UNITED NATIONS ECONOMIC COMMISSION FOR EUROPE

CONFERENCE OF EUROPEAN STATISTICIANS

Expert Meeting on Statistical Data Confidentiality

26-28 September 2023, Wiesbaden

The Potential of Differential Privacy Applied to Detailed Statistical Tables Created Using Microdata from the Japanese Population Census

Shinsuke Ito (Chuo University, Japan)[*], Masayuki Terada (NTT DOCOMO, INC., Japan)[**],

Shunsuke Kato (Statistics Bureau of Japan)

[*] ssitoh@tamacc.chuo-u.ac.jp, [**] teradam@nttdocomo.com, [***] shun911k@gmail.com

Abstract

In numerous countries, perturbative methods are used as a privacy protection method for official statistics. The U.S. Census Bureau has studied the applicability of perturbation based on differential privacy for official statistics, and empirically investigated the mechanism of differential privacy for the publication of statistical tables created based on data from the 2020 Census. In particular, the U.S. Census Bureau has examined the applicability of differential privacy for 2010 Census data in order to create and publish statistical tables for smaller geographical areas and as a protection against “database reconstruction attacks”.

Several empirical studies on the effectiveness of perturbative methods such as additive noise, data swapping and PRAM for Japanese official microdata have been conducted (ex. Ito et al. (2018)). Other studies have investigated the possibility of adapting differential privacy for detailed geographical data from the Japanese Population Census, and examined the potential of differential privacy as an anonymization method for Japanese statistical data (Ito and Terada (2019) and Ito et al. (2020)).

When discussing future directions for the creation and publication of statistical tables, it is important to consider the potential of differential privacy. Towards this objective, this paper conducts a comparative study of the effectiveness of differential privacy for Japanese Population Census data while taking into account the actual situation regarding the application of differential privacy to official statistics in other countries. Specifically, this research conducts a comparative analysis of data usability for various differential privacy methods (with PRAM as a traditional disclosure limitation method) for statistical tables at different geographical levels created using individual data from the 2015 Japanese Population Census.

1 Introduction

Recent international trends in privacy-protecting techniques applied to official statistics include the active use of perturbative methods. For example, the U.S. Census Bureau (hereinafter, "Census Bureau") has investigated the applicability of perturbative methods based on the methodology of differential privacy, which was originally developed in the field of computer science. In particular, the Census Bureau has examined ways to create statistical tables that use differential privacy for the 2020 United States Census. In addition, for creating and publishing tract/block-level Census tables, the Census Bureau has explored the practicality of using differential privacy as a way to prevent "database reconstruction attacks" (Abowd (2018), Garfinkel et al. (2019), Garfinkel (2022)), in which perpetrators attempt to identify personal information by combining multiple published statistical tables.

These international trends suggest that the application of differential privacy techniques is potentially effective not only in creating and publishing statistical tables, but also in constructing synthetic data.

Several empirical studies on the effectiveness of perturbative methods such as additive noise, data swapping and PRAM for Japanese official microdata have been conducted in Japan (ex. Ito et al. (2018)). Other studies have investigated the possibility of adapting differential privacy for detailed geographical data from the Japanese Population Census, and examined the potential of differential privacy as an anonymization method for Japanese statistical data (Ito and Terada (2019) and Ito et al. (2020)). Therefore, investigating the applicability of differential privacy techniques to Japanese official statistics is worthwhile not only from the standpoint of discussing the future creation and publication of official statistical tables, but also the future direction of secondary use of official statistics.

In this context, this paper explores the applicability of differential privacy to data from the Japanese Population Census (hereinafter, "Population Census") not only by empirically demonstrating the characteristics of data obtained under various perturbative methods, but also by comparing and examining the utility of these methods for Japanese official statistical data.

2 Application of Differential Privacy to Census Statistics

2.1 Definition and Interpretation of Differential Privacy

Differential privacy is a privacy protection framework aimed at achieving comprehensive (ad omnia) data security against arbitrary attacks including unknown attacks. A consistent safety index ($\varepsilon \geq 0$) is quantitatively provided for various privacy protection methods (Dwork, 2007)[1]. As the value of the index becomes lower, the level of privacy protection becomes higher.

If privacy loss from a privacy protection method M is guaranteed to be less than or equal to ε, M is said to satisfy ε-differential privacy, which is defined more rigorously below.

