

RWN: A Novel Neighborhood-Based Method for Statistical Disclosure Control

Noah Perry*, Norman Matloff**, Patrick Tendick***

*Department of Statistics, University of California, Davis

**Department of Computer Science, University of California, Davis. Corresponding author.

***Ingredion, Inc.

E-mail: njperry@ucdavis.edu, nsmatloff@ucdavis.edu

Abstract. A novel variation of the data swapping approach to statistical disclosure control is presented, aimed particularly at preservation of multivariate relations in the original dataset. A theorem is proved in support of the method, and extensive empirical investigation is reported.

Keywords.

1 Introduction

The field of statistical disclosure control (SDC)—maintaining privacy of individual records in a dataset while retaining the statistical utility of the data—has long been the subject of arcane technical analysis, conducted mainly by statisticians. The advent of *differential privacy* (DP) methods in 2006 [7] brought computer scientists into the field, and the arcane nature of the field changed with the well-publicized adoption of DP by the United States Census Bureau in 2017. That move by the Bureau has become somewhat controversial [3] [20] and though we do not address that controversy, we note the salutary effect of greatly increasing public awareness of SDC.

Different SDC tools may be appropriate for different databases in different settings, not just in terms of numeric degree of protection afforded by a tool, but also in terms of usability, interpretability and transparency for end users [19] [11]. Here we develop new methodology that we believe database administrators (DBAs) will find useful in a variety of settings.

Our proposed method, Randomization within Neighborhoods (RWN), to be presented below, is inspired by data swapping, a classic approach to SDC. However, RWN differs from previous data swapping approaches, in that it exploits a certain statistical independence property, to be described in detail in the next section.

This paper is organized as follows. We give a brief overview of SDC methods in Section 2, followed by a discussion of considerations in SDC specific to databases used for statistical analysis in Sections 3 and 4. The RWN method is presented in Section 5. The underlying theory is given in Section 6. Tuning parameter selection is covered in Section 7. Our empirical investigation is discussed in Section 8, and computational issues are discussed in Section 9.

Note by the way that while we refer for convenience at some points to census data as a concrete example, our methodology is meant to be general, not specific to the census. It could be applied to medical data, employee data and so on.

2 SDC Methods

Good surveys of SDC methods are in [6] [12] [4]. In order to contrast with RWN, we give a brief overview here.

Previously the Census Bureau had used the popular *data swapping* method for SDC [10] [6] [9]. A set of *key variables* is defined that may render certain individual records in the data vulnerable to disclosure. Some records, especially those deemed most at risk to disclosure, will have the values of their key variables swapped with those in other records, say drawn from the same geographic region.

There are variants, notably *data shuffling* [18]. These methods are typically applied one variable at a time, thus creating the concern of attenuation of multivariate relations.

The Bureau has also used *cell suppression*, in which any query concerning a very small number of database records is denied. However, in 2017, the Bureau, after performing various simulations, decided that the swapping approach was in danger of reconstruction attacks and turned to DP.

Another standard SDC approach is *data perturbation*, in which random noise is added to achieve privacy. The Census Bureau has used this method in the past as well,

DP is also a perturbation method. Its novelty, though, is in its ability to be able to quantify the degree of privacy, in a manner having certain mathematical traits, such as composability. Another difference from classic noise addition methods is that in DP, the noise is typically added to a statistic of interest, say a mean (*global DP*), rather than to the microdata itself (*local DP*).

Two of the present authors, NM and PT, have a data perturbation background background [16] [14] [23] [22], and were interested in modernizing that approach, making a proposal in 2016 [15] for a novel SDC method inspired by data swapping. The present work, joined by author NP, develops those ideas.

3 Statistical Views and Goals

Leo Breiman’s famous essay [5] on predictive modeling described a “cultural” difference between researchers in statistics and computer science, the former viewing the world more in terms of probabilistic behavior, the latter in terms of algorithms. Curiously, as pointed out by statistician Larry Wasserman [24], a similar difference later arose in the SDC field. Referring to the need in (global) DP to develop a separate DP-compliant method for each statistical procedure used (mean, regression analysis etc.), he noted the contrasting views:

CS view: Receive a query for a [specific statistical procedure], return a private answer.

Statistics view: Give me data. Then I can: draw plots, fit models, test fit, estimate parameters, make predictions ...

An amusement park metaphor will be useful. Under the CS approach, one must buy a separate ticket for each ride. Some rides won’t be available at all, pending development of tickets tailored to those rides. With statistics, one purchases a day pass, good for all rides. Our focus here is on settings in which one wishes to have a “statistical day pass.” We protect the data in some way, say perturbation, then let users conduct whatever types of statistical analyses they wish in an open-ended way.

Good arguments can be made for either view. As noted, the one-query-at-a-time nature of the DP/CS approach enables the setting of precise guarantees of privacy, which would be difficult or impossible in the statistics approach. On the other hand, this means the user is restricted to only the types of queries for which a DP version has already been developed and implemented. Linear regression may be available, say, but not quantile regression.

(We note in passing that unlike most applications of DP, the US Census Bureau’s DP methodology does amount to a “day pass.” This is because they add noise to the raw data, which are cell counts in a huge contingency table, rather than to the output of a statistical analysis, such as a mean.)

Due to our goal of developing a statistical *analysis* “day pass,” in the sense Wasserman described—estimating parameters, fitting models and so on, open-endedly—we look at the data in the usual *statistical* manner, i.e. as a sample from a population. This is in contrast to many SDC applications in which the data themselves are of primary interest, say the total count of people in a given income range for a given census block.

A typical example might be that of a medical database, in which the privacy of individual patients is required, but with which medical researchers can still conduct statistical analyses, making population inferences.

We suppose here that there is some population value θ for which we wish to obtain a sample estimate $\hat{\theta}$, performing statistical operations such as inference (confidence intervals, hypothesis tests). These operations will be conducted on the perturbed data obtained by applying RWN to our original microdata. Typically θ will be vector valued, such as a vector of regression coefficients.

4 Desirable Statistical “Day Pass” Characteristics

In developing a new SDC procedure, such as our proposed RWN, these goals are key:

- Ability to handle mixed continuous and discrete/categorical data.
- Preservation, to the degree possible, of not only univariate but also multivariate distribution-
s/relations.
- Limiting the increase in size of the standard errors of $\hat{\theta}$.
- Preservation, to the degree possible of statistical inference levels related to $\hat{\theta}$.

On the other hand, as noted, we do not take as a goal the preservation of marginal totals as in Census data.

Let’s elaborate a bit on these goals.

4.1 Handling Mixed Continuous and Categorical Data

A major obstacle to data perturbation methods is their inability to handle discrete/categorical data. Consider a variable such as Number of Children in Family. After noise addition, a value may become negative, an unacceptable situation. A similar difficulty arises with categorical variables, after they are converted to dummy (*one-hot*) form.

Indeed, the vast majority of the Census Bureau’s TopDown algorithm [1] is devoted to making adjustments to negative values, and satisfying certain constraints involving marginal totals.

RWN will be seen to handle mixed continuous and discrete/categorical data in a simple, natural manner.

4.2 Preservation of Multivariate Relations

Absent some compensating feature, any change to the data arising from applying a “day pass” SDC procedure, say perturbation or swapping, will result in distortions of the relations between variables

in the data. This will also occur with cell suppression methods. Since analysis of multivariate relations comprise the very core of statistics, we take as a major goal at least approximately preserving such relations.

