implies a more relaxed sorting approach called *partial sorting* or partial selection.

In computer science, total sorting is extensively implemented through the quicksort algorithm [42] while partial sorting through the quickselect algorithm [43]. The computational complexity of the sorting task when using quickselect may be significantly reduced, since it finds the  $k$ th smallest number in an unordered list, which does not require total order. This relaxation makes quickselect's problem a much easier one.

As expected, our proposal in this paper is to resort to the use of an implementation of partial sorting (e.g., quickselect) in MDAV when the phase of microcell assignment is performed.

Quickselect is a selection algorithm by which a single element, the  $k$ th smallest, is found from a list. Its approach starts by randomly selecting a pivot element that will partition the elements in two; the elements smaller than the pivot on the left and the larger ones on the right. Then, this same approach is recursively implemented only into the side where the element being searched lies up until a single element is obtained. On the other hand, total sorting implemented through quicksort applies the aforementioned approach on



Table 1: Summary of computational improvements for MDAV

MDAV task	Improvement strategy	Description
Distance calculation	Algebraic modification, precomputing	We propose using a property of the inner product to calculate and compare distances so that less operations are needed for microaggregation. Being these algebraic operations, their implementation is usually even optimized in multiple computing libraries.
Sorting	Use partial sorting (or selection)	Since MDAV does not strictly require total order when finding the $k - 1$ shortest points to a reference, a more relaxed version called partial sorting can help MDAV save computing resources.
Assignment of microcell	Reuse of distance calculations	Given that much of the running time of MDAV is devoted to calculating distances among the tuples of a data set and a reference point (to obtain its shortest and furthest points), such distances could be precomputed and then reused to prevent redundant operations.
Centroid calculation	Precomputation and reuse of calculations	We propose modifying the calculation of centroids in MDAV so that redundant addition operations are eliminated.
All calculations in general	Use less precision in the calculations	Since MDAV operations may not require much precision to build microcells, we propose changing the numerical representation to single precision so that less bits be processed, thus implying a reduction in MDAV's running time.

standard version of MDAV performs a series of mathematical operations over numerical values, we propose to use single precision for them in order to accelerate the microaggregation process.

This is the only modification that might imply a change in the results of the calculations performed by MDAV. Notwithstanding, since this algorithm might not require extremely precise operations (in terms of the number of decimal points considered), we expect no significant changes in the structure of microcells obtained with respect to the original version of MDAV, but a faster  $k$ -anonymous microaggregation.

#### 4. Experimental evaluation

In this section, we evaluate the efficiency of the proposed computational enhancements to MDAV. The objective of evaluation through experimentation of our approach is mainly to determine its impact on microdata. As mentioned in previous sections, such impact can be measured in terms of the algorithm’s speedup and the resulting data distortion (although in our case it is unlikely to occur) spawned by F-MDAV. Another objective is finding out whether such effect is independent of the data set employed.

##### 4.1. Evaluation criteria and data sets

We conduct this evaluation across two dimensions: *speedup* and *performance*. Speedup is measured as time gain, while performance is measured as the additional distortion incurred by the adapted versions of MDAV. However, since most of the proposed modifications do not change the internal operations of the algorithm, there is no additional distortion in the data and thus we mainly focus on measuring speedup. Below, we describe the experimental setup and our results from systematic tests over a variety of data sets. Such results are depicted for each enhancement and data set in Figs. 6, 7, and 8, while the overall speedup is illustrated in 11.

The evaluation of the computational performance of our methods has been conducted with three *standardized* data sets. These real data sets include “Large Census”, “Quant Forest” and “USA House”, which were previously used in [49, 12]. The “Large Census” data set has 149,642 records and includes 13 numerical attributes; “Quant Forest” has 581,012 records, from which we use a random sample of 150,000 records, and 10 numerical attributes. The “USA House” data set has 5,967,303 records and 13 numerical attributes. We use the “Large Census” data set since it is extensively employed in SDC, whereas “Quant Forest” and “USA House” data sets are used to validate the results obtained in “Large Census”. For our study, all attributes are considered to be quasi-identifiers.