Definition 1: For any adjacent databases D_1 and D_2 ($D_1, D_2 \in D$), the randomization function M satisfies ε-differential privacy if $M: D \to R$ satisfies the following inequality, where S is any subspace of the output space R of M (S⊆R).

$$Pr[M(D_1) \in S] \leq e^{\varepsilon} \cdot Pr[M(D_2) \in S].$$

An intuitive interpretation of this definition is that if the result of applying M to database D_1 containing data on individual A is indistinguishable from the result of applying M to database D_2 not containing data on the individual, then the output of M does not violate privacy of individual A; and the lower the value of ε, the more indistinguishable the former result is from the latter, meaning that greater security is provided by the privacy protection method M. Put differently, ε is an index indicating how much privacy is lost due to the output of M. For this reason, ε is also referred to as privacy loss or the privacy loss budget.

As Definition 1 shows, differential privacy does not refer to a certain privacy protection method, but rather is a framework for defining the level of data security provided. The specific method that protects privacy (M in Definition 1) based on differential privacy is called a mechanism. The Laplace mechanism (described below) is a typical mechanism for achieving differential privacy. Some traditional statistical disclosure control (SDC)

[1] For a deterministic method, ε →∞ (no security provided).

methods (anonymization methods), such as PRAM (Post RAndomization Methods), are also known to provide data security based on differential privacy.

2.2 Challenges in Applying Differential Privacy to Japanese Population Census Data

There are various types of mechanisms that achieve differential privacy, and their suitability for a given dataset depends on the nature of the dataset, the type of desired data output (the type of queries), and other factors. Therefore, not all mechanisms are suitable for statistical tables from the Japanese Population Census. Also, even if some mechanisms are suitable, using the wrong mechanism would significantly reduce the utility of the output statistics.

As an example, applying the Laplace mechanism - which is a well-known mechanism that satisfies differential privacy - to a contingency table is quite easy. Specifically, a random number from the Laplace distribution centered on zero (Laplace noise) can be added independently to each cell (even if its value is 0) in the contingency table[2].

It is known, however, that simply applying the Laplace mechanism to a large-scale contingency table, such as those from the Population Census, causes the practical issues listed below and reduces the utility of statistics (Terada et al. (2015), Ito and Terada (2020)).

1. Deviation from the nonnegative constraint: After applying the Laplace mechanism, the output may contain many negative values that are unlikely to be observed in actual census data.
2. Loss of sparseness: Since almost all cells that initially had unstructured zeros receive non-zero values, applying the mechanism to large sparse data, such as a Population Census dataset, significantly increases the volume of output data.
3. Loss of accuracy in partial sums: Since noise is added to the value of each cell, the error resulting from summing the values of multiple cells (e.g., to obtain the total population of a certain region) becomes large, reducing the accuracy of the partial sum.

The Laplace mechanism satisfies differential privacy by adding a Laplace noise to each cell as described above. Since this noise can take a negative value (half of the population receives negative values), if the Laplace mechanism is applied to cells for population size with a value of zero or a small value, the cells can end up with negative values. Such deviation from the nonnegative constraint makes the data unnatural, and also means that conventional analytical methods and tools, in which nonnegative input values are assumed, cannot be used. Therefore, violation of the nonnegative constraint is hard to accept from a practical standpoint. This can be easily remedied by simply replacing negative cell values with zero. This operation also increases the number of cells with zeros and thereby alleviates the problem of increased data volume. However, the simple “zeroing out” of negative values causes a large error (overbias) in partial sums.

Contingency tables with detailed geographical divisions, such as small area statistics from the Population Census, often contain many zero-valued cells (i.e., they are highly sparse). Since sparse data are often expressed in a form that omits zero-valued cells for the purpose of computation from the standpoint of memory efficiency, data volume is roughly proportional to the number of cells with non-zero values (and not to the raw number of cells). However, probabilistically the noise to be added by the Laplace mechanism is almost never zero, and all values in the output of the Laplace mechanism are most likely to be non-zero, which significantly increases the volume of the data.

Regarding the problem of lost accuracy in partial sums, in practical applications of small area statistics, such as trade area analysis, what is important is the sum of the values of multiple cells included in the scope set for a given analysis (e.g., the trade area of a retailer), but not the value in a cell representing the smallest geographical unit. With the Laplace mechanism, the noise added to a partial sum is the sum of the noise added to the values of the relevant cells, and its variance increases as the number of relevant cells increases. In other words, as the range of the cells used for a partial sum expands, the error of the partial sum increases.

For statistical tables consisting of integers only, such as those from the Population Census, the geometric mechanism, which is based on random numbers from a two-sided geometric distribution (a discretized form of the Laplace distribution), can be used instead of the Laplace mechanism. While the geometric mechanism has the feature that its output always consists of integers, its other properties are roughly the same as those of the Laplace mechanism, and the above discussions directly apply to the geometric mechanism.

[2] The Laplace distribution is also called the double exponential distribution. The scale of the distribution depends on the value of the privacy loss budget ε and the so-called global sensitivity, the value of which is determined by the query type.