We are of course willing to let those relations be one aspect of the utility/privacy tradeoff that is necessary to any disclosure avoidance technique. Let's call this property Multivariate Relations Attenuation Resistance (MRAR).

The goal then is to develop an SDC method that includes MRAR, with the method providing the user a "lever" that she can use to choose her desired utility/privacy tradeoff level.

Comparatively little work in the SDC field has focused on MRAR. It is mentioned only briefly in [12] and [6] — and no wonder, as MRAR is a challenging condition to meet.

Consider noise addition methods. One actually can preserve second-order moment structure by setting the covariance matrix of the noise to that of the data [13] [16] [22]. But higher-order moments are lost and other distortions can occur. And there are no obvious techniques for extending this property with noise addition in mixed continuous/categorical variable settings.

5 RWN: Randomization within Neighborhoods

The method works roughly as follows. For each record in the data, we define a neighborhood using either a Euclidean distance-based radius or k -nearest neighbors. Then, for each record r we randomly choose a subset of the variables to perturb. For each such variable, we replace its original value by its counterpart in a randomly chosen record in the neighborhood of r . A key point is that a different random neighbor record may be used for each of the variables to be perturbed in r .

More formally:

Let $W = (w_{ij}), i = 1, \dots, n, j = 1, \dots, p$ denote our original data on n individuals and p variables and $W' = (w'_{ij}), i = 1, \dots, n, j = 1, \dots, p$ be the released (i.e. perturbed) data.

Choose neighborhood radius $\epsilon > 0$, number of nearest neighbors k , and modification probability q . Then we form our released data W' as follows:

For $i = 1, \dots, n$:

1. Consider record i in the database:

$$r_i = (w_{i1}, \dots, w_{ip}) \quad (1)$$

2. Find the set S_i of records within the neighborhood of r_i other than r_i itself. Each neighborhood is defined to be either the k -nearest neighbors of r_i or the set of neighbors within ϵ distance of r_i , whichever set is larger.

3. For $j = 1, \dots, p$:

With probability $1 - q$, leave w_{ij} unmodified, but with probability q , modify it. For variables that are chosen to be modified, we replace w_{ij} with the value in variable j of a random record in the neighborhood S_i (sampled from S_i with replacement).

As noted, there may be a different such random record for each j . This results in a perturbed data point w'_{ij} . For unmodified variables, $w'_{ij} = w_{ij}$.

4. Store the released, modified version of r_i as

$$r'_i = (w'_{i1}, \dots, w'_{ip}) \quad (2)$$

A key point is that in Step 3, the p actions here are *taken independently of each other*. In other words, the process acts as if the p variables in the data are statistically independent of each other. This would at first seem to violate our goal of MRAR, but it is all resolved in the theorem in Section 6 below.

We call this technique Randomization Within Neighborhoods (RWN).

5.1 Comparison to Rank Swapping

Rank swapping [17] does data swapping using ranks rather than data values, in order to facilitate dealing with discrete variables. It's implemented as the function `rankSwap()` in the popular R package for SDC, `sdcMicro` [21]. It too is neighborhood-based. RWN differs from rank swapping in several respects:

- Rank swapping's MRAR feature is focused on bivariate relations, while RWN is fully multivariate, exploiting the theorem in Section 6.
- RWN's neighborhoods are more complex than rank swap's swapping ranges due the following features:
 - The size of the RWN neighborhoods can differ for each record.
 - The distances are computed based on all variables in a record instead of one variable at a time, so the records within a given neighborhood can be "similar" in many different ways.
 - RWN duplicates data values, rather than actually swapping them. This actually helps reduce *k-anonymity* issues [21].
 - Setting $q < 1$ allows for the possibility that some values may be the original ones.

6 Rationale and Theoretical Basis

Since we are using neighborhoods, one might ask, "Why not just replace entire data rows—a given row is replaced by a neighboring row—rather than do replacement component by component, taking different components from different neighboring rows?" That would achieve our MRAR goal, but at the expense to too much increase in $Var(\hat{\theta})$.

By independently perturbing each element of a row, we can limit the increase in standard errors of $\hat{\theta}$. However, we then must ask whether our method has the all-important MRAR property. The following theorem shows that it does.

Let f_X be a density function for the p -variate vector $X = (X_1, \dots, X_p)$. Consider the conditional distribution of X , given that X is in a small neighborhood of t . Then the X_i are approximately independent in this distribution.

For expositional convenience, the theorem and proof will be stated for the case $p = 2$.

Theorem: Consider a bivariate random vector (X, Y) having a joint density, and set $\epsilon > 0$. For any t in \mathcal{R}^2 , let $A_{t,\epsilon}$ denote the ϵ neighborhood of t . Let F denote the joint cdf of (X, Y) . Given $(X, Y) = t$, define $G_{t,\epsilon}$ to be the cdf of (X, Y) , given that that vector is in $A_{t,\epsilon}$. Finally, given $(X, Y) = t$, define independent random variables U_ϵ and V_ϵ to be drawn randomly from the first- and second-coordinate marginal distributions of $G_{t,\epsilon}$, respectively. Then

$$\lim_{\epsilon \rightarrow 0} P(U_\epsilon \leq a \text{ and } V_\epsilon \leq b) = F(a, b) \quad (3)$$

for all $-\infty < a, b < \infty$.

In other words, as ϵ goes to 0, the bivariate distribution of (U_ϵ, V_ϵ) goes to that of (X, Y) , *even though U_ϵ and V_ϵ are independent while X and Y are not independent.*

□

Note that (3) concerns the unconditional distribution of (U_ϵ, V_ϵ) . The latter is a random vector in $A_{(X,Y),\epsilon}$.

Intuitively, although U_ϵ and V_ϵ are conditionally independent, the vector (U_ϵ, V_ϵ) is close to (X, Y) in the unconditional distribution and thus has approximately the same joint distribution.

Proof:

First,

$$\lim_{\epsilon \rightarrow 0} U_\epsilon = X \quad (4)$$

and

$$\lim_{\epsilon \rightarrow 0} V_\epsilon = Y \quad (5)$$

Using the Bounded Convergence Theorem, we have

$$\lim_{\epsilon \rightarrow 0} P(U_\epsilon \leq a \text{ and } V_\epsilon \leq b) = \lim_{\epsilon \rightarrow 0} E[P(U_\epsilon \leq a \text{ and } V_\epsilon \leq b \mid X, Y)] \quad (6)$$

$$= \lim_{\epsilon \rightarrow 0} E[P(U_\epsilon \leq a \mid X, Y) \cdot P(V_\epsilon \leq b \mid X, Y)] \quad (7)$$

$$= E[1_{X \leq a} \cdot 1_{Y \leq b}] \quad (8)$$

$$= E[1_{X \leq a \text{ and } Y \leq b}] \quad (9)$$

$$= P(X \leq a \text{ and } Y \leq b) \quad (10)$$

$$= F(a, b) \quad (11)$$

7 Neighborhoods and Tuning Parameters

The neighborhoods are formed using both the Euclidean distance-based radius ϵ and the number of nearest neighbors k , which must be specified by the user. In short, ϵ provides control over the similarity of the data points within a neighborhood, while the nearest-neighbor parameter k guarantees that the neighborhood will have sufficiently many data points.