##### 4.2. Experimental methodology

The experiments described in this section have been run in an Intel Core i7 CPU 3.4 GHz with 32 GB RAM. The microaggregation algorithm MDAV, its adapted versions, and the measurement tests are implemented entirely in Matlab R2017b, where, for the sake of fairness, we disable any form of parallelization.

In general, all versions of MDAV are parameterized with  $k = 10$ , which implies a reasonable level of privacy without incurring a significant distortion of data. Moreover, before microaggregation is applied, we follow the common practice of normalizing each column of the data set to have zero mean and unit variance.

To find the speedup obtained by our improvements, we measured the running time of MDAV before [26] and after implementing our refinements. Our reference MDAV is the algorithm specified in Algorithm 1. We shall refer to it as *traditional MDAV*. Furthermore, the modifications proposed in this work are applied to MDAV *individually*, with the aim of measuring their separate contribution to the speed of the microaggregation algorithm. Also, to show the combined effect of the five improvements, we implement them in a version of the algorithm we call *fast MDAV*.

Our experiments rely on a speedup to show how faster MDAV may become due to the proposed computational improvements. Let  $t_0$  be the running time of traditional MDAV and  $t$  the running time of any improved version of MDAV (including fast MDAV). Essentially, the speedup factor  $s = \frac{t_0}{t}$  tells us how fast this version is with respect to traditional MDAV. For instance, consider  $t_0 = 15$ ; if, after adapting the algorithm, its running time were reduced to  $t = 5$ , we would have gotten a speedup factor of  $s = 3\times$ , i.e., an MDAV that is 3 times faster.

Regarding our experimental *methodology*, we have a few final remarks. First, we assess MDAV over a varying number of records  $n$  with the aim of verifying the impact of our methods when the size of the data is increased. Thus, from each data set, we extract portions of data of varied sizes (different values of  $n$ ). For each value of  $n$ , the running time we measure is the averaged time that it takes MDAV to microaggregate  $n$  records. To that end, for every data set, we systematically obtain 3 random samples of  $n$  records each and then average the corresponding running times of MDAV. The running times for every improvement are registered and then compared with the time of traditional MDAV through the aforementioned speedup factor  $s$ .




Figure 6: Speedup factor ( $s$ ) of each of the five proposed improvements, i.e., when applied individually on the Large Census data set.




Figure 7: Speedup factor ( $s$ ) of each of the five proposed improvements, i.e., when applied individually on the Quant Forest data set.

The results of our computational strategies are presented in 4 bar charts; the first 3 illustrate, for each tested data set, the speedup factor obtained by every method. Although we experimented with several lengths, for the sake of visibility, the results are shown only for 3 representative values of  $n$  (10,000, 70,000 and 150,000). In the same line, the last bar chart exhibits the speedup factor reached by the fast MDAV, i.e., when all the improvements are consolidated within the same MDAV implementation. Although in essence the modifications we propose to MDAV do not imply a change in the numerical results of its internal operations, we verify whether or not each improvement leads to a variation of the built microcells or an increased distortion with respect to traditional MDAV.

#### 4.3. Experimental results

In the following subsections, we depict and explain the results of our experiments.




Figure 8: Speedup factor ( $s$ ) of each of the five proposed improvements, i.e., when applied individually on the USA House data set.

#### 4.3.1. Algebraic improvement

As already explained in §3.1, this method reduces the number of operations needed to calculate distances by taking advantage of a property of the inner product. Remarkably, numerical libraries are usually optimized for these algebraic operations.

Consequently, the results of our experiments show a significant speedup of MDAV that could reach a factor of 1.54×. This is depicted in Figs. 6, 7, and 8, for the three data sets we use, over which a homogeneous computational improvement is revealed.

