Disclosure-Driven Suppression in Frequency Tables

Daniel P. Lupp and Øyvind Langsrud

Abstract. The abstract should briefly summarize the contents of the paper in 15–250 words.

1 Introduction

Cell suppression as currently implemented in popular tools consists of two parts: (1) identifying sensitive cells to be suppressed (primary suppression) and (2) preventing too precise recalculation of the primary cell values (secondary suppression). When the disclosed information is the cells' value or a unit's contribution (as is the case with magnitude tables), avoiding recalculation of cell values is a natural approach. However, in frequency tables, where a unit's membership in a cell should not be disclosed (as opposed to a unit's contribution to the cell value, which is 1 for frequency tables), disclosure can happen independent of the actual value. Consider the following table:

	Uninjured	Light	Serious	**Total**
Oslo	0	0	153	**153**
Trondheim	0	0	2	**2**
Total	**0**	**0**	**155**	**155**

With regards to disclosure of *membership* in a cell, both rows in the above table are equally disclosive: all units in Oslo and Trondheim, respectively, were seriously injured. However, if one primary suppresses only small cells greater than zero, the suppressed table will still disclose the level of injury in Oslo. To address this, it is well-discussed within the community that primary suppression of small values is not sufficient to protect against disclosure, and that the presence of zeros in fact leads to disclosure. Indeed, primary suppressing all zeroes as well as small counts will often provide sufficient protection against disclosure. However, we argue this can lead to over-protection. Consider the following table:

	Uninjured	Light	Serious	**Total**
Oslo	3	0	3	**6**
Trondheim	3	3	3	**9**
Total	**6**	**3**	**6**	**15**

If one suppresses all cells less than or equal to 3 (including zeroes), then no inner cells and not all marginals are allowed to be published. However, a publishing entity might consider this grossly over-protected: for instance, if it is deemed highly improbable that an attacker has knowledge of more than one unit (e.g.,

	Uninjured	Light	Serious	**Total**
Oslo	0	1	5	**6**
Trondheim	4	1	3	**8**
Bergen	0	0	2	**2**
Stavanger	1	4	4	**9**
Total	**5**	**6**	**14**	**28**

Table 1. Example frequency table, describing counts for various degrees of injury for car accidents in major Norwegian cities.

themselves), then the unsuppressed table might be considered safe to publish. In such cases, where the table consists of primarily small cells, one can argue that one should rather choose other protection methods than suppression—or not publish the table at all. However depending on users' needs and requirements, publishing with suppression might be the only feasible option.

We posit that the existing rule of thumb relying on suppressing small cell values (with or without zeros) can be poorly suited for protecting against disclosure: either over-protecting by suppressing too many cells, or under-protecting by letting certain disclosures slip through. In this work, we present an approach to cell suppression which is directly aimed at protecting against such categorical attribute disclosure in frequency tables. In particular, we claim that in this case it is not cell values that are sensitive and need protection, but rather certain differences between marginal cells and their contributors. This is formalized in Sections 2 and 3. In Section 4, we discuss how protection can be further relaxed in the presence of *non-disclosive unknown values*. We present a general algorithm for solving this version of the cell suppression problem, and describe in Section 5 how existing tools such as sdcTable and tauArgus can be adjusted to protect against these disclosures. In particular, most features of this method have been implemented in the `GaussSuppression` R package [1], which we briefly describe in Section 6. Finally, we discuss possible extensions and generalizations in Section 7.

2 *k*-Disclosure

In this section, we wish to formalize assumptions on an attacker's knowledge and use this to define the notion of k-disclosure. Intuitively, k-disclosure represents membership disclosure in frequency tables based on an attacker's knowledge of contributing units. We begin with some basic definitions to serve as a foundation for discussion, as well as discuss the types of disclosure this work is concerned with.

A *table* is a vector of values $T = [a_1, \ldots, a_n]$. This paper is concerned with frequency tables, so we assume that the values a_i are natural numbers greater than or equal to 0. In the following we refer to the indices $i = 1, \ldots, n$ as *cells*. A *marginal cell* in T is a cell that is equal to the sum of other cells. Let j_r for $r = 1, \ldots, m$ denote the marginal cells in the table T.

2.1 Types of Disclosure

In general, this paper is concerned with protecting against disclosure of categorical information that can be associated with individual statistical units. In particular, we do not consider purely existential disclosures (e.g., "there exists only one unit in Oslo that was lightly injured"), as these do not disclose information about specific units.[1] Rather, we are interested in universally quantified disclosures ("all units in Bergen were seriously injured"), as these lead to attribute disclosure (e.g., the level of injury) of any unit within the disclosure's scope. This distinction raises the question of negative disclosure: Existential disclosures can give rise to restricted universally quantified negative disclosures. Consider the previous example "there exists only one unit in Oslo that was lightly injured". In this circumstance, the one unit in Oslo who was lightly injured can infer that *all other units were not lightly injured*. Thus, the lightly injured unit can infer more specific information regarding the level of injury for all other statistical units in Oslo.

