Assessing the disclosure risk of CTA-like methods

Jordi Castro

jordi.castro@upc.edu

Dept. of Statistics and Operations Research
Universitat Politècnica de Catalunya
Barcelona

This work has been supported by grants MTM2012-31440 of the Spanish research program and SGR-2009-1122 of the Government of Catalonia

Contents

Contents

Motivation I

- Tabular protection method: map $F : \mathit{Tables} \to \mathit{Tables}$ such that $F(T) = T'$ and T' is "safe".

Motivation I

- Tabular protection method: map $F : \mathit{Tables} \to \mathit{Tables}$ such that $F(T) = T'$ and T' is "safe".
- "safe": tight estimates of sensitive cells can not be recomputed.

Motivation I

- Tabular protection method: map $F : \mathit{Tables} \to \mathit{Tables}$ such that $F(T) = T'$ and T' is "safe".
- "safe": tight estimates of sensitive cells can not be recomputed.
- Therefore, the inverse map $T = F^{-1}(T')$ should not be available or difficult to compute by any attacker, otherwise disclosure risk is high.

Motivation I

- Tabular protection method: map $F : \mathit{Tables} \to \mathit{Tables}$ such that $F(T) = T'$ and T' is "safe".
- "safe": tight estimates of sensitive cells can not be recomputed.
- Therefore, the inverse map $T = F^{-1}(T')$ should not be available or difficult to compute by any attacker, otherwise disclosure risk is high.
- CTA: post-tabular approach which looks for safe table T' closest to original T. It solves a minimum-distance optimization problem.

Motivation I

- Tabular protection method: map $F : \mathit{Tables} \to \mathit{Tables}$ such that $F(T) = T'$ and T' is "safe".
- "safe": tight estimates of sensitive cells can not be recomputed.
- Therefore, the inverse map $T = F^{-1}(T')$ should not be available or difficult to compute by any attacker, otherwise disclosure risk is high.
- CTA: post-tabular approach which looks for safe table T' closest to original T. It solves a minimum-distance optimization problem.
- CTA-like methods have low disclosure risk if no attacker can obtain a good estimate $\hat{T} = \hat{F}^{-1}(T')$, $\hat{F}^{-1}$ being an estimate of F^{-1}.

Motivation II

This talk:

- considers 4 different attacker scenarios;

Motivation II

This talk:

- considers 4 different attacker scenarios;
- provides exhaustive empirical evaluation of the disclosure risk of CTA;

Motivation II

This talk:

- considers 4 different attacker scenarios;
- provides exhaustive empirical evaluation of the disclosure risk of CTA;
- reports results for the solution of more than 2500 optimization attacker problems;

Motivation II

This talk:

- considers 4 different attacker scenarios;
- provides exhaustive empirical evaluation of the disclosure risk of CTA;
- reports results for the solution of more than 2500 optimization attacker problems;
- summarizes results in the paper:
  Castro, J. (2012). On assessing the disclosure risk of controlled adjustment methods for statistical tabular data, *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 20, 921–941.

Contents

CTA: Example

ORIGINAL TABLE. Protection levels: $x_{23} \geq 45$ or $x_{23} \leq 35$

	Z_1	Z_2	Z_3	TOTAL
E_1	20	24	28	72
E_2	38	38	40	116
E_3	40	39	42	121
TOTAL	98	101	110	309

CTA: Example

ORIGINAL TABLE. Protection levels: $x_{23} \geq 45$ or $x_{23} \leq 35$

	Z_1	Z_2	Z_3	TOTAL
E_1	20	24	28	72
E_2	38	38	40	116
E_3	40	39	42	121
TOTAL	98	101	110	309

PROTECTED TABLE: either ... or ...

	Z_1	Z_2	Z_3	
E_1	25	24	23	72
E_2	33	38	45	116
E_3	40	39	42	121
TOTAL	98	101	110	309

Z_1	Z_2	Z_3	
15	24	33	72
43	38	35	116
40	39	42	121
98	101	110	309

General formulation of CTA

- Set of cells $a_i, i = 1, \ldots, n$.
- Set $\mathcal{S} = \{i_1, i_2, \ldots, i_s\} \subseteq \{1, \ldots, n\}$ of indices of sensitive cells.
- Linear relations $Aa = b$.
- Lower and upper protection level for each sensitive cell $i \in \mathcal{S}$: lpl_i and upl_i.
- Lower and upper bound for each cell: l_{a_i} and u_{a_i}.
- Cell weights w_i for cost of adjustment of each cell.