For many datasets, there are typical records as well as records with more unusual or extreme values. Typical records will have many neighbors even for small values of ϵ while unusual records may have zero neighbors unless ϵ is large. If ϵ alone were used to form neighborhoods, RWN would exclude from the released dataset these records with empty neighborhoods. This would protect their privacy, which is important considering that these unusual records may correspond to more identifiable individuals, but would result in a complete loss in their utility.

To avoid this, the user may choose to increase ϵ until these more extreme records have non-empty neighborhoods. However, as seen in Figure 1, an increase in ϵ can substantially increase the size of the neighborhood for typical records as well, causing the values in a typical record to be mixed with

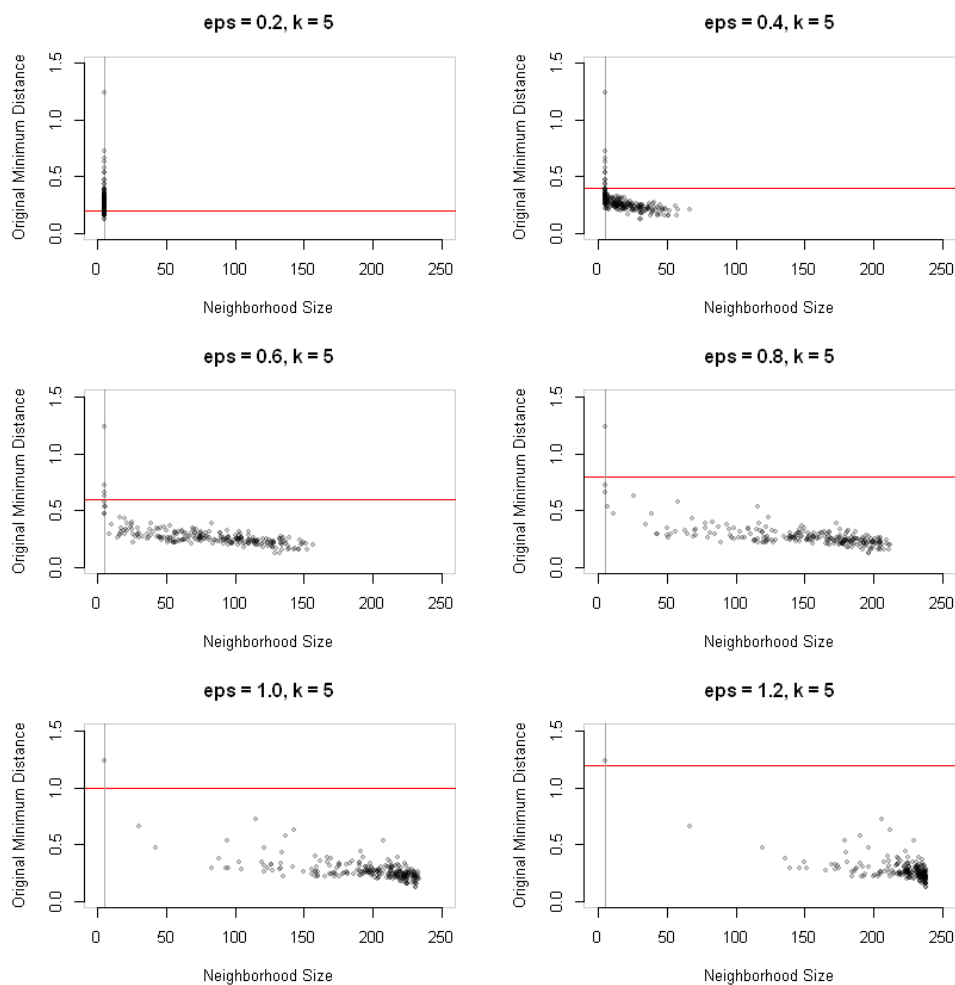


Figure 1: Body Fat Data, Plots of Minimum Distance vs. Neighborhood Size

very dissimilar records in the perturbation process, potentially leading to a decrease in utility of the perturbed data.

On the other hand, using k alone would impose both an upper and lower bound on the neighborhood size. For instance, for small k , the neighborhood size may be suitable for unusual records but unnecessarily small for typical records.

Thus, having two neighborhood size parameters ϵ and k gives the database administrator finer control over the perturbation of the data to balance utility with privacy for a specific dataset.

In Figure 1 we illustrate how RWN interacts with a specific dataset and choice of tuning parameters to form neighborhoods for perturbation. Using the **bodyfat** data from the R package **mfp** [2], we calculate the minimum distance for each record (i.e. the distance to the closest record) and plot it against the neighborhood size (i.e. the number of records within the neighborhood) for multiple choices of ϵ while holding constant $q = 1$ and $k = 5$. In the charts, the red horizontal line depicts the value of ϵ and the gray vertical line depicts the value of k .

8 Empirical Investigations

We will present the results of our empirical investigation shortly. But first, we describe our criteria.

Any SDC method is a balance of statistical utility and degree of privacy. Empirical investigation of the method must then define measures for these two criteria.

8.1 Statistical Utility

As noted, statistical analysis is at its core a matter of identifying relations between variables. The question to consider in the SDC context is whether the relationships that exist in the original data tend to remain intact in the released data. Here we follow [8], who note that a reasonable measure is to assess whether the released dataset “can obtain approximately the same substantive [relational] results while simultaneously protecting the privacy.” For instance, from the color correlation plots in Figure 2, we find that we can perturb data while retaining broad correlation structure, including to a large extent the strength of the correlations.

Another aspect of utility is validity of standard errors for statistical inference purposes. For SDC methods affecting only a small portion of the data, this is less of an issue. For RWN, one can prove, say, that the standard errors are asymptotically valid, as $q \rightarrow 0$. (See [8] for a DP solution, under certain assumptions. As usual, the problem of discrete/categorical variables remains a challenge.)

Since we emphasize MRAR, a utility measure is needed toward that end. We address this by investigating how correlations, regression coefficients, and principal components are affected by perturbation.

Most SDC methodology is aimed at estimating relationships, rather than using those relationships for prediction of new cases. Here, we investigate that latter aspect.

8.2 Privacy

A number of measures of privacy have been used in the SDC literature. The reader is referred to the references on this, but here our choices were guided primarily by two examples. (As already mentioned, RWN by its nature is resistant to k-anonymity problems.)

- (a) The **bodyfat** dataset has several outliers. If an intruder to the database knows, for instance, that a certain individual has the highest Body Mass Index of all the subjects in the data, then this individual is at risk of the intruder identifying sensitive variables (say if this dataset contained private health records).
- (b) Problems may arise involving “inliers,” points in dense portions of the data but which may still be at high risk of identification. A hypothetical example that has been used in the SDC literature is that of an employee database, in which an intruder knows that there is just one female electrical engineer. The intruder then queries the total salaries of all female electrical engineers, and thereby illicitly learns her salary. Here the record of concern may, in terms of distance metric, be close to many other data points yet still be at substantial risk of identification.

A number of measures could be used to gauge potential problems of these sorts. We address outlier issues such as in Example (a) by using Mahalanobis distance, and via Cook’s distance for linear regression analysis [25].

8.3 Experiments

The first dataset in our empirical experiments is the **bodyfat** dataset, consisting of measurements of 252 adult males. Two estimates of their body fat percentage are calculated using the Brozek and Siri equations. For our analysis, we use only the body fat percentage based on Siri, eliminate 11 observations that contain values that appear to be erroneous or biologically implausible such as body fat percentages less than four percent, and calculate Body Mass Index (BMI) for each individual.