2.3 Methods for Achieving Differential Privacy Applicable to the Japanese Population Census

As an approach to solving the problems discussed above, a method based on wavelet transform has been shown to be useful for mesh population statistics. The Privelet method (Xiao et al. (2011)) introduces the Haar wavelet transform in the process of noise injection so that the noises of neighboring cells offset one another, and thereby increases the accuracy of the partial sum for a continuous domain. However, the price of using the Privelet method is that it requires injection of larger amounts of noise than the Laplace mechanism. In addition, the Privelet method solves neither the problem of deviation from the nonnegative constraint nor the problem of loss of sparseness. Though the technique can solve the former problem by zeroing out the cells with negative values as in the case of the Laplace mechanism, the technique is still subject to the overbias problem.

Terada et al. (2015) propose a method based on the Morton order mapping and the Wavelet transform with nonnegative refinement (hereinafter, "nonnegative wavelet method"). In this method, noise is injected over the wavelet space as in the Privelet method. For two-dimensional data such as mesh statistics, the wavelet transform is applied after conversion to one-dimensional data via Morton order mapping (a type of locality-preserving mapping), which prevents an increase in the magnitude of noise associated with the use of multidimensional wavelets. Also, by applying an inverse wavelet transform while correcting coefficient values to prevent the output from deviating from the nonnegative constraint (i.e., while performing nonnegative refinement), the nonnegative wavelet method produces population data that satisfies the nonnegative constraint and guarantees differential privacy. Population data obtained through this method has the characteristics that the accuracy of partial sums can be controlled based on the properties of the wavelet transform, and that the sparseness of data can be restored in the process of nonnegative refinement. In other words, this method can be expected to simultaneously solve the three aforementioned problems.

An empirical experiment (Ito and Terada (2020)), in which the nonnegative wavelet method is applied to mesh statistics from the 2010 Population Census, indeed shows that the nonnegative wavelet method solves the three problems discussed above. Also, Ito et al. (2020) show that for mesh statistics, the nonnegative wavelet method is more useful than the top-down construction approach with constrained optimization, which is discussed below. For these reasons, the nonnegative wavelet method is considered an effective method for applying differential privacy to mesh statistics from the Population Census. However, it is not evident how it could be applied to other statistical tables.

Another method is based on constrained optimization. Specifically, noise is injected via the Laplace mechanism or a geometric mechanism, and optimization is performed subject to a total-number constraint and nonnegative constraint; and the solution becomes the output.

Lee et al. (2015) propose an algorithm for such constrained optimization that uses ADMM (alternating direction method of multipliers) which is a type of numerical optimization method. For the 2020 Census, a constrained optimization method was implemented using the commercial solver Gurobi. It should be noted that constrained optimization involves high computational costs. For example, it took 21 s for the method by Lee et al. (2015) to be applied to a dataset containing 4,096 cells. For the 2020 Census, a large-scale system was constructed based on Amazon EMR (Elastic Map-Reduce, a distributed computing system) offered on Amazon Web Services, a commercial cloud computing infrastructure, which uses up to 100 high-performance machines (each with 96 virtual CPUs and 768GB of RAM).

Terada et al. (2017) present a high-speed method for large-scale data, taking advantage of the fact that an optimization problem with a total-number constraint and nonnegative constraint can be reduced to the problem of projection onto a canonical simplex in a multidimensional vector space. This method is shown to process data with 100,000 cells in 12.6 ms on a typical laptop computer at that time, which makes it a suitable method for Population Census statistics and other large-scale statistics.

Two approaches to data construction are considered in this study. In one approach, one of the aforementioned methods is applied to a statistical table from the Population Census as follows: in general, the method is applied to only the population of the smallest geographical unit (the basic unit district, in the case of the Japanese Population Census); and the resulting district-level population data are summed to obtain the population at the municipal or prefectural level in a bottom-up manner. Another approach works in a top-down, recursive manner as follows: one of the aforementioned methods is applied to the prefecture-level population data with the total national population setting the total-number constraint in order to obtain privacy-protected prefecture-level population data; privacy-protected municipality-level population data is then obtained with the prefecture-level population setting the total-number constraint; and so on. In this paper, the former is called the bottom-up data construction approach, and the latter is called the top-down data construction approach. The top-down algorithm (TDA) used for the 2020 U.S. Census is an algorithm based on the top-down data construction approach.

In both the bottom-up and top-down data construction approaches, constrained optimization satisfies the nonnegative constraint, and data sparseness is expected to be restored. In addition, the top-down data construction approach is expected to solve the problem of lost accuracy in partial sums.

Both data construction approaches are applicable not only to mesh statistics, but also to other statistics. However, there have been few empirical studies that apply either approach to non-mesh statistics from the Japanese Population Census. Therefore, the practicality and quantitative properties of these approaches are unclear. In theory, the top-down approach is expected to produce better output than the Laplace mechanism in terms of the accuracy of partial sums, but the extent of the superiority has not been quantitatively clarified.