Whether or not this type of disclosure deserves protection is up to the publishing institution's policies. If it is deemed necessary to protect against any form of negative disclosure, one must at the very least suppress all zeros (and possibly all small cell values), as any zero would allow an attacker to exclude a category. Depending on the use case, however, protecting against all forms of negative disclosure might not be required. As such, we wish to provide more flexibility. We provide formalizations for specifying (1) an attacker's knowledge and (2) which (combinations of) cells need protection, as well as discuss an adaptation of the cell suppression problem tuned towards protecting against this disclosure. In Section 2.2, we define the notion of a k-coalition, describing the knowledge an attacking entity might have. Furthermore, we introduce the notion of k-disclosure, which corresponds to attribute disclosure a k-coalition can make. These sections are primarily focused on preventing disclosure of membership in single cells. In Section 3, we discuss how combinations of cells can be protected, thus providing a means for protecting *meaningful combinations* of cells as well as only specified negative disclosures.

2.2 Disclosure by k-Coalitions

We are concerned with the scenario where an attacker has knowledge of a certain number of statistical units contributing to a frequency table. In this case, the attacker might be able to use this information to disclose information about other units. The following definition formalizes this intuition.

Definition 1. *Let $T = [a_1, \ldots, a_n]$ be a frequency table with values a_i natural numbers greater than or equal to zero. A* k-coalition *is a set of k units contributing to T.*

[1] Note that we do not consider it problematic if a unit can identify themselves in a frequency table.

	Uninj.	Light	Serious	**Total**
Oslo	0	1	5	**6**
Trondheim	4	1	3	**8**
Bergen	0	0	2	**2**
Stavanger	*1*	4	4	**9**
Total	**5**	**6**	**14**	**25**

	Uninj.	Light	Serious	**Total**
Oslo	0	1	5	**6**
Trondheim	4	1	3	**8**
Bergen	0	0	2	**2**
Stavanger	*0*	4	4	**9**
Total	**4**	**6**	**14**	**24**

Table 2. Example of a frequency table (left) and a 1-coalition's difference table (right), where the attacker has knowledge of the uninjured unit in Stavanger. The differing cell is blue and italicized.

Thus, a k-coalition intuitively represents an attacker with full knowledge of the k units that are part of the coalition. Such a k-coalition might attack a frequency table by incorporating its knowledge into the table. They can do this by constructing a *difference table* from the published table: given $T = [a_1, \ldots, a_n]$, they construct a new table $T' = [a'_1, \ldots, a'_n]$ by removing its member units from T (see Table 2 for an example).

Definition 2. *Let $T = \{a_1, \ldots, a_n\}$ be a frequency table, where the values a_i are natural numbers greater than or equal to 0. A cell $i_s \in \{1, \ldots n\}$ is said to be* k-disclosive w.r.t. *j_r if it satisfies the following condition:*

$$i_s \text{ contributes to } j_r \text{ and } a_{j_r} - a_{i_s} \leq k.$$

We omit specifying the marginal j_r and refer to a cell as k-disclosive *where the marginal is not important in the context. If k is clear from context, we simply refer to k-disclosive cells as* disclosive.

With this definition, the following proposition provides the foundation for protection against k-disclosure.

Proposition 1. *Let i_s be a cell contributing to a marginal j_r. Then i_s is k-disclosive w.r.t. j_r if and only if there exists a k-coalition A such that all statistical units in j_r are either in A or in i_s.*

Thus, in the cases where an attacker can know a unit's membership in a marginal, the problem of protecting against membership disclosure reduces to protecting against k-disclosure. This in turn reduces to ensuring that an attacker cannot deduce that the difference between marginals and their contributing cells are less than k. Therefore, table cells themselves (independent of their value) are not sensitive. Instead, the differences between marginals and their contributing cells are sensitive when they are too small. Thus, to protect against k-disclosure, one must prevent attackers from knowing that such differences are less than k.