For the experiments in this section, we use the cleaned **bodyfat** dataset as well as several perturbed versions of this dataset created using RWN with tuning parameters $k = 5$, $q = 1$, and varying values of ϵ . In each figure where a horizontal blue line is present, it denotes the result for the unperturbed data.

Correlation:

We calculated Pearson correlation coefficients pairwise for all variables in the **bodyfat** dataset. The color correlation plots in Figure 2 show that for small ϵ , the correlation coefficients in the perturbed and unperturbed data for a given variable are of the same sign and of very similar magnitude. This provides evidence of the MRAR property of RWN.

Regression:

We regressed the body fat percentage on BMI and the neck, chest, abdomen, and hip measurements. After estimating the model on each datasets, we compared the estimated coefficients, standard errors, and Cook's distances.

As seen in Figure 4, there is non-trivial variation in the estimated coefficients on the intercept, BMI, and neck variables over varying ϵ . We also note that there are numerous cases where the sign of the estimated coefficient on the perturbed data differs from the sign of the corresponding estimated coefficient in the model estimated on the unperturbed data. To some extent, these results are unsurprising. It is clear from the correlation plot that many of the body measurements have a high positive correlation with each other. Consequently, we expect some instability in the estimated coefficients due to multicollinearity.

As shown in Figure 5, with the exception of the intercept for larger ϵ values, the standard errors in the perturbed data appear to be relatively stable and similar to the standard errors calculated using the unperturbed data.

The cleaned bodyfat dataset contains one individual who is substantially larger than all the others. Consequently, this individual would be one of the most easily identifiable in the unperturbed data. In Figure 6, this individual corresponds to the largest Cook's distance in the unperturbed data. However, even with minimal perturbation, the maximum Cook's distance becomes much lower, suggesting that this individual is no longer as easily identifiable in the perturbed data.

Mahalanobis distance and minimum Euclidean distance:

Figures 7 and 8 show a similar result as the Cook's distance plots; privacy is provided to identifiable data points for even small values of ϵ .

Principal component analysis:

After scaling the data, we performed principal component analysis on both the unperturbed and perturbed datasets. For small values of ϵ , we find that the standard deviation of the first principal component and proportion of variance corresponding to the first principal component are slightly higher than in the unperturbed data. As ϵ increases the variation appears to be spread over more principal components.

8.4 Prediction-Oriented Assessment

Statistical applications tend to fall into one of two general categories, which we will refer to as Description and Prediction. The former has the goal of understanding some entity, process, effect

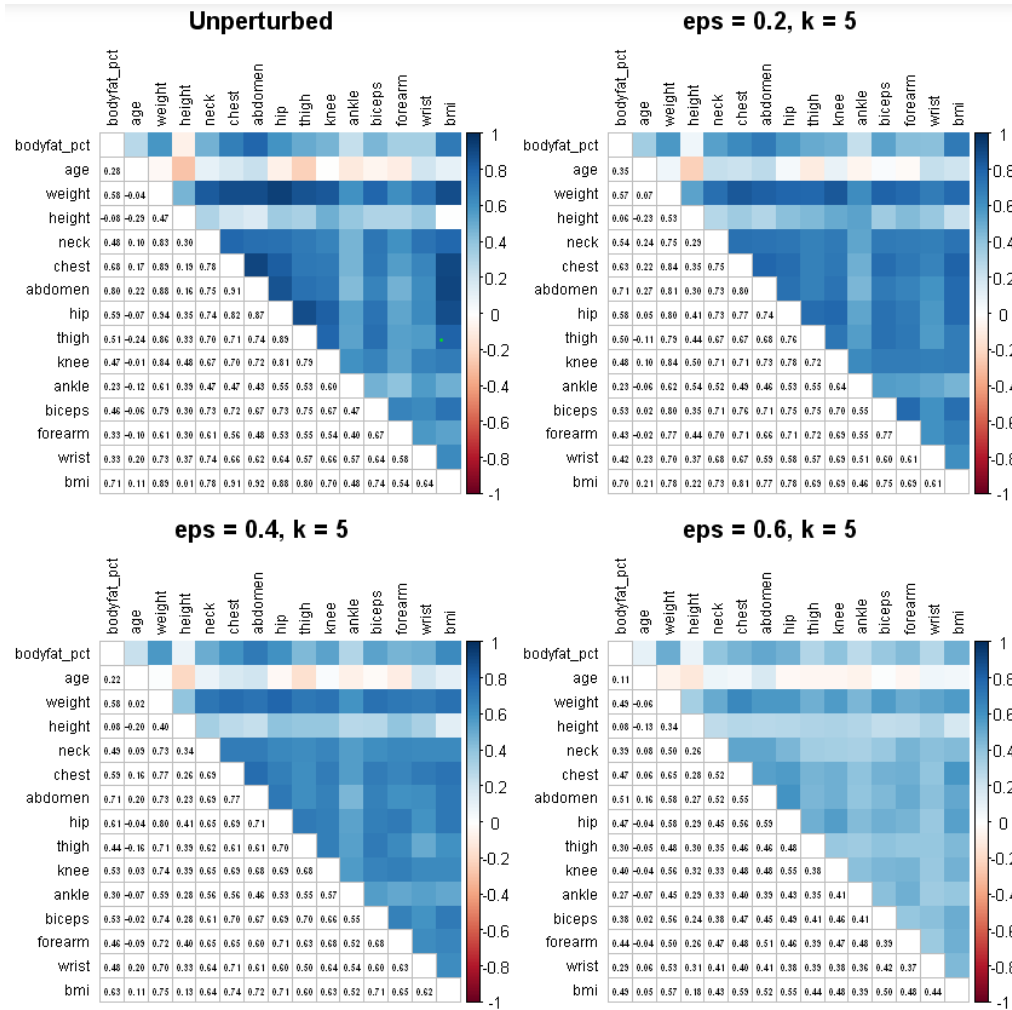


Figure 2: Body Fat Data, Correlation Plot

and so on, while the latter concerns predicting new data.

Much of the SDC literature has been aimed at the Description side of things, estimating means, totals, regression coefficients and the like. The Prediction side has rarely been the focus, and we now turn to that aspect in this section. We are interested particularly classification settings, the form of many modern applications.

Intuitively, SDC methods should do fairly well in classification settings. Consider 2-class problems, for instance, where we are predicting $Y = 0, 1$ from a vector of covariates X , say with a continuous distribution. Assuming equal misclassification costs, the set of t for which $P(Y = 1|X = t) = 0.5$ forms the decision boundary. What is the effect of RWN’s data perturbation?

