Attribute disclosure risk in smart meter data[⋆]

Guillermo Navarro-Arribas[1][0000−0003−3535−942X] and Vicenç Torra[2][0000−0002−0368−8037]

[1] Dep. of Information and Communications Engineering, Universitat Autonoma de Barcelona,
guillermo.navarro@uab.cat
[2] Department of Computing Science, Umeå University,
vtorra@cs.umu.se

Abstract. This paper studies attribute disclosure risk in aggregated smart meter data. Smart meter data is commonly aggregated to preserve the privacy of individual contributions. The published data shows aggregated consumption, preventing the revelation of individual consumption patterns. There is, however, a potential risk associated to aggregated data. We analyze some datasets of smart meter data consumption to show the potential risk of attribute disclosure. We observe that, even if data is aggregated with the most favorable aggregation approach, it presents this attribute disclosure risk.

Keywords: Smart meter data · k-anonymity · Attribute disclosure.

1 Introduction

Smart meter data provides a detailed insight on home energy consumption and can be used for different purposes: prediction of future consumption, redesign of power grids, or simply for marketing and public knowledge. Several approaches exist to anonymized smart meter data, and most of them are based on providing some sort of aggregation [1,4,11]. That is, data is aggregated to preserve the privacy associated to each smart meter, and consequently to each home power consumption.

There are different approaches to aggregate smart meter data with privacy purposes. One approach is to globally aggregate all the data from all available sources. In other schemes the data is aggregated geographically in distribution points. In any case, aggregated data is usually considered safe. Several smart meters (homes) readings are aggregated together forming an anonymity set, this

[⋆] Preprint of the publication:

Navarro-Arribas, G., Torra, V. (2024). Attribute Disclosure Risk in Smart Meter Data. In: Domingo-Ferrer, J., Önen, M. (eds) Privacy in Statistical Databases. PSD 2024. Lecture Notes in Computer Science, vol 14915. Springer, Cham. https://doi.org/10.1007/978-3-031-69651-0_18

is commonly extended to ensure k-anonymity [17] in the released dataset. There are however some potential risks associated to aggregated smart meter data. In this paper we investigate a potential attribute disclosure attack, where a user, whose data is included in the dataset, can estimate some partial information from the anonymized dataset.

The idea is to check if the aggregation of data could lead to situations where the aggregated value can be used to approximate the original value of a respondent. To that end, we investigate smart meter data aggregated using microaggregation, which can be seen as the most favorable aggregation to avoid this kind of attacks.

The goal of the paper is to outline the problem and make aware the community about it. This paper complements our previous results [3,2] discussing some specific attacks to smart grid data. More particularly, we used Non-Intrusive Load Monitoring (NILM) with the goal of detecting individual appliances in aggregated data.

The paper is organized as follows. Section 2 discusses smart meter data and its aggregation, Section 3 describes the attribute disclosure problem in aggregated data, and Section 4 analyzes the attribute disclosure in smart meter data. Section 5 discusses some aspects of e attribute disclosure attack, and Section 6 concludes the paper.

2 Aggregation of smart meter data

We consider a generic case of smart meter data, where the data provided by a smart meter can be seen as time series. A time series is a sequence of values taken at some time interval. In our case, we will consider regular time periods, and for each time period we will have the specific power consumption reading provided by the smart meter. We denote the time series of smart meter i as $x_i = (x_i^t : t \in T)$, where T is the time based index.

In our case, we will consider $T = t_0, \ldots, t_m$, for all time series in the dataset. Each time series corresponds to a different smart meter, and we consider readings for a single day, having time periods of 30 minutes, thus $m = 48$. In some sense the dataset could be seen as a microdata file where each record is a time series and the attributes are the consumption for each time period.

As previously stated, it is quite common to aggregate smart meter data in order to provide some degree of privacy regarding individual consumption patterns. Moreover, this aggregation could be done to provide k-anonymity guarantees in the protected data.