3 Applying Differential Privacy to the 2015 Japanese Population Census Data

This study performs a comparative experiment concerning the application of differential privacy to Population Census data. This section presents the data and procedures used in the experiment and presents the experiment's results.

3.1 Data Used in the Experiment

The experiment is based on three small area statistics with different aggregation categories that were created from the 2015 Population Census individual data (full data). For each aggregate data table, the focus is on population size; the level of aggregation is at the basic unit district (minimum geographic district) level; and the three aggregation categories considered are "all", "males and females", and "males and females in 5-year age groups". In other words, the experiment is based on three aggregate data tables for (1) basic unit district-level total population, (2) basic unit district-level population by gender, and (3) basic unit district-level population by gender and 5-year age group. These tables are hereinafter referred to as Aggregation Tables 1, 2, and 3, respectively.

3.2 Experimental Methods

This empirical study aims to gain knowledge relevant to the aforementioned three research questions by implementing various differential privacy methods for Aggregation Tables 1 to Aggregation Table 3, and evaluates the utility of the data. The specific procedures for method implementation and indices for comparative evaluation are explained below.

PRAM, the Laplace mechanism, the top-down data construction method, and the bottom-up data construction method, which are discussed in Section 3, were used as methods for achieving differential privacy. While methods based on wavelet transform are effective for mesh statistics, they were excluded from consideration as they are difficult to apply to data tables with other geographical aggregation levels. As mentioned earlier, the output of the Laplace mechanism can include negative population values (which violate the nonnegative constraint). These were zeroed out as a post-processing adjustment. The constraint optimization method used with the top-down data construction approach or the bottom-up data construction approach was the one proposed by Terada et al. (2017), which are based on projection onto a canonical simplex. To apply the top-down approach to Aggregation Tables 2 and 3, the results for the (attribute-based) aggregation categories were merged. For example, in the case of applying the top-down approach to Aggregation Table 2, the results of applying the top-down approach to the male population data and the results of applying this approach to the female population data were merged into one table, and this table was treated as the result of applying the top-down approach to Aggregation Table 2.

In this experiment, the highest geographical level was not the national level, but the prefectural level. This is due to limitations of the computational environment available at the on-site data access facility used for the experiment. Therefore, instead of creating an aggregate data table for Japan as a whole and applying each method to that table, a total of 47 aggregate data tables were created for all the prefectures, each method was applied to each of those tables independently, and the resulting tables were then merged to calculate evaluation indices (discussed below)[3].

The aforementioned four methods were applied with each of the following eight values for the privacy loss budget (ε) set for the experiment: 0.1, 0.2, 0.7, 1.0, 1.1, 5, 10, and 20. The values 0.7 and 1.1 were chosen as approximations of $log_e 2$ and $log_e 3$, respectively, which are frequently used as values of the privacy loss

[3] Ideally, in the top-down data construction approach, for example, the national level should be above the prefectural level. In this experiment, however, the prefecture level is highest of the four geographical levels considered, followed by municipality level, town/village level, and basic unit district level.

budget). Though PRAM and the top-down data construction approach can be configured to allocate different privacy loss budgets to different geographical levels or attribute categories, in this experiment, each value of the privacy loss budget was evenly allocated.

Utility of the data was evaluated for population data at the basic unit district level (the most detailed population data), and for population data at higher geographical levels (partial sums) in consideration of real-world use of aggregate data. Specifically, the errors in the prefecture-level, municipality-level, town/village-level, and basic unit district-level population data were quantitatively compared. The mean absolute error (MAE) was used as an error index and was calculated for each of the three aggregation tables, the four methods applied, the eight values of the privacy loss budget, and the four geographical levels used for partial sums.

3.3 Experimental Results

Tables 1 to 3 show the evaluation results. In each table, (a) PRAM, (b) Laplace, (c) BottomUp, and (d) TopDown refer to PRAM, the Laplace mechanism (plus zeroing out of negative values), the bottom-up data construction method, and the top-down data construction method, respectively. Tables 1, 2, and 3 summarize the evaluation results for Aggregation Tables 1, 2, and 3, respectively.

It should be noted that in Table 1, the errors for the prefecture-level population data are zero for the three methods other than the Laplace mechanism. These three methods (PRAM, the bottom-up data construction method, and the top-down data construction method) have the characteristic that the total number of records in the input data is preserved in the output data (which is referred to as satisfaction of the total-number constraint). As mentioned above, the prefectural level was the highest geographical level in this experiment (due to limitations of the computational environment). Therefore, when a method that satisfies the total-number constraint is applied to Aggregation Table 1 (which does not have attribute-based aggregation categories), the errors for the population data at the highest geographical level will be zero (because of the constraint). In other words, the result that errors at the prefectural level in Table 1 are zero is attributable to the conditions of this experiment, and the same result would not be obtained if the highest geographical level were the national level (instead, errors for total population at the national level would become zero).