General formulation of CTA

- Set of cells $a_i, i = 1, \ldots, n$.
- Set $\mathcal{S} = \{i_1, i_2, \ldots, i_s\} \subseteq \{1, \ldots, n\}$ of indices of sensitive cells.
- Linear relations $Aa = b$.
- Lower and upper protection level for each sensitive cell $i \in \mathcal{S}$: lpl_i and upl_i.
- Lower and upper bound for each cell: l_{a_i} and u_{a_i}.
- Cell weights w_i for cost of adjustment of each cell.
- CTA formulation:

$$\begin{array}{ll} \min\limits_{x} & ||x - a||_{\ell(w)} \\ \text{s. to} & Ax = b \\ & l_{a_i} \leq x_i \leq u_{a_i} \quad i \in \mathcal{N} \\ & (x_i \leq a_i - lpl_i) \text{ or } (x_i \geq a_i + upl_i) \quad i \in \mathcal{S}. \end{array}$$

CTA formulation fixing protection senses

- Defining $z = x - a$, and fixing the protection senses we have the convex problem:

$$\begin{array}{ll} \min\limits_{z} & ||z||_{\ell(w)} \\ \text{s. to} & Az = 0 \\ & l(a) \leq z \leq u(a), \end{array}$$

- For ℓ_1 norm we have the linear problem

$$\begin{array}{ll} \min\limits_{z^+, z^-} & \sum_{i=1}^{n} w_i(a_i)(z_i^+ + z_i^-) \\ \text{s. to} & A(z^+ - z^-) = 0 \\ & l^+(a) \leq z^+ \leq u^+(a) \\ & l^-(a) \leq z^- \leq u^-(a), \end{array}$$

- For ℓ_2 norm we have the convex quadratic problem

$$\begin{array}{ll} \min\limits_{z} & \sum_{i=1}^{n} w_i(a_i) z_i^2 \\ \text{s. to} & Az = 0 \\ & l(a) \leq z \leq u(a). \end{array}$$

Contents

Attacker information and attacker problem

Attacker wants to obtain a *good* estimate $\hat{z}$ of z from released T'.

- Attacker complete information:
  - Released values x.
  - Structure of the table, i.e., constraints matrix A.
- Attacker partial information:
  - Distance used.
  - Cell weights $w(a)$: depend on function of a
  - Bounds $l^+(a), l^-(a), u^+(a), u^-(a), u(a), l(a)$: depend on a, $\mathcal{S}$ and protection senses.

Optimization problem to be solved by attacker

$$\begin{array}{ll} \min\limits_{\hat{z}} & ||\hat{z}||_{\ell(x)} \\ \text{s. to} & A\hat{z} = 0 \\ & \hat{l}(x) \leq \hat{z} \leq \hat{u}(x) \end{array}$$

where $\hat{l}(x)$ and $\hat{u}(x)$ are estimates of the bounds.

Attacker scenarios

B. The attacker has incomplete information about both the bounds and objective function, but he/she knows the subset $\mathcal{S}$ of sensitive cells, and the original cell bounds l_{a_i} and u_{a_i}, $i \in \mathcal{N}$ (which are quite strong assumptions). We have three subscenarios:

   B1. The attacker neither knows the protection levels $upl_i, lpl_i, i \in \mathcal{S}$, nor the protection sense.

   B2. The attacker knows the protection sense, but not the protection levels upl_i, lpl_i, $i \in \mathcal{S}$.

   B3. The attacker knows both the protection sense and protection levels upl_i, lpl_i, $i \in \mathcal{S}$. The only unknown terms to reproduce the real bounds are then $a_i - l_{a_i}$ and $u_{a_i} - a_i$, $i \in \mathcal{N}$.

C. The attacker has complete information about the bounds, i.e, he/she knows all the parameters but the objective function w.