We can observe two different approaches to perform this aggregation:

- Global aggregation: the data publisher has the data from all smart meters and can then aggregate them at once. This allows to perform, for example, microaggregation of the time series. The advantage of this approach is that the aggregation can better preserve information loss.

- Local aggregation: in this case, the data is aggregated locally, usually by some power grid distribution center. Records are aggregated together based on e.g. geographical distribution and not the actual consumption pattern.

In this study, we will consider the global aggregation case since it is the one expected to have a better behavior regarding attribute disclosure risk. If attribute disclosure is possible in globally aggregated data, it will surely be possible in locally aggregated data. This claim is based on the intuition that global aggregation will provide more homogeneous groups, where consumption patters will be more similar and thus, more difficult to differentiate.

2.1 Datasets

We have used 3 datasets from two different sources. From each dataset we took only one day, that is, 48 readings for each smart meter. We chose the day as the one that had more readings from the dataset. The three datasets are:

- *banes*: *BANES Energy Data Electricity* [5] is a dataset with electricity energy usage data in Council buildings from Bath and Nord East Somerset. The dataset shows consumption in 30 minutes slots. We have taken the consumption for 2019-11-13, which consists of 79 different buildings.
- *cer*: which contains electricity and gas consumption data from the Commission of Energy Regulation, as provided by the *Irish Social Science Data Archive* [10], also for 30 minutes slots. Here we consider:
  - *cer-elec*: electricity consumption for 2009-08-20, with consumption from 983 houses.
  - *cer-gas*: gas consumption for 2009-12-03, with consumption from 1493 houses.

We also assume that the values of each reading will be positive, since we are only considering energy consumption.

2.2 Smart meter data microaggregation

We have considered a global microaggregation of the smart data. Microaggregation [7,8,13] is a well known method for data privacy that is commonly used to provide k-anonymity. It builds small clusters and then replaces each or the records in the cluster by the cluster center. As each cluster has at least k records, when all the records are replaced by the same cluster center, k-anonymity is satisfied.

Microaggregation is formulated as an optimization problem with specific constraints. The objective function resembles that of k-means, where cluster centers are considered, and records are assigned to the nearest cluster center. Constraints ensure that each record is assigned to exactly one cluster, with each cluster containing at least k records (and at most $2k$).

When considering more than one variable, that is, multivariate microaggregation, the problem becomes NP-hard [16] so heuristic methods have to be used.

MDAV [7,20] is one of such methods and has been extensively used in the literature. In this work we use MDAV for microaggregation. We have used values of $k = 2, \ldots, 41$.

Moreover, to microaggregate time series we need to define a distance function to form the clusters, and an aggregator operator to compute the cluster representative. Given that the time series are aligned, we use the Euclidean distance and the average. The average is commonly used in smart meter data aggregation.

3 Attribute disclosure in aggregated smart data

Regarding k-anonymity, attribute disclosure is commonly associated to categorical confidential attributes. Common attacks such a homogeneity, similarity, or skewness attacks [18] are considered on the distribution of confidential attributes. To prevent such attacks there are well known proposals such as p-sensitivity [21], l-diversity [15], or t-closeness [14].

Less known are attacks on the masked numeric attributes. In a smart meter dataset, all attributes are masked, and no confidential attributes are considered. In this case, we show that some types of attacks are possible, usually considering an internal attacker.

In [19] some metrics are proposed to estimate the likelihood of this kind of attacks. These metrics are based on sensitivity rules commonly used in tabular data (see e.g. [6,9,12] for details). We adapt them here to measure the sensitivity of aggregated smart meter data.

We consider a given smart meter dataset X with n time series or records, with m consumption readings. We assume the data has been protected resulting in a protected dataset $X' = \rho(X)$, where ρ is the protection method, in our case, microaggregation.

Such dataset X' consists of clusters of record, each cluster with equal time series. In each one we can observe different cells, one for each time reading. Let us consider that for a given cell we have t contributors which provide the original values $c_1, \ldots, c_t$. If $\bar{c}$ is the average, $\bar{c} = \frac{1}{t}\sum_i c_i$, and the protected cell will have this value for all the records.