Attention should also be paid to the interpretation of the result for PRAM for the basic unit district level, especially for small values of the privacy loss budget. For example, in Table 3, the errors caused by PRAM to the basic unit district-level data vary very little across different values of the privacy loss budget (ε = 0.1 to 20).

The result reflects the fact that the degree of perturbation by PRAM required for a given privacy loss budget is so great that the output population data seem to follow a uniform distribution. In other words, since the basic unit district-level aggregate data table used for Table 3 is quite sparse, with most of the values being either 0 or 1, even if the output data follows a random uniform distribution, at first glance its accuracy does not appear undesirable at the basic unit district level. However, this is just a false accuracy and is not statistically meaningful. Similarly, the errors caused by PRAM at the municipality level and the town/village level shown in Table 3 reveal that the accumulation of errors greatly degrades the accuracy of the partial sums and there is significant degradation of the characteristics of the original aggregate data table.

4 Discussion

The experimental results show that the errors vary significantly across the differential privacy methods applied and also across the different values of the privacy loss budget (ε) considered. As discussed in Section 3, for a given privacy loss budget, differential privacy guarantees the same level of privacy protection regardless of the differential privacy method used, but the utility of the resulting data depends on the method and the use of the data. The results obtained in the experiment support this description. It is therefore necessary, in discussing the practicality of applying differential privacy and the utility of the output data, to use various differential privacy methods and different values of the privacy loss budget, as in this experiment, and to examine the results quantitatively.

The experimental results also show that the characteristics of the errors associated with data for the smallest geographical unit (the basic unit district) and the characteristics of the errors associated with the partial sums for larger geographical units (e.g., municipalities) are not necessarily the same. If the errors at the basic unit district

Table 1: Evaluation results for Aggregation Table 1 (basic unit district-level total population)

ε	Method	Prefecture	Municipality	Town/Village	Basic Unit District
0.1	(a)PRAM	0.00	14408.44	520.10	48.65
	(b)Laplace	98607.22	2490.96	83.50	17.60
	(c)BottomUp	0.00	855.53	72.06	17.38
	(d)TopDown	0.00	79.09	73.35	49.83
0.2	(a)PRAM	0.00	14399.53	520.26	48.64
	(b)Laplace	30844.00	817.25	39.15	9.23
	(c)BottomUp	0.00	367.38	36.86	9.21
	(d)TopDown	0.00	41.12	37.78	30.42
0.7	(a)PRAM	0.00	14401.57	520.09	48.65
	(b)Laplace	5433.01	157.62	11.06	2.75
	(c)BottomUp	0.00	97.37	10.81	2.75
	(d)TopDown	0.00	11.62	11.26	10.42
1	(a)PRAM	0.00	14406.80	520.17	48.65
	(b)Laplace	3483.00	104.64	7.65	1.92
	(c)BottomUp	0.00	65.78	7.52	1.91
	(d)TopDown	0.00	7.84	7.83	7.38
1.1	(a)PRAM	0.00	14400.56	520.19	48.64
	(b)Laplace	3117.37	94.48	6.96	1.74
	(c)BottomUp	0.00	57.97	6.84	1.74
	(d)TopDown	0.00	7.13	7.12	6.73
5	(a)PRAM	0.00	14351.58	518.16	48.47
	(b)Laplace	609.34	19.08	1.54	0.39
	(c)BottomUp	0.00	13.10	1.51	0.38
	(d)TopDown	0.00	1.60	1.57	1.52
10	(a)PRAM	0.00	10215.78	359.65	34.59
	(b)Laplace	314.63	9.65	0.77	0.19
	(c)BottomUp	0.00	6.43	0.76	0.19
	(d)TopDown	0.00	0.84	0.79	0.76
20	(a)PRAM	0.00	3.53	0.23	0.02
	(b)Laplace	152.12	4.80	0.38	0.10
	(c)BottomUp	0.00	3.15	0.38	0.10
	(d)TopDown	0.00	0.42	0.39	0.38
100	(a)PRAM	0.00	0.00	0.00	0.00
	(b)Laplace	30.92	0.96	0.08	0.02
	(c)BottomUp	0.00	0.64	0.08	0.02
	(d)TopDown	0.00	0.08	0.08	0.08

Table 2: Evaluation results for Aggregation Table 2

(basic unit district-level total population by gender)