3 Meaningful Combinations

In addition to protecting against membership disclosure in table cells, one might wish to additionally protect certain combinations of cells, which we refer to as

meaningful combinations.[2] A meaningful combination represents a combination of published cells for which unit membership disclosure should be protected. In our running example in Table 1, let us assume one wishes to protect disclosure of whether a unit was injured or not. Since no "injured" marginals are published, this disclosure does not correspond to membership disclosure of a table cell. Thus, we introduce meaningful combination cells to be protected as the combination of "Light" and "Serious". The first row in Table 1, containing accidents in Oslo, illustrates a 0-disclosive meaningful combination: an attacker can deduce that any statistical unit in Oslo was injured. Likewise, the fourth row containing accidents in Stavanger illustrates a 1-disclosive meaningful combination: an attacker with knowledge of the uninjured unit in Stavanger can disclose that all other units in Stavanger were injured.

Thus, in addition to the table cells in T, we define new indices $\mathrm{c}_i, i = 1, \ldots, q$ representing *meaningful-combination cells*.

Definition 3. *Let $T = [a_1, \ldots, a_n]$ be a frequency table with $a_i \geq 0$. We define new indices c_i, $i = 1, \ldots, q$, which we call* meaningful combination cells*. Furhermore, we define* mcontr *to be a function associating meaningful-combination cells with their contributing table cells, i.e.,* $\mathrm{mcontr}(\mathrm{c}_i) \subseteq \{1, \ldots, n\}$*. Then the value a_{c_i} of a meaningful-combination cell c_i is defined as follows:*

$$a_{\mathrm{c}_i} := \sum_{s \in \mathrm{mcontr}(\mathrm{c}_i)} a_s.$$

Example 1. In the table in Table 1, if we wish to define the meaningful combination "injured" as the combination of "light" and "serious" then we introduce new meaningful-combination cells (Oslo, Injured), (Trondheim, Injured), (Bergen, Injured) and (Stavanger, Injured), with

$$\mathrm{mcontr}(X, \mathrm{Injured}) = \{(X, \mathrm{Light}), (X, \mathrm{Serious})\}$$

for $X \in \{\mathrm{Oslo}, \mathrm{Trondheim}, \mathrm{Bergen}, \mathrm{Stavanger}\}$.

Intuitively, one can view a meaningful combination as a marginal cell that is not to be published. Disclosure occurring in conjunction with meaningful combinations can naturally be divided into two categories: disclosure *of* a meaningful combination and disclosure *within* a meaningful combination.

Membership disclosure of a meaningful combination In this form of disclosure, we are concerned with the disclosure of a unit's membership in a meaningful combination cell c_i, or equivalently, disclosure that a unit must be a member in one of the cells in $\mathrm{mcontr}(\mathrm{c}_i)$.

Let j_r be a marginal cell containing a meaningful combination cell c_i. Using Definition 3, the statement in Proposition 1 hols for meaningful combinations:

[2] The naming of this term is borrowed from our colleagues in Statistics Netherlands, whom we thank for their much more suitable term than our original *shadow cell*.

An attacking k-coalition A can disclose that all units not in A that contribute to j_r must contribute to c_i if and only if c_i is k-disclosive w.r.t. j_r. Thus, as was the case with table cells, small differences between marginals and meaningful combination cells can lead to k-disclosure.

Example 2. To illustrate disclosure of meaningful combinations as well as how to protect against it, consider the following row in an extended version of our example table:

	U1	U2	Light	Serious	Critical	**Total**
Kristiansand	1	1	1	1	1	**5**

For the sake of discussion, U1 and U2 are two distinct characterizations of "uninjured". Assume we wish to protect the meaningful combination "injured", which is the sum of "light", "serious" and "critical", against disclosure by 2-coalitions, i.e., an attacker has full knowledge of up to two statistical units. In particular, assume the attacker has knowledge of the two units in U1 and U2. Consider the following suppression pattern:

	U1	U2	Light	Serious	Critical	**Total**
Kristiansand	1	×	×	1	1	**5**

Given this suppressed table and their knowledge, the attacking 2-coalition's estimate for the difference Total − Injured = U1 + U2 is 2 or 3. Since there is a possibility that the difference is 3, the attacker cannot deduce with certainty that all other units in Kristiansand are injured.

Membership disclosure within a meaningful combination So far, we have focused on using published marginal cells to disclose the membership of a statistical unit in a table cell or meaningful combination. There are, however, situations where a published marginal cell does not play a vital role in the disclosure.

Consider the second row in Table 1 concerning accidents in Trondheim. In terms of both "classical" k-disclosiveness as well as disclosure of the meaningful combination "Injured", this row does not exhibit k-disclosiveness for $k < 4$. However, it exhibits disclosure *within* the meaningful combination for $k = 1$. Consider the 1-coalition consisting of the lightly injured person in Trondheim. Assume further that they know (or can deduce through other means) that a specific statistical unit X in Trondheim that is represented in the table is injured. They can then use this knowledge and the table in Table 1 to disclose that X must be seriously injured.