Along with the partial sorting improvement, this algebraic strategy presents the best performance in terms of running time. In addition, the results of the cell assignment function  $c(j)$  and the resulting distortion of this new version of MDAV remain unchanged with respect to traditional MDAV.

#### 4.3.2. Distance reuse

Given the recurrence of distance calculations in MDAV, its running time can be reduced by precomputing some of such distances as theoretically explained in §3.2. Accordingly, after testing this improvement, when using the experimental setup described in §4.2, we observe a speedup factor between 1.14 and 1.31×. This execution performance is illustrated in Figs. 6, 7, 8 for the three data sets previously mentioned.

Once again, implementing this distance reuse does not reflect any variation neither in the structure of the microcells obtained nor in the distortion imposed to the data sets.

#### 4.3.3. Partial sorting

##### 4.3.3.1. Preliminary test.

As described in §3.3, the impact of sorting operations on the running time of MDAV could be reduced by using partial sorting, given its lower complexity with respect to total sorting. To illustrate the potential improvements due to partial sorting, we first perform an experiment in Matlab comparing two applications of both problems. Although not explicitly stated in the documentation, it is reasonable to assume that the functions `sort` and `mink` of Matlab R2017b implement variants of quicksort (total sorting) and quickselect (partial sorting) algorithms, respectively, as explained in §2. In fact, this experiment confirms that these functions follow the behavior of quicksort and quickselect in terms of computational complexity.

For this initial test, we do not only measure how long `sort` and `mink` take to execute over a list of random generated numerical values, but we try to mimic the sorting operations performed by MDAV through a few simple steps: sorting real-valued numbers and allocating and returning the indices of sorted values. It turns






Figure 9: Running time of different variants of sorting implemented in Matlab R2017b. Extensive testing is performed for several values of  $n$  (number of elements in the sorted list) and  $k$  (here representing the number of elements to be selected and sorted from the list, when partial sorting is tested). For the sake of clarity, double logarithmic scales are used.

that returning indices noticeably slows down the execution of sorting functions for short lengths. Moreover, when necessary, we preallocate outputs to exclude the time for memory allocation from our measurements. Particularly for the function `sort`, we also consider the time of trimming off the shortest values from the list.

The test involves more than 300 experiments, each of which consists in measuring the time it takes to obtain the  $k$  shortest values from a randomly generated list of length  $n$ . This process, along with the considerations of the last paragraph, emulate the role of sorting within MDAV. To evaluate the benefits of partial sorting over total sorting, we obtain the running times when using `mink` and `sort` functions to find the shortest values; we test their performance for several values of  $n$ , which ranges from 10 to 10 million, and for  $k = \{5, 10, 20, 50\}$ .

For the sake of reliability, we measure the running times for several repetitions in each experiment. Then, we compute the mean, as an estimate of average performance. Also, while the length of the data used is the same for every experiment, the values of the list are randomly sampled for every repetition. After a one-hour experiment, we found that our measurements were extremely reliable: the worst coefficient of variation, calculated as the standard deviation divided by the mean, was observed to be 1.63%.

Fig. 9 shows how our experiments take longer (the running time  $t$  gets higher) as the length  $n$  of the list increases. We use double logarithmic scales since we have very wide ranges of values for  $n$  and  $t$ , and thus extremely low and high values may appear. We can see that the running time for `mink` grows linearly with  $n$ , regardless of the value of  $k$  used, in line with the complexity  $\mathcal{O}(n)$  of the quickselect algorithm; this is important evidence that `mink` would be implementing a variant of this algorithm. For `sort`, the corresponding running times are certainly higher. However, the magnitude of the difference with respect to `mink` is not very clear. For that reason, we depict in Fig. 10 the running time per element  $\frac{t}{n}$  for every experiment. Using a semilogarithmic scale, this figure does show that `mink` (partial sorting) is much more efficient than `sort` (total sorting) since while for `mink` the running time per element is constant, said time grows logarithmically with  $n$  for `sort`. As a reference, we also plot the running time of the function `min` that retrieves only the lowest value from each list.