ε	Method	Prefecture	Municipality	Town/Village	Basic Unit District
0.1	(a)PRAM	35301.70	7277.54	261.95	24.80
	(b)Laplace	158482.08	3940.62	96.57	16.10
	(c)BottomUp	4800.07	977.40	68.02	15.55
	(d)TopDown	60.08	79.80	69.48	36.40
0.2	(a)PRAM	34432.49	7273.24	261.97	24.81
	(b)Laplace	50430.20	1271.48	41.92	8.77
	(c)BottomUp	2081.00	443.80	36.11	8.68
	(d)TopDown	24.09	39.48	36.56	24.83
0.7	(a)PRAM	29972.11	7260.04	261.72	24.79
	(b)Laplace	7016.17	193.09	11.17	2.72
	(c)BottomUp	432.39	100.55	10.83	2.71
	(d)TopDown	9.27	11.83	11.16	9.67
1	(a)PRAM	27463.17	7255.24	261.68	24.80
	(b)Laplace	4239.02	120.71	7.69	1.90
	(c)BottomUp	310.38	68.29	7.50	1.89
	(d)TopDown	7.14	8.10	7.77	7.00
1.1	(a)PRAM	26432.21	7247.65	261.63	24.80
	(b)Laplace	3728.69	107.08	7.00	1.73
	(c)BottomUp	271.13	61.45	6.84	1.73
	(d)TopDown	5.63	7.30	7.07	6.43
5	(a)PRAM	5492.60	7215.53	261.27	24.78
	(b)Laplace	653.28	19.92	1.54	0.38
	(c)BottomUp	59.69	12.58	1.51	0.38
	(d)TopDown	1.19	1.58	1.57	1.51
10	(a)PRAM	530.02	7197.73	260.39	24.70
	(b)Laplace	315.57	9.95	0.77	0.19
	(c)BottomUp	26.27	6.56	0.76	0.19
	(d)TopDown	0.54	0.82	0.78	0.76
20	(a)PRAM	7.51	5115.43	180.82	17.74
	(b)Laplace	159.90	4.92	0.39	0.10
	(c)BottomUp	14.20	3.23	0.38	0.10
	(d)TopDown	0.25	0.40	0.39	0.38
100	(a)PRAM	0.00	0.00	0.00	0.00
	(b)Laplace	32.34	1.00	0.08	0.02
	(c)BottomUp	2.49	0.64	0.08	0.02
	(d)TopDown	0.08	0.08	0.08	0.08

Table 3: Evaluation results for Aggregation Table 3

(basic unit district-level total population by gender and 5-year age group)

ε	Method	Prefecture	Municipality	Town/Village	Basic Unit District
0.1	(a)PRAM	15899.43	587.27	18.50	2.05
	(b)Laplace	376314.96	9323.57	173.38	10.77
	(c)BottomUp	14008.77	544.73	25.62	3.23
	(d)TopDown	81.64	76.50	30.75	3.44
0.2	(a)PRAM	15874.93	586.70	18.49	2.05
	(b)Laplace	175973.88	4359.93	81.69	5.69
	(c)BottomUp	11884.21	447.52	19.48	2.80
	(d)TopDown	41.25	39.09	20.72	3.28
0.7	(a)PRAM	15703.82	584.13	18.46	2.05
	(b)Laplace	40261.21	997.68	19.59	1.90
	(c)BottomUp	5618.71	203.17	8.85	1.55
	(d)TopDown	11.51	11.42	8.38	2.66
1	(a)PRAM	15573.72	582.26	18.44	2.05
	(b)Laplace	25349.47	628.32	12.71	1.38
	(c)BottomUp	4077.82	146.72	6.56	1.20
	(d)TopDown	7.94	7.93	6.20	2.37
1.1	(a)PRAM	15540.75	581.12	18.43	2.05
	(b)Laplace	22484.80	557.35	11.37	1.27
	(c)BottomUp	3725.80	133.89	6.05	1.12
	(d)TopDown	7.26	7.28	5.73	2.29
5	(a)PRAM	12827.43	541.94	17.99	2.04
	(b)Laplace	3506.22	87.20	2.10	0.31
	(c)BottomUp	795.49	28.69	1.50	0.30
	(d)TopDown	1.61	1.60	1.45	0.96
10	(a)PRAM	6345.74	469.39	17.20	2.02
	(b)Laplace	1684.23	41.92	1.03	0.16
	(c)BottomUp	395.63	14.25	0.76	0.15
	(d)TopDown	0.77	0.80	0.74	0.55
20	(a)PRAM	356.10	433.32	16.58	1.98
	(b)Laplace	839.77	20.90	0.52	0.08
	(c)BottomUp	197.79	7.14	0.38	0.08
	(d)TopDown	0.40	0.40	0.38	0.29
100	(a)PRAM	0.00	0.00	0.00	0.00
	(b)Laplace	167.76	4.17	0.10	0.02
	(c)BottomUp	39.47	1.42	0.08	0.02
	(d)TopDown	0.08	0.08	0.08	0.06