Consider a nonparametric regression method, say random forests. If a data point in the training set is far from the boundary, perturbation will have little or no effect on future predictions; only points very close to the boundary would be likely to experience change in prediction. One may conjecture, then, that SDC methods will not compromise prediction ability much, at least for nonparametric

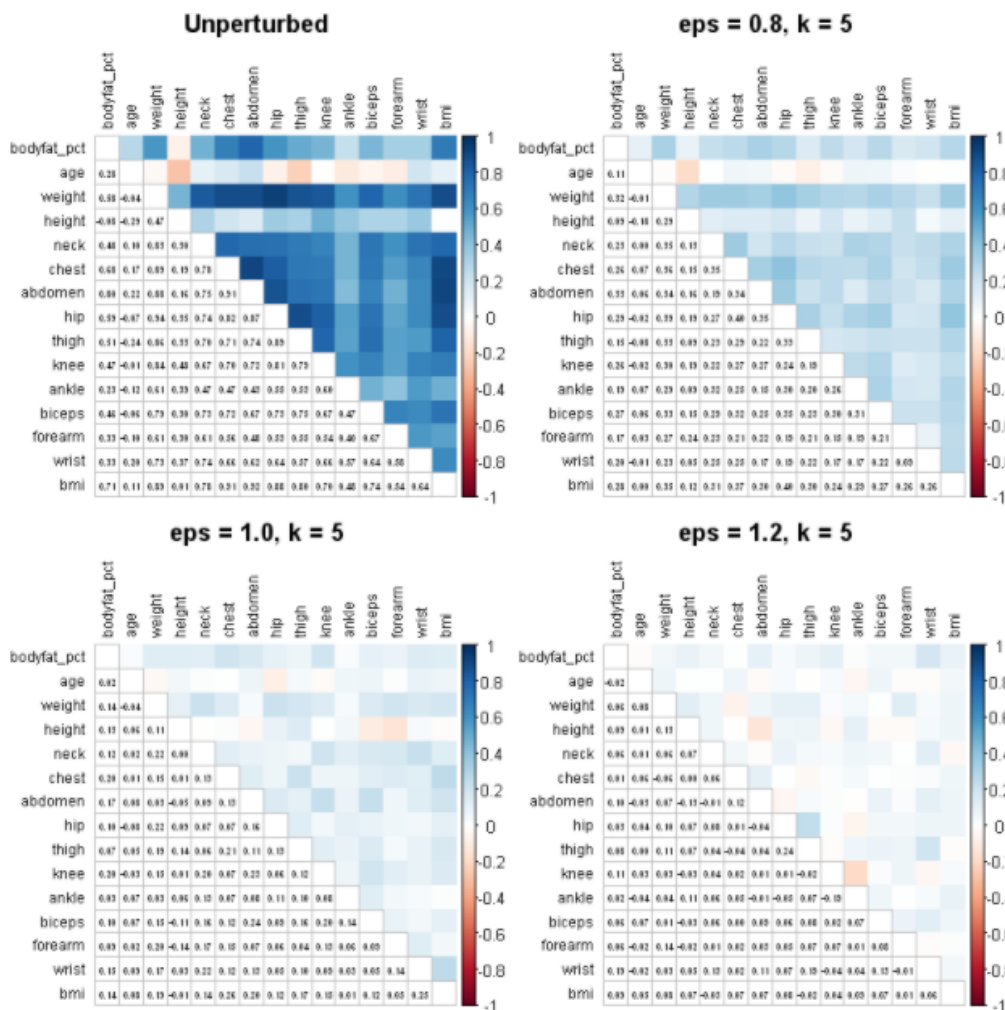


Figure 3: Body Fat Data, Correlation Plot (Large eps)

regression methods. This was confirmed in our experiments.

We first considered **pef**, a dataset included in the **R regtools** package. This is data on programmers and engineers, from the 2000 US census. We predict a variable **occ**, which codes one of six occupations, from variables such as age, income and education. Here we took $q = 0.5, \epsilon = 0$, and varied $k = 5, 10, 25, 50$. Misclassification rates, using random forests, are as follows:

method	no perturb.	k = 5	k = 10	k = 25	k = 50
RF	0.628	0.627	0.645	0.682	0.659

This dataset does not lend itself to strong predictability, with an error rate about about 63%. However, that rate increases only slightly under RWN. The above results were based on 25 replications, with a holdout set size of 1000.

Here is the same analysis on a second dataset, the well-known Pima diabetes study. It's quite different from the census data, in that it is much smaller (768 rows, vs. 20090 for **pef**), thus requiring

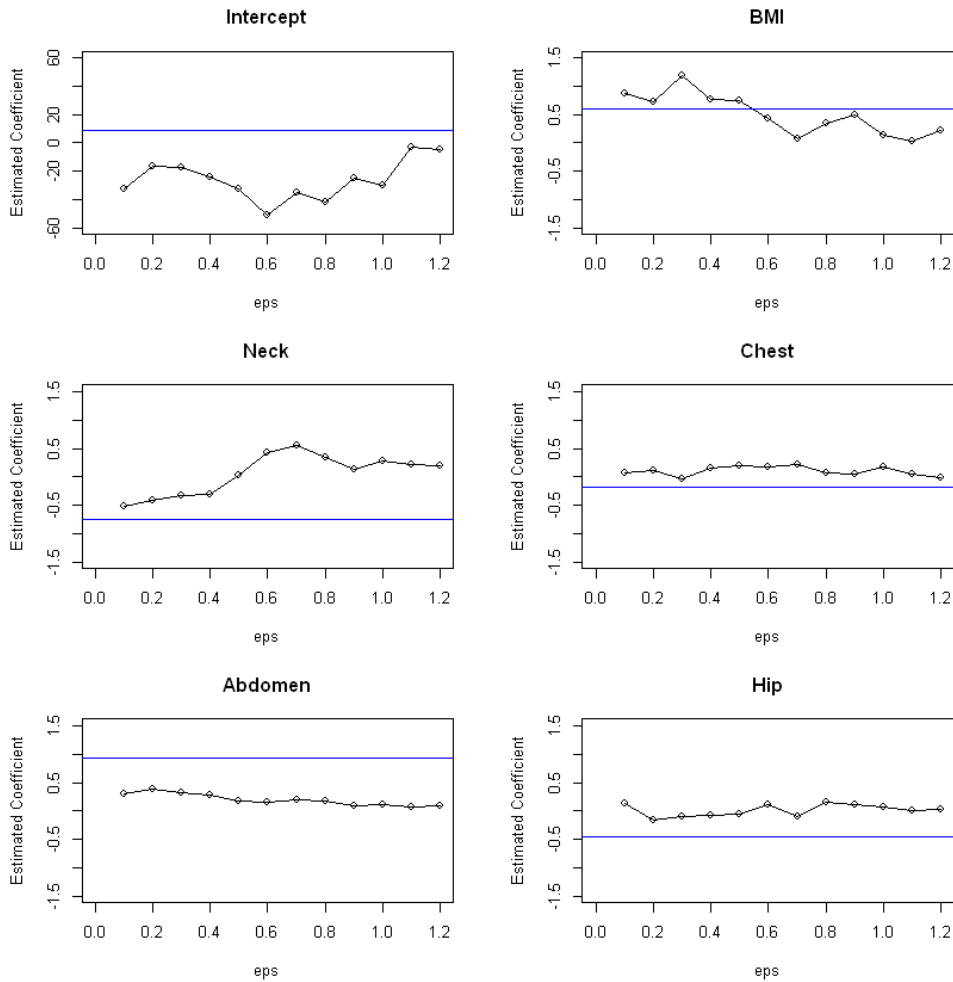


Figure 4: Body Fat Data, Regression Coefficients

more privacy protection. On the other hand, greater predictive accuracy is possible for this data.

method	no perturb.	k = 5	k = 10	k = 25	k = 50
RF	0.241	0.242	0.233	0.245	0.241

Again, performance appears not to decline due to the privacy action, and may even help, due to salutary smoothing effects.