Figure 10: Running time of different variants of sorting implemented in Matlab R2017b. Here, we depict the time taken per element  $\frac{t}{n}$  ( $t$  is the time taken to sort a list of  $n$  elements) to have a clearer illustration of the remarkable performance of partial sorting implementations compared to those of total sorting. Again  $k$  represents the number of elements to be selected and sorted from the list in the case of partial sorting. Briefly, the running time of partial sorting remains constant for large values of  $n$ , while for total sorting time grows logarithmically. Also, for clarity, a semilogarithmic scale is used.

For the sake of verification, we also perform this experiment by using actual quicksort and quickselect algorithms implemented in C++. The tests for this experiment are also run in Matlab, where the corresponding C++ code is called through a gateway function in a MEX file. The results coincide with the ones described in the last paragraph.

#### 4.3.3.2. MDAV with partial sorting.

To estimate the speedup of microaggregation due partial sorting, we run the experiments according to the setup proposed in §4.2, now using a version of MDAV that relies on the function `mink` for microcell assignment. We then compare the resulting running times with those of traditional MDAV that uses `sort` by default.

As expected, partial sorting introduces interesting computational improvements. In fact, a speedup factor of almost  $2\times$  is reached for the Large Census data set when  $n = 150,000$ , as depicted in Fig. 6. However, the degree of improvement is not uniform for the three data sets, as it is shown in Fig. 7 and Fig. 8 for the partial sorting method where the maximum speedup factor does not attain  $1.5\times$ . This seems reasonable since the complexity of the variants of sorting may depend on the intrinsic structure of the data.

#### 4.3.4. Centroid by subtraction

It is clear from §3.4 that the operations for the calculation of centroids in MDAV are subject to redundancy since the tuples of a microdata set have to be added recurrently to find a representative mean. Our centroid by subtraction strategy, which uses precomputing and subtraction, obtains a speedup factor of up to  $1.15\times$  as shown in Figs. 6, 7 and 8. Although performance gain is more moderate for MDAV, it is still important, considering that its implementation does not represent any additional cost in terms of distortion.

#### 4.3.5. Single precision

Being MDAV an algorithm whose calculations may not require extreme precision, our last method is based on using single precision for the corresponding mathematical operations. Strikingly, this modification allows

a computational improvement that is even better than that of centroid by subtraction, i.e., a speedup factor of up to  $1.35\times$ , as depicted in Figs. 6, 7, 8, for our three data sets.

As anticipated in §3.5, the results from using this strategy show a slight variation in the structure of the microcells built by MDAV. However, the resulting distortion remains unchanged.

#### 4.3.6. Fast MDAV

Our last series of experiments analyze the case when all proposed modifications are combined in a single version of the microaggregation algorithm, fast MDAV. The tests, following the setup in §4.2, show remarkable results, as reflected in Fig. 11. We confirm that fast MDAV is up to 4 times faster than the original version and that, as expected, no additional distortion is introduced in the three microaggregated data sets.




Figure 11: Overall speedup factor  $s$  of our fast MDAV obtained over the three data sets. The five strategies are consolidated in a single version and it is tested for several values of  $n$ . Due to space considerations, only the results of tests for  $n = 10\,000, 70\,000, 150\,000$  are depicted in this figure.

#### 4.3.7. Information loss with F-MDAV

As mentioned when describing each computational improvement in §3, the rationale behind our work is simplifying redundant operations when implementing MDAV. Such simplification is based on finding alternate algebraic expressions, reusing and precomputing (repetitive) calculations, adapting a more relaxed sorting strategy, and even using less precise calculations.

When single precision is used for calculations, there is a risk of obtaining a structure of microcells different from that of MDAV. However, microaggregation does not require extremely precise calculations since those are used only for comparing distances between points. Namely, in practice, our proposal does not imply any modification of the resulting  $k$ -anonymous groups built by MDAV, so there might not be a price in distortion. As a matter of fact, we have verified that microcell allocation remains invariable after F-MDAV is implemented. Consequently, our strategies do not incur in additional distortion or information loss with respect to that provoked by original MDAV.