Considering the contribution of each record to the cell average, we can check the following sensitivity rules to denote if a cell is sensitive. In this context, a sensitive cell is a cell where attribute disclosure can take part.

***(n, r)*-dominace**: The rule (n, r)-dominance determines that the cell is sensitive when n contributors represent more than the r fraction of the total. If we consider the values c_i ordered in decreasing order, $c_{\sigma(1)} \geq c_{\sigma(2)} \geq \cdots \geq c_{\sigma(t)}$, this rule will detect a cell as sensitive when

$$\frac{\sum_{i=1}^{n_r} c_{\sigma(i)}}{\sum_{i=1}^{t} c_i} > r. \tag{1}$$

***p*% rule**: The rule $p\%$ is stated as follows. A cell is sensitive when an intruder can estimate the contributor within p percent, taking into account the released table. It can be proven that the best estimation is the one of the second-largest contributor (i.e., the one which contributes with $c_{\sigma(2)}$) on the largest one. Then a cell is sensitive when

$$\sum_{i=3}^{t} c_{\sigma(i)} < p c_{\sigma(1)}. \quad (2)$$

In this expression we use p as a value in [0,1] instead of a percentage.

Hundepool et al. [12] recommend the use of $p' = (1-r_r)/r_r$ (and $p\% = 100p'$) as providing a risk assessment similar to the $(2, r_r)$ rule. E.g., for $n_r = 2$ and $r_r = 0.6$, we would have $p = 66\%$. In general, for the dominance rule, parameterization with $n = 1$ or 2 and $r > 0.6$ have been considered in the literature. For the rule $p\%$, a parameter larger than 60% has also been considered in the literature. We will use $n = 1$, $n = 2$, $r = 0.6$, and $p = 66\%$ in our experiments.

4 Attribute disclosure risk in smart meter data

In this section, we analyze the sensitivity of smart meter data according to the sensitivity rules commented in Section 3. We have considered the three datasets commented in Section 2.1. Each dataset has been protected with microaggregation for $k = 2, \ldots, 41$ using the MDAV algorithm. Table 1 shows a summary of the number of cells for each dataset and some values of k for each dataset. We include the number of sensitive cells according to the (n, r)-dominance for $n = 2$, $n = 1$, $r = 0.6$, and the $p\%$ rule for $p = 66\%$ (See 3).

	banes				*cer-elec*				*cer-gas*			
k	*cells*	*dom2*	*dom1*	*p100*	*cells*	*dom2*	*dom1*	*p100*	*cells*	*dom2*	*dom1*	*p100*
2	1872	0	669	1410	23568	0	8671	23556	35808	0	4728	19306
4	912	184	623	384	11760	1426	10861	5068	17904	980	5582	9650
6	624	325	436	129	7824	5566	7683	723	11904	2283	4787	6425
8	432	274	312	109	5856	5349	5818	153	8928	2313	4187	4838
10	336	224	243	83	4704	4508	4688	62	7152	2218	3873	3789
15	240	195	196	44	3120	3097	3118	6	4752	1954	3103	2234
20	144	100	105	44	2352	2341	2352	2	3552	1722	2565	1454
30	96	96	96	0	1536	1536	1536	0	2352	1403	1920	743

Table 1: Number of cells (*cells*), dominance for $n = 2$, $r = 0.6$ (*dom2*), and for $n = 1$, $r = 0.6$ (*dom1*), and $p\%$ for $p = 66\%$ (*p100*) for each dataset masked with different values of k.

More detailed results are shown in Figures 1, 2, 3. Both sensitivity rules are show in absolute value and percentage over the total number of cells for each dataset.