9 Large Data Sets and Computational Complexity

9.1 Current Trends

Data sets are growing rapidly in both size and dimension. This is driven by many factors, including:

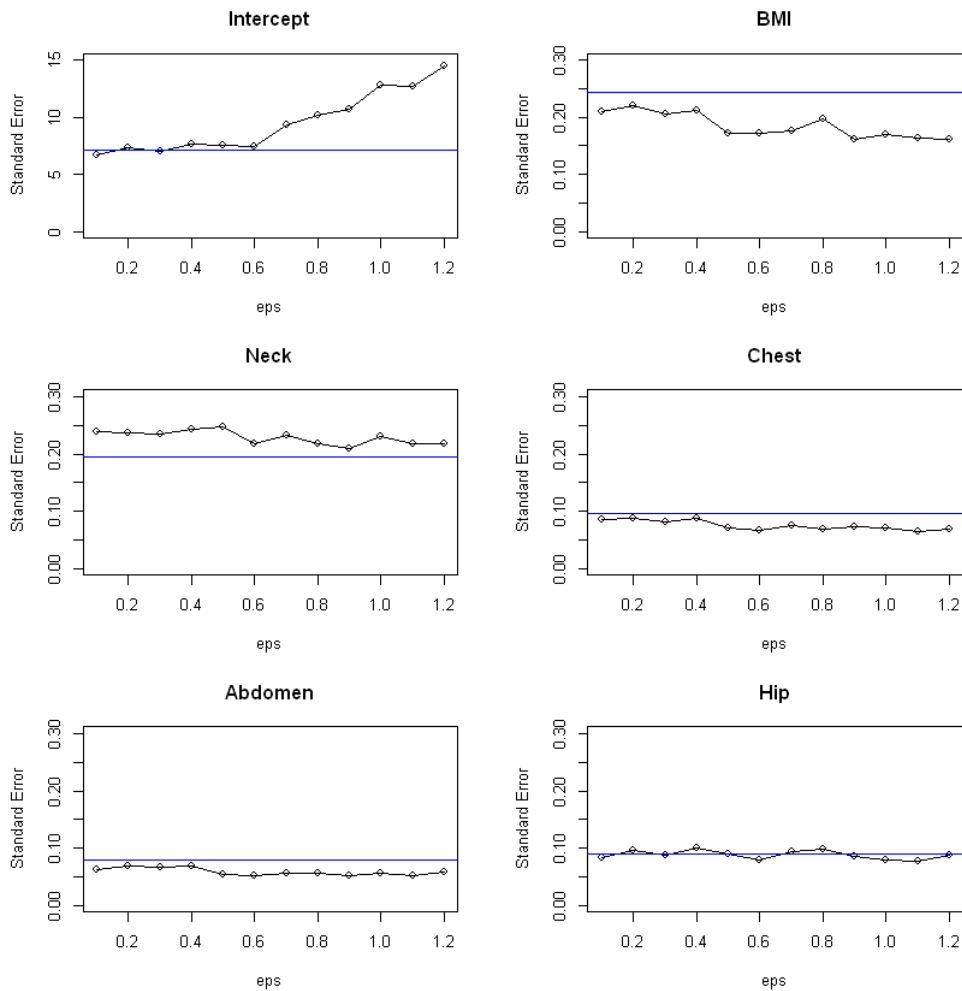


Figure 5: Body Fat Data, Regression Standard Errors

- A proliferation of data sources, including applications, back-office systems like ERP, smart devices, and sensors.
- The growth of the systems themselves, as businesses and other organizations achieve global scale.
- Increases in data collection and storage capacity through networks and mass storage.
- Larger computing devices and cloud scale computing.
- The transition from recording data about entities (such as people) to storing transactions like purchases to storing events like clicks. Rather than understanding people, the goal now is often is now on how people behave in terms of transactions, mobility and so on.
- The transition from structured to semi-structured to unstructured data. Structured data, like that found in traditional database tables, is strictly constrained in terms of dimension. Semi-

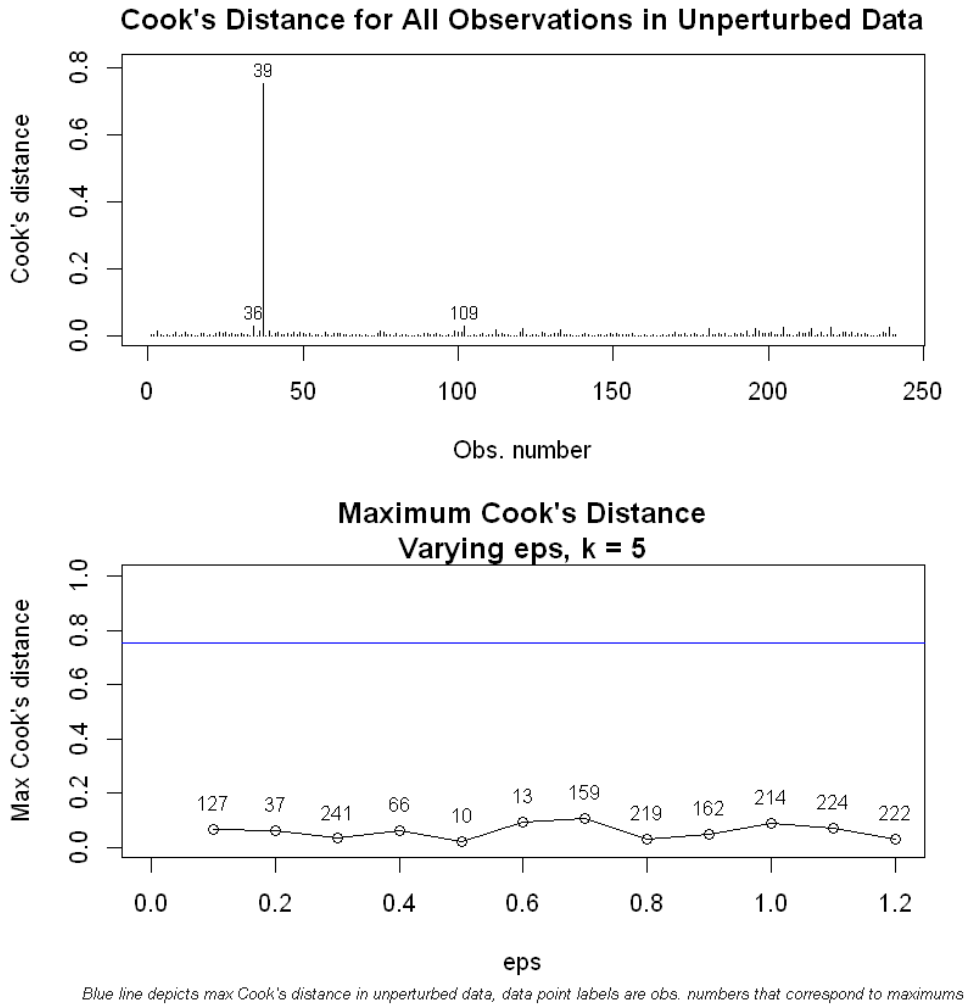


Figure 6: Body Fat Data, Cook's Distances

structured data like JSON objects, which is often captured from running applications, is unconstrained in dimension. Unstructured data, like documents, photos, or videos, is of almost unlimited dimension. Unstructured data might not seem like a candidate for the method described in this paper, but documents, photos, and videos are often reduced to a set of features described by categorical, binary, or numeric variables.