Contents

Instances and experiments

- 33 standard instances from the literature.
- Procedure:
  - Tables first protected with ℓ_1-CTA and ℓ_2-CTA.
  - Next, attacker problems solved for the four scenarios, 10 random replications for different values of $\tilde{x}$.
- This amounts to: 33 instances $\times$ 2 distances $\times$ 4 scenarios $\times$ 10 replications = 2640 optimization problems.
- Following plots show the distribution of percentage differences between estimates and true values of sensitive cells $|\hat{a}_i - a_i|/a_i \cdot 100$, $i \in \mathcal{S}$ for all the attacker problems, and the different scenarios.

Results for scenario B1

ℓ_1

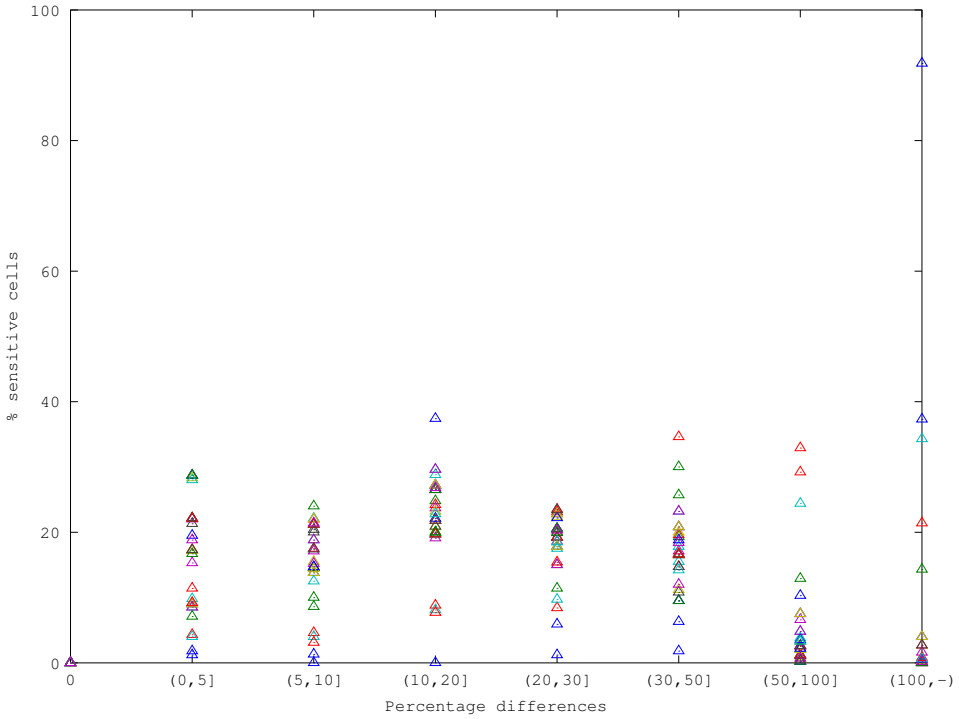

ℓ_2