In times when the world revolves around big data, processing time quickly becomes a bottleneck with respect to the potential applications of large-scale databases. Moreover, domains as critical as health, vehicular traffic, or network intrusion detection are currently using tons of data to help computational systems make real-time, and even life-or-death decisions. Due to such demanding requirements, privacy issues related to data processing are commonly overshadowed. Thus, from the perspective of privacy, we feel that any improvement in (computing) efficiency is not negligible, particularly when the strategy does not

entail additional data distortion, and even more when its multiplicative effect may turn a privacy mechanism feasible for a critical application.

## 5. Conclusions and future work

This work addresses the problem of computational complexity of  $k$ -anonymous microaggregation for large datasets with a substantial amount of numerical records. Our approach strives to obtain a more usable privacy mechanism in contexts brought by the current big data era. We have proposed five algorithmic and algebraic strategies that significantly reduce the running time of MDAV, a widely-known microaggregation algorithm, while leaving untouched the resulting utility of the anonymized data.

Our improvement strategies mainly span the reuse of calculations, precomputing, algebraic modifications, and a relaxed approach of sorting; all of them are implemented in the main tasks performed by MDAV, e.g., distance calculation, sorting, assignment of microcell, and centroid calculation. As expected, the strategies focused on the most repetitive operations in MDAV (e.g., such as distance calculation and sorting) lead to the highest performance, making the modified MDAV up to twice faster than the original one. However, the specific improvement may vary with the data set used since the performance of some of these operations (particularly sorting) depends on the internal structure of the numerical data.

The results of systematic experimentation in three data sets show that the combined effect of our strategies produce a version of MDAV whose running time is four times shorter than its original version. The remarkable overall effect in our fast MDAV seems to be uniform for the three data sets tested and is conveniently attained without incurring in additional distortion on the data. A similar result is observed when testing a synthetic data set. Our concluding message aims at calling attention to put more emphasis on the multidimensional trade off among privacy, utility and usability brought by the emerging context of big data when dealing with old and new approaches of privacy protection. Future directions to further improve the performance of microaggregation may involve, e.g., reducing its inherent complexity by resorting to eliminate its potential statistical redundancy, or applying other lossy techniques that give up some data utility in exchange for usability of the privacy mechanism.

## Acknowledgment

This manuscript presents some of the results developed through the collaboration of the Universitat Politècnica de Catalunya (UPC) and Scytl Secure Electronic Voting S.A. (Scytl) in the context of the project “Data-Distortion Framework”, and in accordance with the guidelines therein. This work is thus partly supported by the Spanish Ministry of Industry, Energy and Tourism (MINETUR) through the “Acción Estratégica Economía y Sociedad Digital (AEESD)” funding plan, through the aforementioned project, “Data-Distortion Framework (DDF)”, ref. TSI-100202-2013-23.

Additional funding supporting this work has been granted to UPC and Universitat Rovira i Virgili (URV) by the Spanish Ministry of Economy and Competitiveness (MINECO) through the “Anonymized Demographic Surveys (ADS)” project, ref. TIN2014-58259-JIN, under the funding program “Proyectos de I+D+i para Jóvenes Investigadores”, through the projects “INRISCO”, ref. TEC2014-54335-C4-1-R, “Smart-Glaciés”, ref. TIN2014-57364-C2-1-R, “Sec-MCloud”, ref. TIN2016-80250-R, as well as by the Government of Catalonia, under grant 2014 SGR 1504.

Partial support to this work has also been granted to URV by the European Commission through projects H2020-644024 “CLARUS” and H2020-700540 “CANVAS”. J. Parra-Arnau is the recipient of a Juan de la Cierva postdoctoral fellowship, IJCI-2016-28239, from the MINECO. J. Parra-Arnau is with the UNESCO Chair in Data Privacy, but the views in this paper are his own and are not necessarily shared by UNESCO.