As a result, datasets can now easily be in the billions of rows with hundreds or thousands of variables or features. But with the growth of datasets comes a growth in risk. Larger datasets put more people at risk, and increased dimensionality increases damage per person. Also, the increase in dimensionality makes it easier to identify someone's record in a dataset. We need to be able to apply the method to these huge datasets, so the algorithm must be reasonably efficient from a computational complexity perspective. As we will see in the next section, the basic algorithm is computationally expensive, but with slight modification can handle large datasets.

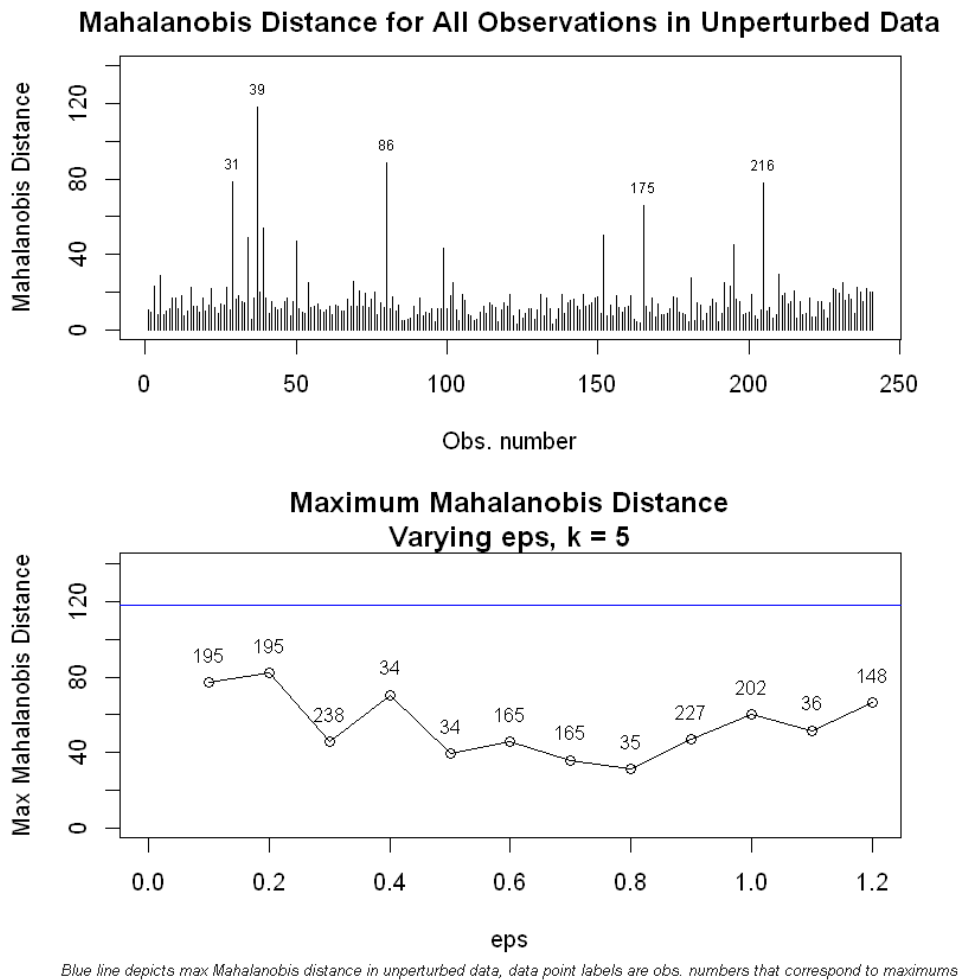


Figure 7: Body Fat Data, Mahalanobis Distances

9.2 Computational Complexity of RWN

To assess computational complexity of the method, going forward we will assume p is large but fixed, $q = 1$, and n is increasing. We will also assume that the distance metric used to define closeness is arbitrary. This is reasonable, since we need wide latitude to define closeness in different ways for very different datasets. Under these conditions, whether we are selecting nearby points based on epsilon neighborhood or k -nearest neighbor, the computational complexity of the basic method is $O(n^2)$, since we have to calculate the distance between all possible points. This is probably not tenable for large datasets, e.g., $n = 10^9$, for which the number of calculations would be of the order 10^{18} .




Figure 8: Body Fat Data, Boxplots of Minimum Euclidean Distances Before and After Perturbation

9.3 Alternative Methods

9.3.1 Method 1: Draw Neighbors from a Sample of Points

For the first method, we will just take a smaller sample of data points to use as neighbors. That is, we will take a sample S of size $m \ll n$. For each data point in the original data set, select the neighbors from the sample, then apply the method as before. For this method, the complexity is $O(mn)$, considerably better than the original method. For $n = 10^9$, $m = 10^4$, complexity is of order 10^{13} instead of 10^{18} , an improvement of five orders of magnitude. Assuming that all the data is in memory, the cost of generating a sample of size m is $O(m)$, so if we were to use a distinct random sample for each point, the complexity would still be $O(mn)$. For $n = 10^9$, $m = 10^4$, complexity is still order 10^{13} . The advantage is that we potentially get slightly richer data and better protection.

9.3.2 Method 2: Sample Randomly from the Distance Matrix

In Method 2, we are actually sampling twice as many points as we need to, since when we sample point x_j to be a possible neighbor of the point x_i , we can also use x_i as a potential neighbor of x_j . This is due to the symmetry of distance metrics and the distance matrix $D = [d_{ij}]$ where d_{ij} is the distance between points i and j . So instead of drawing a new sample for each data point, we could simply draw a sample of elements from the entire distance matrix. Then for each data point, we could look at the distances in the sample and find those that are sufficiently small or find the smallest k elements.

We will now describe the method more formally and derive the complexity. Let $W = (x_{ij})$ be the original data set. We want to sample pairs (x_i, x_j) , $i < j$ at random with a sample size of $n_s = nm/2$ where $m \ll n$ is the desired sample size per data point. There are $n(n-1)/2$ such pairs total, and we are sampling $n_s \ll n(n-1)/2$. To generate the sample, we can use a mixed congruential random number generator to generate numbers in $[1, n(n-1)/2]$.

Once we have n_s random integers in the range $[1, n(n-1)/2]$, we will map them to (i, j) pairs, $1 \leq i < j \leq n$ to obtain the set $S_d = \{(i, j)\}$. We then calculate the distances $d_{ij} = \langle x_i, x_j \rangle$ and store them in an undirected graph G_d with nodes representing the n data points and edges $\{e_{ij} : (i, j) \in S_d\}$. To map a random integer r to an (i, j) pair, we will use the following algorithm:

1. Calculate $v = \sqrt{(8r+1)}$ as an integer operation. That is, the calculation of v produces an integer plus a remainder that indicates whether v is the exact square root.
2. If v is odd and the exact square root,
 - $j \leftarrow (1+v)/2$
 - $i \leftarrow j-1$
 - and we're done.
3. Else
 - (a) If v is odd but not the exact root,
 - $j \leftarrow (3+v)/2$
 - (b) Else (v even)
 - $j \leftarrow (2+v)$
 - $i \leftarrow r-(j-1)(j-2)/2$

Since the complexities of generating the random numbers, mapping the random numbers to (i, j) pairs, calculating distances, and finding neighbors are all $O(nm)$, the complexity of the entire algorithm is $O(nm)$.