Once again, we can adapt the statement of Proposition 1 and apply the same principle: Assume we have a meaningful combination cell c_i. Then there exists a k-coalition A that can disclose than another unit contributing to c_i is in cell i_s if and only if i_s is k-disclosive w.r.t. c_i. Note that here c_i plays the role of the marginal in Proposition 1. Thus, in order to protect against k-disclosure of i_s within a meaningful combination cell c_i, attackers must not be able to estimate the difference $a_{c_i} - a_{i_s}$ to be less than or equal to k.

4 Disclosive and Non-Disclosive Unknowns

5 *k*-Disclosure Suppression

Here, we describe the algorithm for k-disclosure suppression (work in progress):

Algorithm 1 Find disclosive difference cells

Require: Table $T = \{a_1, \ldots, a_n\}$, meaningful combinations cells $M = \{\mathrm{c}_1, \ldots, \mathrm{c}_n\}$, $k \geq 0$

1: **function** FIND-DIFF-CELLS(T, M, k)
2: $dcells = \emptyset$
3: **for** each cell $i \in \{1, \ldots, n, \} \cup M$ **do**
4: **for** each marginal j_r to which i contributes **do**
5: **if** $a_{j_r} - a_i \leq k$ **then**
6: create new difference cell $\langle j_r, i \rangle$ with $a_{\langle j_r, i \rangle} = a_{j_r} - a_i$
7: add $\langle j_r, i \rangle$ to $dcells$
8: **if** $i \in M$ **then**
9: **for** each cell $l \in \mathrm{mcontr}(i)$ **do**
10: **if** $a_i - a_l \leq k$ **then**
11: create new difference cell $\langle i, l \rangle$ with $a_{\langle i,l \rangle} = a_i - a_l$
12: add $\langle i, l \rangle$ to $dcells$
return $dcells$

Lines 4–7 produce the difference cells for k-disclosive cells i w.r.t. a marginal j_r (for i a table cell or meaningful combination cell.) Lines 8–12 add those difference cells corresponding to k-disclosive cells w.r.t. meaningful combination cells, i.e., those cells that can lead to disclosure within a meaningful combination.

Using Algorithm 1, we can describe a general cell suppression method for protecting against k disclosure, as seen in Algorithm 2.

Algorithm 2 k-disclosure suppression

Require: Table $T = \{a_1, \ldots, a_n\}$, meaningful combinations cells $M = \{\mathrm{c}_1, \ldots, \mathrm{c}_n\}$, $k \geq 0$

1: $dcells =$ FIND-DIFF-CELLS(T, M, k)
2: Run secondary suppression algorithm where
- $dcells$ as primary suppressed,
- Only table cells $\{1, \ldots, n\}$ are candidates for secondary suppression.

In practice, the difference cells will correspond to sums of table cells: The difference between a (marginal or meaningful combination) cell j and a table cell i is the sum of the other table cells contributing to j, and the difference between a marginal j_r and a meaningful combination cell c_i is the sum of table cells contributing to j_r not in $\mathrm{mcontr}(\mathrm{c}_i)$.

In theory, any cell suppression tool that supports specification of such difference cells can be used to protect against k-disclosure.[3] As such, support for this functionality should be fairly straightforward to implement in popular solutions such as `sdcTable` [2] and tauArgus [3]. An implementation of many of the features discussed above is available in the `GaussSuppression` R package [1], which we briefly describe in Section 6.

6 Implementation

7 Generalizations

Acknowledgments We would like to thank Peter Paul de Wolf from Statistics Netherlands for his insightful questions at and after the UNECE Expert Meeting on SDC 2021, in particular those regarding meaningful combinations, which helped make this work more mature.

References

1. Langsrud, Ø., Lupp, D.P.: GaussSuppression: Tabular Data Suppression using Gaussian Elimination (2022), https://CRAN.R-project.org/package=GaussSuppression, r package version 0.4.0
2. Meindl, B.: sdcTable: Methods for Statistical Disclosure Control in Tabular Data (2021), https://CRAN.R-project.org/package=sdcTable, r package version 0.32.2
3. de Wolf, P.P., Hundepool, A., Giessing, S., Salazar, J.J., Castro, J.: tau-ARGUS user's manual, version 4.1. Tech. rep., Statistics Netherlands (2014)

[3] Indeed, one could "hack" this support by simply adding all difference cells to the table and primary suppressing those difference cells corresponding to k-disclosive cells. This naive approach will however lead to very large instances to be solved, greatly impacting run time.