In general, the dominance rule highlights more cells as sensitive than the $p\%$ rule. Specially, the dominance rule for $n = 1$ gives a very high percentage of sensitive cells. The attacker could estimate the highest consumer with a 60% of

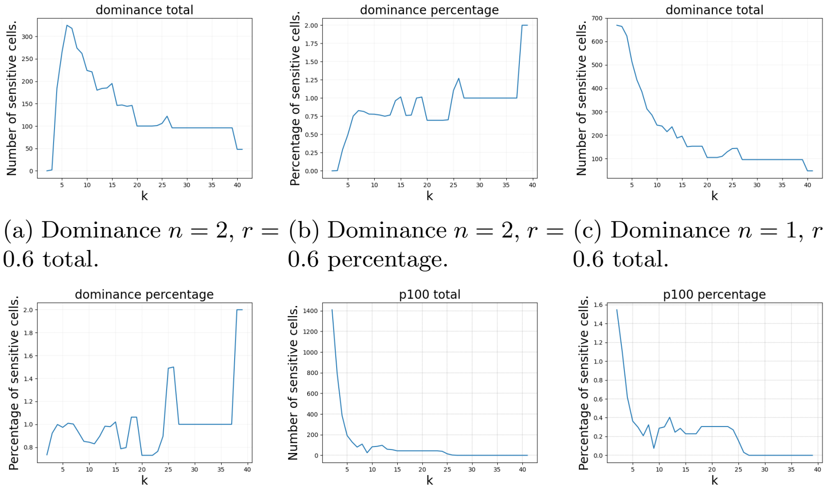

(a) Dominance $n = 2$, $r = 0.6$ total.

(b) Dominance $n = 2$, $r = 0.6$ percentage.

(c) Dominance $n = 1$, $r = 0.6$ total.

(d) Dominance $n = 1$, $r = 0.6$ percentage.

(e) p% rule $p = 0.66$ total.

(f) p% rule $p = 0.66$ percentage.

Fig. 1: Dominance and p% rule for the dataset *banes*.

confidence. These results show that even masked data can have sensitive cells, yielding attribute disclosure. This means cases in which a user can have a good estimation of the consumption of the rest of the users. It is also important to note that we have employed a global masking approach. All records are protected globally resulting in a more homogeneous microaggregation, a geographical aggregation will lead to a higher number of sensitive cells.

5 Internal attacks on aggregated data

In this section, we discuss the potential attack performed by an internal user attempting to estimate consumption from the masked dataset. As an example scenario, suppose that a user knows its own consumption for the day, and has the protected dataset with the consumption from all users aggregated. This user can attempt to estimate the consumption of users that fall in its own anonymity set (microcluster). To do that, the attacker needs to:

1. Identify the microcluster where its consumption has been aggregated.
2. Estimate the average consumption of the other users.

For the first step, the attacker can estimate the closest aggregated time series consumption to its own consumption and assume that it will be its own microcluster or anonymity set.

We denote the aggregate dataset X' as composed of s microclusters of time series: $C_1, C_2, \ldots C_s$, with their respective cluster centers $\bar{C_1}, \bar{C_2}, \ldots, \bar{C_s}$. The

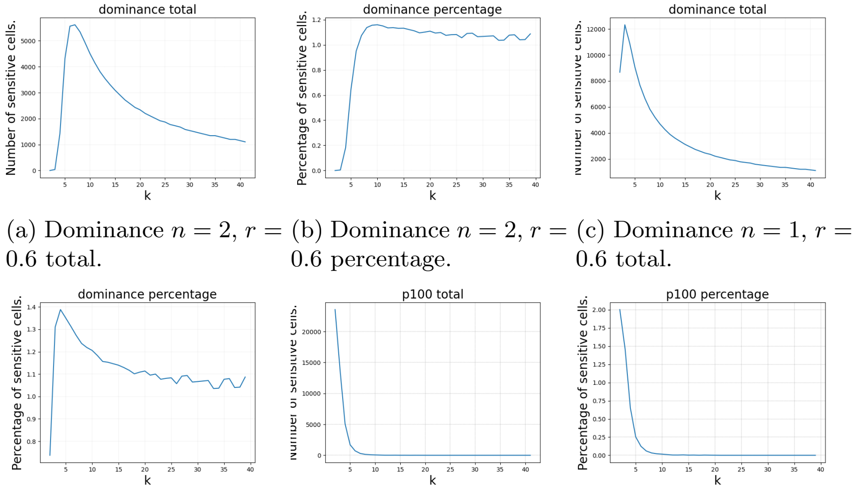

(a) Dominance $n = 2$, $r = 0.6$ total.