level are taken as indices of the utility of the relevant data, then the output data from the bottom-up data construction method and the Laplace mechanism seem superior[4]. However, for partial sums at the municipality level and town/village level, the errors tend to be larger for both the bottom-up method and the Laplace mechanism, and the tendency is particularly noticeable with the Laplace mechanism. The reason is that zeroing out negative values to satisfy the nonnegative constraint in the experimental implementation of the Laplace mechanism introduces small positive biases throughout all of the output data[5]. These biases rarely show up as large errors for the basic unit district-level output data; however, they can cause serious overestimations in the output data for higher geographical levels.

The top-down data construction method is inferior to the bottom-up data construction method in terms of the errors at the basic unit district level. However, for the top-down method, the errors are not accumulated at higher geographical levels. In other words, the errors are largely unchanged across different geographical levels and indicate high levels of data utility. Specifically, the results for a privacy loss budget ($\varepsilon = 20$), which is close to the value used for the 2020 U.S. Census, show that for every type of output data based on a single variable or a combination of variables, the errors are smaller than 1 for the partial sums at all geographical levels, which indicates relatively high levels of data utility. However, in comparing the output data, it should be noted that the number of cells in an aggregate data table increases as the number of variables increases, or as more detailed categories are used, and that the errors for partial sums based on an aggregate data table with a large number of cells will tend to be large.

Comparing PRAM with the other methods from this viewpoint shows that, under a given set of conditions, in most cases PRAM is significantly inferior in terms of privacy protection efficiency. For small values of the privacy loss budget, the results of applying PRAM at the basic unit district level seem to be superior to the results of other methods. However, as discussed in Section 3, this is attributable to false accuracy and does not have real meaning.

In addition, while data utility associated with the methods other than PRAM tend to improve as the value of ε increases, this tendency is hardly seen for PRAM. This result implies that the privacy protection efficiency of PRAM is far worse than that of the other methods. Perturbation in PRAM is performed as follows: values of a certain attribute in individual data remain unchanged with a given probability ρ or randomly changed with probability 1-ρ. If PRAM achieves differential privacy, the value of ρ is determined as a function of ε. However, the actual calculated value of ρ is often close to zero unless the value of ε is very large. In other words, even if privacy protection is sacrificed by moderately increasing the value of ε, ρ remains close to zero, which does not lead to improved data utility. This suggests that PRAM can achieve differential privacy, but entails quite low privacy protection efficiency, and that the use of other methods should be considered.

Then, what method is suitable for applying differential privacy to Japanese Population Census statistics? As mentioned before, satisfying the nonnegative constraint is a problem for a simple Laplace mechanism. Even if an attempt is made to satisfy the constraint by zeroing out negative values as in this experiment, it is still difficult to obtain a practically usable aggregate data table because of the large overestimation bias affecting partial sums. Also, PRAM clearly fails to achieve both a reasonable level of privacy protection and data utility.

Regarding the bottom-up data construction method and the top-down data construction method, the errors at the basic unit district level show that the bottom-up method provides higher data utility, but its errors for partial sums increase as the range of cells used for the partial sums expands, which indicates decreasing data utility. In contrast, for the top-down method the errors for partial sums remain small. Therefore, when partial sums are calculated for a higher geographical level, the degree of data utility is maintained.

In summary, judging from the errors for the output data at the basic unit district level, data utility is relatively high for the bottom-up method, but data utility deteriorates for partial sums since the errors associated with them tend to increase significantly. For the top-down method, errors at the basic unit district level are larger than those for the bottom-up method, and if the level of privacy protection is properly set, the utility of different output data considered in this study is maintained, even when the effect on partial sums is taken into account.

[4] In Table 2, PRAM seems to be the best for small values of ε, but this observation is attributed to a false accuracy and does not bear real meaning, as discussed in the previous section.

[5] In other words, without zeroing-out of negative values, errors would not expand significantly, but, at the same time, the nonnegative constraint would not be satisfied. For this reason, the resulting tables could contain a large number of negative population values.

It should be noted, however, that the variables used in this empirical experiment are limited, and that only the errors in output data, including partial sums, in relevant cells are used as indices for evaluating data utility and no other statistics are used for evaluation purposes. Further investigation into the applicability of various methods to the Population Census data should be carried out using a variety of statistics for analysis.