9.3.3 Method 3: Partitioning

With the advent of cloud computing, it has become much more feasible to harness the power of many computers or virtual machines, each with gigabytes of memory. Under this scenario, the problem (and data) are partitioned and spread across instances. There are several ways this approach could be applied to the problem at hand:

1. Simply partition the dataset into u equal partitions and then apply the original method. In this case, the size of each partition is n/u , so the complexity is $O(n^2/u^2) = O(n^2)/u^2$, a reduction by a factor of u^2 over the original method.

2. Partition the dataset into u equal partitions and then apply the Method 2 above. In this case, the size of each partition is n/u , so the complexity is $O(mn/u^2) = O(n^2)m/nu^2$, a reduction by a factor of nu^2/m over the original method.

Partitioning has two advantages: It both spreads the work across multiple instances, thereby increasing the computing power that can be brought to bear on the problem, and it also reduces the complexity of the problem being solved on each node. However, the resulting combined dataset will be somewhat different, since it is a combination of perturbed subsets.

10 Conclusion

We have presented a new SDC method for data release, distinguished by its ability to preserve multivariate relations as well as handle mixed continuous and discrete/categorical data, and provided a theoretical basis for the role of neighborhoods in achieving MRAR.

In addition to its theoretical appeal, RWN possesses practical features needed by DBAs. RWN's tuning parameters k , ϵ , and q provide DBAs multiple levers to adapt the method to a specific dataset and context. The DBA must select and tune an SDC method according to the special characteristics of the database at hand: Prevalence of outliers, distributional traits of the variables, presence of small cells, a need to make some parts of the data more secure than others, and so on. In deciding what kind of privacy must be provided, the DBA may need to take into account requirements of public policy specific to their domain or their organization's legal contracts. Hence, the importance of a particular measure of utility or privacy may differ by context. In light of this, we performed a variety of statistical analyses to examine privacy and utility from multiple angles. For example, a DBA who aspires to provide special privacy to individuals who are easily identifiable due to having unique or rare attributes may choose to adopt our Cook's distance approach. Our empirical results demonstrate RWN's ability to balance the twin goals of preserving data utility while providing privacy.

References

- [1] ABOWD, J., ASHMEAN, R., GARFINKEL, S., KIFER, D., LECLERC, P., MACHANAVAJIHALA, A., MORAN, B., AND SEXTON, W. Census topdown: Differentially private data, incremental schemas, and consistency with public knowledge, 2019.
- [2] AMBLER, G. *mfp: Multivariable Fractional Polynomials*, 2015. R package version 1.5.2.
- [3] BHRAMPUR, T., AND LANG, M. J. New system to protect census data may compromise accuracy, some experts say. *Washington Post* (6 2021).
- [4] BRAND, R. *Microdata Protection through Noise Addition*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2002, pp. 97–116.
- [5] BREIMAN, L. Statistical Modeling: The Two Cultures (with comments and a rejoinder by the author). *Statistical Science* 16, 3 (2001), 199 – 231.
- [6] DUNCAN, G., ELLIOT, M., AND SALAZAR, G. *Statistical Confidentiality: Principles and Practice*. Statistics for Social and Behavioral Sciences. Springer New York, 2011.
- [7] DWORK, C. Differential privacy. In *Proceedings of the 33rd International Conference on Automata, Languages and Programming - Volume Part II* (Berlin, Heidelberg, 2006), ICALP'06, Springer-Verlag, p. 1–12.
- [8] EVANS, G., KING, G., SCHWENZFEIER, M., AND THAKURTA, A. Statistically valid inferences from privacy protected data. <https://gking.harvard.edu/dp>, 2021.

-
- [9] FIENBERG, S. E., AND MCINTYRE, J. Data swapping: Variations on a theme by dalenius and reiss. In *Privacy in Statistical Databases* (Berlin, Heidelberg, 2004), J. Domingo-Ferrer and V. Torra, Eds., Springer Berlin Heidelberg, pp. 14–29.
- [10] FISCHETTI, M., AND SALAZAR-GONZÁLEZ, J.-J. Partial cell suppression: A new methodology for statistical disclosure control. *Statistics and Computing* 13, 1 (Feb. 2003), 13–21.
- [11] GONG, R., GROSHEN, E. L., AND VADHAN, S. Harnessing the Known Unknowns: Differential Privacy and the 2020 Census. *Harvard Data Science Review*, Special Issue 2 (jun 24 2022). <https://hdsr.mitpress.mit.edu/pub/fgyf5cne>.
- [12] HUNDEPOOL, A., DOMINGO-FERRER, J., FRANCONI, L., GIESSING, S., NORDHOLT, E., SPICER, K., AND DE WOLF, P. *Statistical Disclosure Control*. Wiley Series in Survey Methodology. Wiley, 2012.
- [13] KIM, J. J. Method for limiting disclosure in microdata based on random noise and transformation. In *Proceedings of the Section on Survey Research Methods (JSM)* (1986), pp. 303–308.
- [14] MATLOFF, N., AND TENDICK, P. The 'curse of dimensionality' in database security. In *Database Security, II: Status and Prospects. Result of the IFIP WG 11.3 Workshop on Database Security, Kingston, Ontario, Canada, October 5-7, 1988* (1988), C. E. Landwehr, Ed., North-Holland, pp. 225–232.
- [15] MATLOFF, N., AND TENDICK, P. A new method for avoiding data disclosure while automatically preserving multivariate relations, 2015.
- [16] MATLOFF, N. S. Another look at the use of noise addition for database security. In *1986 IEEE Symposium on Security and Privacy* (1986), pp. 173–173.
- [17] MOORE, R. Controlled data-swapping techniques for masking public use microdata sets. https://www.census.gov/srd/CDAR/rr96-04_Controlled_DataSwapping.pdf, 1996.
- [18] MURALIDHAR, K., AND SARATHY, R. Data shuffling: A new masking approach for numerical data. *Management Science* 52, 5 (2006), 658–670.
- [19] OBERSKI, D. L., AND KREUTER, F. Differential privacy and social science: An urgent puzzle. *Harvard Data Science Review* 2, 1 (1 2020).
- [20] RATHINDRA SARATHY, K. M. Evaluating laplace noise addition to satisfy differential privacy for numeric data. *Transactions on Data Privacy* 4 (4 2011).
- [21] TEMPL, M., KOWARIK, A., AND MEINDL, B. Statistical disclosure control for micro-data using the r package `sdcmicro`. *Journal of Statistical Software* 67, 4 (2015), 1–36.
- [22] TENDICK, P. Optimal noise addition for preserving confidentiality in multivariate data. *Journal of Statistical Planning and Inference* 27 (1991), 341–353.
- [23] TENDICK, P., AND MATLOFF, N. A modified random perturbation method for database security. *ACM Trans. Database Syst.* 19, 1 (Mar. 1994), 47–63.
- [24] WASSERMAN, L. A statistical view of differential privacy. www.cs.cmu.edu/afs/cs/usr/wing/www/class/15-895/LarryWasserman.pdf, 2011.
- [25] WEISBERG, S. *Applied Linear Regression*. Wiley Series in Probability and Statistics. Wiley, 2014.




Figure 9: Body Fat Data, PCA Variance Plots