(b) Dominance $n = 2$, $r = 0.6$ percentage.

(c) Dominance $n = 1$, $r = 0.6$ total.

(d) Dominance $n = 1$, $r = 0.6$ percentage.

(e) p% rule $p = 0.66$ total.

(f) p% rule $p = 0.66$ percentage.

Fig. 2: Dominance and p% rule for the dataset *cer-elec*.

attacker with a consumption time series x_a attempts to identify the cluster C_a where its data have been aggregated as:

$$C_a = \{C_i \mid \arg\min_i d(x_a, \bar{C}_i)\} \tag{3}$$

The attack is successful if $c_a \in C_a$. Here, d is a distance function for time series. We use the Euclidean distance (see Section 2.2).

Figure 4 shows the microcluster reidentification success ratio for each dataset. This has been computed by randomly selecting a record from the original dataset and then attempting to identify the cluster in the protected dataset according to Eq. (3). The result is the average of 100 executions for each case.

We can see that the success ratio is mostly above 50%, but in general, one could expect a bigger success ratio given the microaggregation was performed on the whole dataset. It is thus expected that a locally aggregated dataset based on e.g. geographical distribution could yield worse results.

The second step, estimating the average consumption of other users, can obviously be estimated from the published cluster representative. Given the cluster where the attacker data is masked, C_a with its representative $\bar{C}_a$, if we consider that the cluster has r records (time series) we could see a cell of consumption values for a given time period t as $C_a^t = \{c_{a1}^t, c_{a2}^t, \ldots, c_{cr}^t\}$ with an average of $\bar{C}_a^t$. The attacker can estimate the average consumption of the other $t-1$ users. Let c_{a1} be attacker's consumption, with this information can easily set an upper bound of the estimation. The maximum value for any c_{ai} for $i = 2, \ldots, r$ will be:

$$\max_{i=2,\ldots,r} \{c_{ai}\} \leq (r \cdot \bar{C}_a^t) - c_{a1} \tag{4}$$

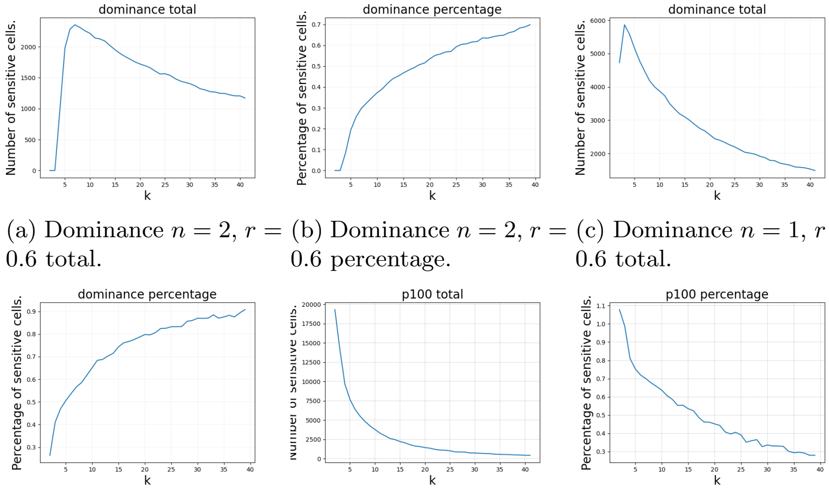

(a) Dominance $n = 2$, $r = 0.6$ total.

(b) Dominance $n = 2$, $r = 0.6$ percentage.

(c) Dominance $n = 1$, $r = 0.6$ total.

(d) Dominance $n = 1$, $r = 0.6$ percentage.

(e) $p\%$ rule $p = 0.66$ total.

(f) $p\%$ rule $p = 0.66$ percentage.