## References

## References

- [1] W. Raghupathi, V. Raghupathi, Big data analytics in healthcare: Promise and potential, Health Inform. Sci., Syst. 2.

- [2] D. Li, X. Shen, L. Wang, Connected geomatics in the big data era, *Int. J. Digit. Earth* (2017) 1–15.
- 555 [3] A. Monaco, *Big data: The measure of humankind*, Times High. Educ. (THE).  
URL [www.timeshighereducation.com/comment/big-data-measure-humankind#survey-answer](http://www.timeshighereducation.com/comment/big-data-measure-humankind#survey-answer)
- [4] L. Sweeney, Uniqueness of simple demographics in the U.S. population, Tech. Rep. LIDAP-WP4, Carnegie Mellon Univ., Sch. Comput. Sci., Data Priv. Lab., Pittsburgh, PA (2000).
- 560 [5] A. Narayanan, V. Shmatikov, Robust de-anonymization of large sparse datasets, in: *Proc. IEEE Symp. Secur., Priv. (S&P)*, Oakland, CA, 2008, pp. 111–125.
- [6] *AOL search data scandal* (Aug. 2006).  
URL [http://en.wikipedia.org/wiki/AOL\\_search\\_data\\_scandal](http://en.wikipedia.org/wiki/AOL_search_data_scandal)
- [7] Assoc. Press, *Divided by citizen wealth tables*, NY Times.  
URL [www.nytimes.com/2009/10/24/business/global/24tax.html?\\_r=2&ref=global](http://www.nytimes.com/2009/10/24/business/global/24tax.html?_r=2&ref=global)
- 565 [8] P. Samarati, Protecting respondents' identities in microdata release, *IEEE Trans. Knowl., Data Eng.* 13 (6) (2001) 1010–1027.
- [9] C. Dwork, Differential privacy, in: *Proc. Int. Colloq. Automata, Lang., Program. (ICALP)*, Vol. 4052 of *Lect. Notes Comput. Sci. (LNCS)*, Venice, Italy, 2006, pp. 1–12.
- [10] L. Sankar, S. R. Rajagopalan, H. V. Poor, Utility-privacy tradeoffs in databases: An information-theoretic approach, *IEEE Trans. Inform. Forensics, Secur.* 8 (6) (2013) 838–852.
- 570 [11] D. Rebollo-Monedero, J. Forné, E. Pallarès, J. Parra-Arnau, *A modification of the Lloyd algorithm for  $k$ -anonymous quantization*, *Inform. Sci.* 222 (2013) 185–202.  
URL <http://doi.org/10.1016/j.ins.2012.08.022>
- [12] E. Pallarès, D. Rebollo-Monedero, A. Rodríguez-Hoyos, J. Estrada-Jiménez, A. M. Mezher, J. Forné, Mathematically optimized, recursive repartitioning strategies for  $k$ -anonymous microaggregation of large-scale datasets, *Expert Systems with Applications* (2019) 113086.
- 575 [13] D. Rebollo-Monedero, A. M. Mezher, X. Casanova Colomé, J. Forné, M. Soriano, Efficient  $k$ -anonymous microaggregation of multivariate numerical data via principal component analysis, *Information Sciences* (2019) 417–443.
- [14] A. Calviño, A simple method for limiting disclosure in continuous microdata based on principal component analysis, *J. Official Stat.* 33 (1) (2017) 15–41.
- 580 [15] R. Mortazavi, S. Jalili, Fine granular proximity breach prevention during numerical data anonymization, *Transactions on Data Privacy* (2017) 117–144.
- [16] M. Laszlo, S. Mukherjee, Minimum spanning tree partitioning algorithm for microaggregation, *IEEE Trans. Knowl., Data Eng.* 17 (7) (2005) 902–911.