5 Conclusion

To explore the applicability of differential privacy to Population Census statistics, this paper evaluated the utility of statistical tables for different geographical levels which were created using individual data from the Population Census and by applying various differential privacy methods. In addition, a comparison was made with the conventional anonymization method PRAM. The results of this study show that in applying differential privacy to Japanese Population Census data, the top-down data construction method yields a higher level of data utility than the other methods. As described above, the Census Bureau has adopted TDA for creating and publishing detailed statistical tables for the 2020 Census. The results of this study are therefore consistent with the Census Bureau's choice of technique used in its statistical work. This suggests that given a hierarchical geographical structure, reasonable results from the standpoint of data utility can be obtained by top-down, consistent allocation of the noise generated based on differential privacy to the cells of a statistical table (rather than injecting noises to the cells at the basic unit district level and performing aggregation for a higher geographical level by summing the values of cells in a town/village-level table).

This experimental study was the first attempt to apply differential privacy to detailed statistical tables (a basic unit district-level statistical table) from the Japanese Population Census. This study's investigation was based on basic unit district-level cross tables with age and gender categories. Since the Population Census collects demographic data (beyond age and gender), employment data (including employment status, industry, and occupation), and residential data, future studies can potentially examine the data utility resulting from applying a differential privacy method to statistical tables created with these other variables used for aggregation categories. Our future research agenda also includes further investigation into the effectiveness of differential privacy based on aggregate data tables created with various Population Census variables.

6 Note

This paper is the revised version of Ito et al. (2023) which were published in Japanese.

7 References

Abowd, J. M. (2018) Staring-down the database reconstruction theorem, Joint Statistical Meetings, Vancouver, BC, Canada.

Dwork, C. (2007) "An Ad Omnia Approach to Defining and Achieving Private Data Analysis", *Proc. 1st intl. conf. Privacy, security, and trust in KDD*, pp. 1-13.

Garfinkel, S. Abowd, J. M., and Martindale, C. (2019) "Understanding Database Reconstruction Attack in Public Data", Communications of the ACM, Vol. 62 No. 3, ACM, pp. 46-53.

Garfinkel, S. (2022) "Differential Privacy and the 2020 US Census", MIT Case Studies in Social and Ethical Responsibilities of Computing, Winter 2022.

Ito, S., Yoshitake, T., Kikuchi, R., Akutsu, F. (2018) "Comparative Study of the Effectiveness of Perturbative Methods for Creating Official Microdata in Japan", Josep Domingo-Ferrer and Francisco Montes (eds.) *Privacy in Statistical Databases: UNESCO Chair in Data Privacy, International Conference, PSD 2018, Valencia, Spain, September 26–28, 2018, Proceedings (Lecture Notes in Computer Science)*, Springer, pp.200-214.

Ito, S. and Terada, M. (2019) The potential of anonymization method for creating detailed geographical data in Japan, Joint UNECE/Eurostat Work Session on Statistical Data Confidentiality, pp. 1–14.

Ito, S. and Terada, M. (2020) "An Evaluation of Anonymization Methods for Creating Detailed Geographical Data" (in Japanese), *Journal of the Japan Statistical Society*, Vol. 50, No. 1, pp.139-166.

Ito, S., Miura, T., Akatsuka, H., and Terada, M. (2020) “Differential Privacy and Its Applicability for Official Statistics in Japan – A Comparative Study Using Small Area Data from the Japanese Population Census”, Josep Domingo-Ferrer and Krishnamurty Muralidhar(eds.) *Privacy in Statistical Databases: UNESCO Chair in Data Privacy, International Conference, PSD 2020, Tarragona, Spain, September 23–25, 2020, Proceedings (Lecture Notes in Computer Science)*, Springer, pp.337-352.

Ito, S., Terada, M., Kato, S. (2023) “An Empirical Study on the Effectiveness of Perturbative Methods Applied to Japanese Population Census Data” (in Japanese) *Research Paper*, No.58, pp.1-26.

Lee, J., Wang, Y., and Kifer, D. (2015) “Maximum Likelihood Postprocessing for Differential Privacy under Consistency Constraints.” Proc. 21st ACM SIGKDD intl. conf. Knowledge Discovery and Data Mining (KDD ’15), pp. 635–644.

Terada, M., Suzuki, R., Yamaguchi, T., and Hongo, S. (2015) “On Publishing Large Tabular Data with Differential Privacy” (in Japanese), *Transactions of the Information Processing Society of Japan*, Vol. 56, No. 9, pp. 1801-1816.

Terada, M., Yamaguchi, T., and Hongo, S. (2017), “On Releasing Anonymized Microdata with Differential Privacy” (in Japanese), *Transactions of the Information Processing Society of Japan*, Vol. 58, No. 9, pp. 1483-1500.

Xiao, X., Wang, G., Gehrke, J. and Jefferson, T. (2011). Differential Privacy via Wavelet Transforms, *IEEE Transactions on Knowledge and Data Engineering*, 23(8), 1200–1214.