Fig. 3: Dominance and $p\%$ rule for the dataset *cer-gas*.

The minimum will be 0 for $r > 2$.

This estimation is somehow expected in aggregated data, but if the cell is sensitive, it means that the attribute disclosure is more critical. Prior knowledge will allow the attacker to estimate the higher consumer with more precision.

5.1 Practical implications discussion

From a practical point of view, an attacker could force the previously described attack if its smart meter is taking part in the anonymized dataset.

We have considered the worst case for the attacker in order to give some insight on the potential problem of attribute disclosure, but the attack is more significant on locally aggregated data. In such cases, the attacker will attempt to gain information about a neighbor (probably known neighbors).

One problem with locally aggregated data is that the attacker might not easily identify its microcluster in the protected dataset. To increase the probability of succeeding in the attack, the attacker can generate a high consumption pick in a specific time period. This will affect the average of this specific time period, making the aggregate easily identifiable. Alternatively, the attacker can induce some particular patterns in the consumption to inject specific patterns (signatures) in the aggregated data. Then, using non-intrusive load monitoring (NILM) techniques, these signatures [3] can be extracted from aggregated data.

We observe that this attack is feasible on smart meter data using real data and should thus be considered for further research.

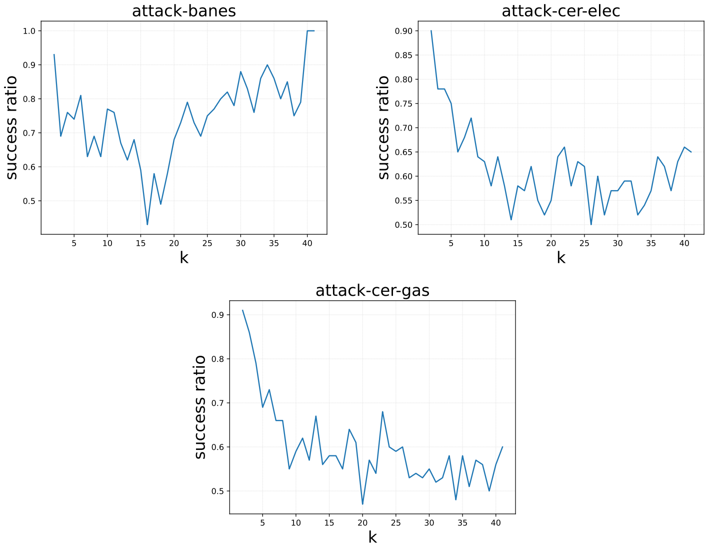

Fig. 4: Microcluster reidentification success ratio for global microaggregation in the *banes*, *cer-elec*, and *cer-gas* dataset.

6 Conclusions

In this paper we have analyzed attribute disclosure in aggregated smart meter data. We have shown that there is potential risk for attribute disclosure by checking sensitivity of aggregated values in a global microaggregation. This leads to assume that locally aggregate data will introduce more risk, leading to higher possibilities of attribute disclosure.

Acknowledgements

This research has been partially supported by the project DANGER C062/23 within the Plan de Recuperación, Transformación y Resiliencia funded with Next Generation EU funds. Spanish Ministry under Grant PID2021-125962OB-C33 SECURING/NET. Catalan AGAUR under Grant SGR2021-00643. The Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by Knut and Alice Wallenberg Foundation. Support also received from the Swedish Research Council under the project Privacy for complex data (VR 2022-04645).

References

1. Adewole, K.S., Torra, V.: DFTMicroagg: a dual-level anonymization algorithm for smart grid data. Int. J. Inf. Secur. (2022). https://doi.org/10.1007/s10207-022-00612-8.