- 585 [17] X. Sun, H. Wang, J. Li, Y. Zhang, An approximate microaggregation approach for microdata protection 39 (2) (2012) 2211–2219.
- [18] R. Mortazavi, S. Jalili, H. Gohargazi, Fast data-oriented microaggregation algorithm for large numerical datasets, *Knowl.-Based Syst.* 67 (3) (2014) 195–205.
- [19] D. Rebollo-Monedero, J. Forné, M. Soriano, J. Puiggalí Allepuz,  *$p$ -Probabilistic  $k$ -anonymous microaggregation for the anonymization of surveys with uncertain participation*, *Inform. Sci.* 382–383 (2017) 388–414.  
URL <http://doi.org/10.1016/j.ins.2016.12.002>
- 590 [20] D. Rebollo-Monedero, A. Mohamad Mezher, X. Casanova Colomé, J. Forné, M. Soriano, Efficient  $k$ -anonymous microaggregation of multivariate numerical data via principal component analysis, *Inform. Sci.* 503 (2019) 417–443, in press.
- [21] J. Parra-Arnau, J. Domingo-Ferrer, J. Soria-Comas, Differentially private data publishing via cross-moment microaggregation 53 (2) (2020) 269–288.
- 595 [22] A. Hundepool, A. V. de Wetering, R. Ramaswamy, L. Franconi, A. Capobianchi, P.-P. de Wolf, J. Domingo-Ferrer, V. Torra, R. Brand, S. Giessing,  *$\mu$ -ARGUS version 3.2 software and user's manual*, Stat. Neth., Voorburg, Netherlands (2003).  
URL <http://neon.vb.cbs.nl/casc>
- 600 [23] J. Domingo-Ferrer, J. M. Mateo-Sanz, Practical data-oriented microaggregation for statistical disclosure control, *IEEE Trans. Knowl., Data Eng.* 14 (1) (2002) 189–201.
- [24] A. Rodríguez-Hoyos, J. Estrada-Jiménez, D. Rebollo-Monedero, J. Parra-Arnau, J. Forné, *Does  $k$ -anonymous microaggregation affect machine-learned macro trends?*, *IEEE Access* 6 (2018) 28258–28277.  
URL <http://doi.org/10.1109/ACCESS.2018.2834858>
- 605 [25] M. Templ, *Introduction to Statistical Disclosure Control (sdc)*, Res. rep. (Jul. 2014).  
URL <https://pdfs.semanticscholar.org/311d/73e8e4b5b1dca49901d929f383bbd0817a7c.pdf>
- [26] J. Domingo-Ferrer, V. Torra, Ordinal, continuous and heterogeneous  $k$ -anonymity through microaggregation, *Data Min., Knowl. Discov.* 11 (2) (2005) 195–212.
- [27] D. Rebollo-Monedero, J. Forné, M. Soriano, *An algorithm for  $k$ -anonymous microaggregation and clustering inspired by the design of distortion-optimized quantizers*, *Data, Knowl. Eng.* 70 (10) (2011) 892–921.  
URL <http://doi.org/10.1016/j.datak.2011.06.005>
- 610 [28] D. Defays, P. Nanopoulos, Panels of enterprises and confidentiality: The small aggregates method, in: *Proc. Symp. Design, Anal. Longit. Surv., Stat. Can.*, Ottawa, Canada, 1993, pp. 195–204.
- [29] J. Domingo-Ferrer, F. Sebé, A. Solanas, A polynomial-time approximation to optimal multivariate microaggregation, *Comput., Math., Appl.* 55 (4) (2008) 714–732.
- 615 [30] J. Domingo-Ferrer, V. Torra, A critique of  $k$ -anonymity and some of its enhancements, in: *Proc. Workshop Priv., Secur., Artif. Intell. (PSAI)*, Barcelona, Spain, 2008, pp. 990–993.

- [31] D. Rebollo-Monedero, J. Forné, J. Domingo-Ferrer, [From  \$t\$ -closeness-like privacy to postrandomization via information theory](#), IEEE Trans. Knowl., Data Eng. 22 (11) (2010) 1623–1636.  
 URL <http://doi.ieeecomputersociety.org/10.1109/TKDE.2009.190>
- [32] D. Rebollo-Monedero, J. Parra-Arnau, C. Díaz, J. Forné, [On the measurement of privacy as an attacker's estimation error](#), Int. J. Inform. Secur. 12 (2) (2013) 129–149.  
 URL <http://doi.org/10.1007/s10207-012-0182-5>
- [33] T. M. Truta, B. Vinay, Privacy protection:  $p$ -Sensitive  $k$ -anonymity property, in: Proc. Int. Workshop Priv. Data Mgmt. (PDM), Atlanta, GA, 2006, p. 94.
- [34] X. Sun, H. Wang, J. Li, T. M. Truta, Enhanced  $p$ -sensitive  $k$ -anonymity models for privacy preserving data publishing, Trans. Data Priv. 1 (2) (2008) 53–66.
- [35] N. Li, T. Li, S. Venkatasubramanian,  $t$ -Closeness: Privacy beyond  $k$ -anonymity and  $l$ -diversity, in: Proc. IEEE Int. Conf. Data Eng. (ICDE), Istanbul, Turkey, 2007, pp. 106–115.
- [36] J. Brickell, V. Shmatikov, The cost of privacy: Destruction of data-mining utility in anonymized data publishing, in: Proc. ACM SIGKDD Int. Conf. Knowl. Discov., Data Min. (KDD), Las Vegas, NV, 2008.
- [37] D. Rebollo-Monedero, J. Forné, J. Domingo-Ferrer, From  $t$ -closeness to PRAM and noise addition via information theory, in: Proc. Int. Conf. Priv. Stat. Databases (PSD), Lect. Notes Comput. Sci. (LNCS), Istanbul, Turkey, 2008, pp. 100–112.
- [38] J. L. Lin, T. H. Wen, J. C. Hsieh, P. C. Chang, Density-based microaggregation for statistical disclosure control, Expert Syst., Appl. 37 (4) (2010) 3256–3263.
- [39] N. Matatov, L. Rokach, O. Maimon, Privacy-preserving data mining: A feature set partitioning approach, Inform. Sci. 180 (14) (2010) 2696–2720.
- [40] J. Domingo-Ferrer, Ú. González-Nicolás, Hybrid microdata using microaggregation, Inform. Sci. 180 (15) (2010) 2834–2844.
- [41] A. Mohamad Mezher, A. García-Álvarez, D. Rebollo-Monedero, J. Forné, Computational improvements in parallelized  $k$ -anonymous microaggregation of large databases, in: Proc. IEEE Int. Conf. Distrib. Comput. Syst. (ICDCS), Workshop Priv., Secur. Big Data (PSBD), Atlanta, GA, 2017, pp. 258–264.
- [42] C. A. R. Hoare, Algorithm 64: Quicksort, Commun. ACM 4 (7) (1961) 321.
- [43] C. A. R. Hoare, Algorithm 65: Find, Commun. ACM 4 (7) (1961) 321–322.
- [44] D. Musser, Introspective sorting and selection algorithms, J. Softw.: Pract., Exp. 27 (8) (1999) 983–993.
- [45] J. Williams, Algorithm 232: Heapsort, Commun. ACM 7 (6) (1964) 347–348.
- [46] R. W. Floyd, R. L. Rivest, Algorithm 489: The algorithm SELECT —For finding the  $i^{\text{th}}$  smallest of  $n$  elements, Commun. ACM 18 (3) (1975) 173.
- [47] R. W. Floyd, R. L. Rivest, Expected time bounds for selection, Commun. ACM 18 (3) (1975) 165–172.
- [48] K. C. Kiwiol, On Floyd and Rivest's SELECT algorithm, Theor. Comput. Sci. 347 (1–2) (2005) 214–238.
- [49] M. Solé, V. Muntés-Mulero, J. Nin, Efficient microaggregation techniques for large numerical data volumes data volumes, Int. J. Inform. Secur. 11 (2012) 253–267.