2. Adewole, K. S., Torra, V.: Energy disaggregation risk resilience through microaggregation and discrete Fourier transform, Information Sciences 662 (2024) 120211. https://doi.org/10.1016/j.ins.2024.120211
3. Adewole, K.S., Torra, V.: Privacy issues in smart grid data: from energy disaggregation to disclosure risk, Proc. DEXA (2022) 71-84
4. Alsaid, M., Slay, T., Bulusu, N., Bass, R.B.: K-anonymity applied to the energy grid of things distributed energy resource management system. In: Proceedings of the 20th Annual International Conference on Mobile Systems, Applications and Services. pp. 581–582. Association for Computing Machinery, New York, NY, USA (2022). https://doi.org/10.1145/3498361.3538794.
5. BANES Energy Data Electricity. (2020) Bathhacked, Bath and North East Somerset, https://data.bathhacked.org/datasets/banes-energy-data-electricity
6. Castro, J. (2006) Minimum-distance controlled perturbation methods for large-scale tabular data protection, European Journal of Operational Research 171 39-52.
7. Domingo-Ferrer, J., Mateo-Sanz, J. M. (2002) Practical data-oriented microaggregation for statistical disclosure control, IEEE Trans. on Knowledge and Data Engineering 14:1 189-201.
8. Domingo-Ferrer, J., Torra, V. (2005) Ordinal, Continuous and Heterogeneous k-Anonymity Through Microaggregation, Data Mining and Knowledge Discovery 11:2 195-212.
9. Duncan, G. T., Elliot, M., Salazar, J. J. (2011) Statistical confidentiality, Springer.
10. Commission for Energy Regulation (CER). (2012). CER Smart Metering Project - Electricity Customer Behaviour Trial, 2009-2010 [dataset]. 1st Edition. Irish Social Science Data Archive. SN: 0012-00, and 0013-00. https://www.ucd.ie/issda/data/commissionforenergyregulationcer/
11. Gerlitz, C., Eriksson, A., Hansson, C.: Anonymisation score for time series consumption data. In: 27th International Conference on Electricity Distribution (CIRED 2023). pp. 428–432. Institution of Engineering and Technology, Rome, Italy (2023). https://doi.org/10.1049/icp.2023.0338.
12. Hundepool, A., Domingo-Ferrer, J., Franconi, L., Giessing, S., Nordholt, E. S., Spicer, K., de Wolf, P.-P. (2012) Statistical Disclosure Control, Wiley.
13. Laszlo, M., Mukherjee, S. (2005) Minimum spanning tree partitioning algorithm for microaggregation, IEEE Transactions on Knowledge and Data Engineering 17:7 902-911.
14. Li, N., Li, T., Venkatasubramanian, S.: t-Closeness: Privacy Beyond k-Anonymity and l-Diversity. In: 2007 IEEE 23rd International Conference on Data Engineering. pp. 106–115 (2007). https://doi.org/10.1109/ICDE.2007.367856.
15. Machanavajjhala, A., Gehrke, J., Kiefer, D., Venkitasubramanian, M. (2006) L-diversity: privacy beyond k-anonymity, Proc. of the IEEE ICDE.
16. Oganian, A., Domingo-Ferrer, J. (2000) On the Complexity of Optimal Microaggregation for Statistical Disclosure Control, Statistical J. United Nations Economic Commission for Europe, 18, 4, 345-354.
17. Samarati, P. (2001) Protecting Respondents' Identities in Microdata Release, IEEE Trans. on Knowledge and Data Engineering, 13:6 1010-1027.
18. Torra, V.: Guide to Data Privacy: Models, Technologies, Solutions. Springer International Publishing, Cham (2022). https://doi.org/10.1007/978-3-031-12837-0.
19. Torra, V., Navarro-Arribas, G.: Attribute disclosure risk for k-anonymity: the case of numerical data. Int. J. Inf. Secur. 22, 2015–2024 (2023). https://doi.org/10.1007/s10207-023-00730-x.

20. Templ, M. (2008) Statistical Disclosure Control for Microdata Using the R-Package sdcMicro, Transactions on Data Privacy 1:2 67-85.
21. Truta, T. M., Vinay, B. (2006) Privacy protection: p-sensitive k-anonymity property. Proc. 2nd Int. Workshop on Privacy Data management (PDM 2006) p. 94.