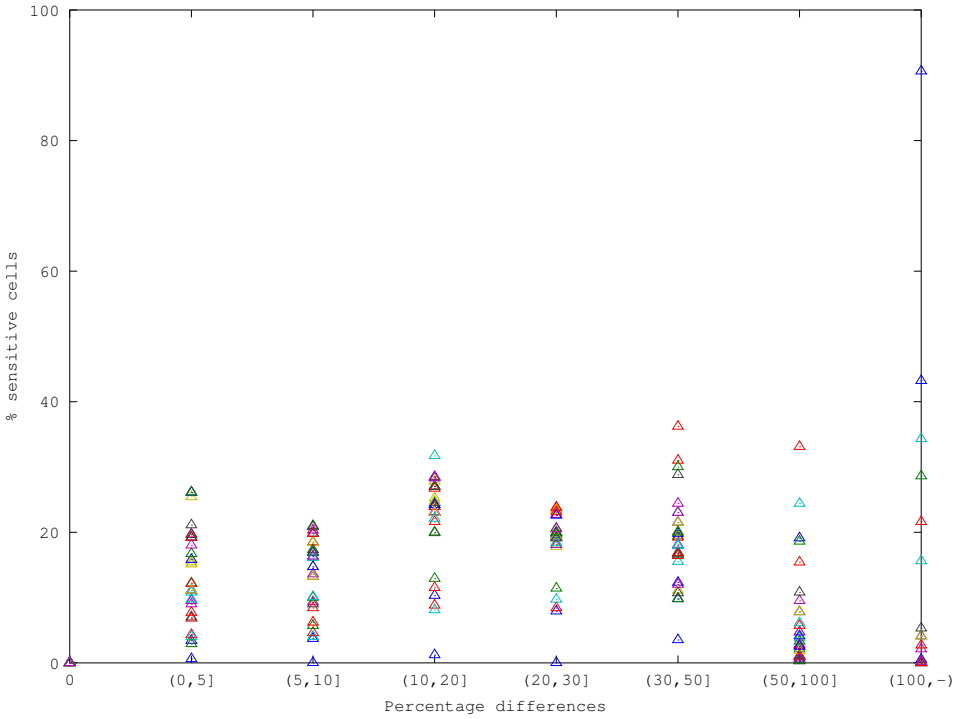

CTA is safe for scenario B1

ℓ_1 and ℓ_2 are similar

Results for scenario B2

ℓ_1

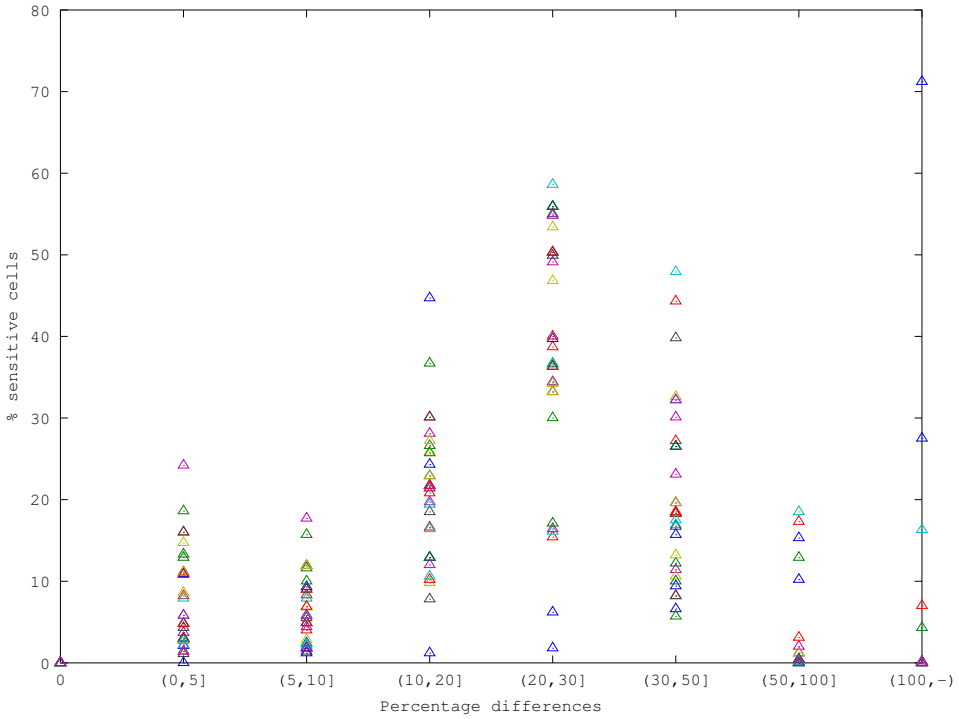

ℓ_2

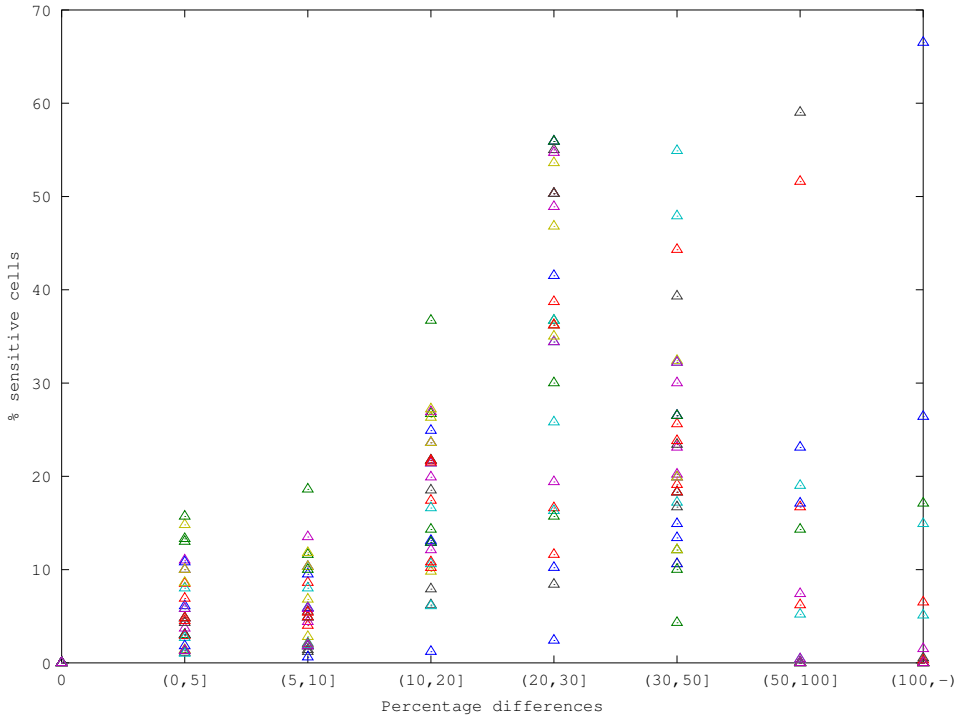

CTA is safe for scenario B2

ℓ_2 seems to be safer than ℓ_1

Results for scenario B3

ℓ_1

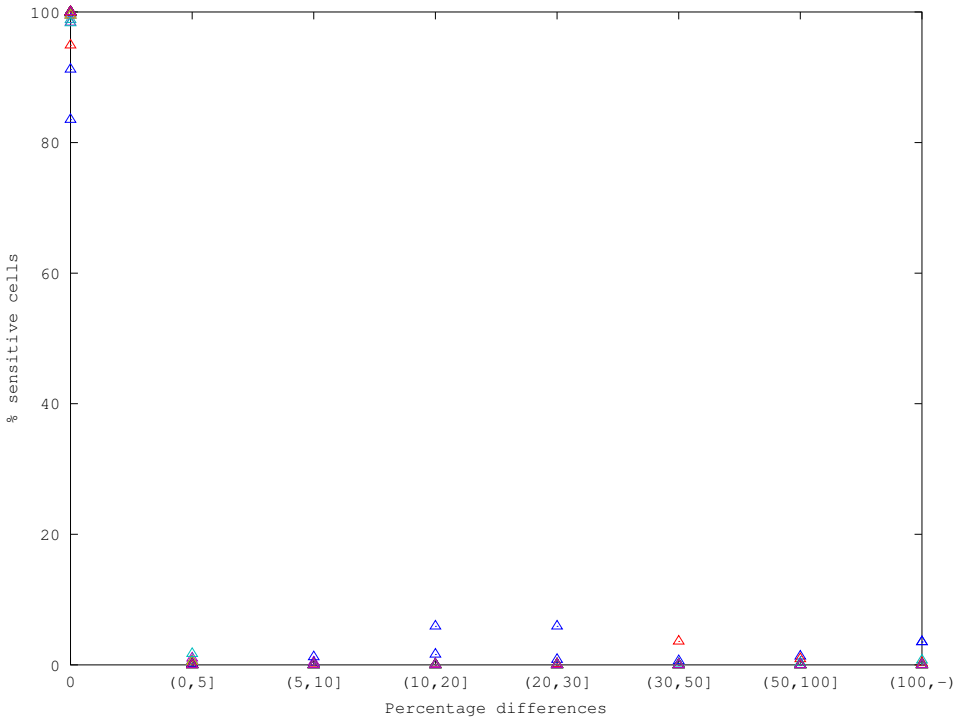

ℓ_2

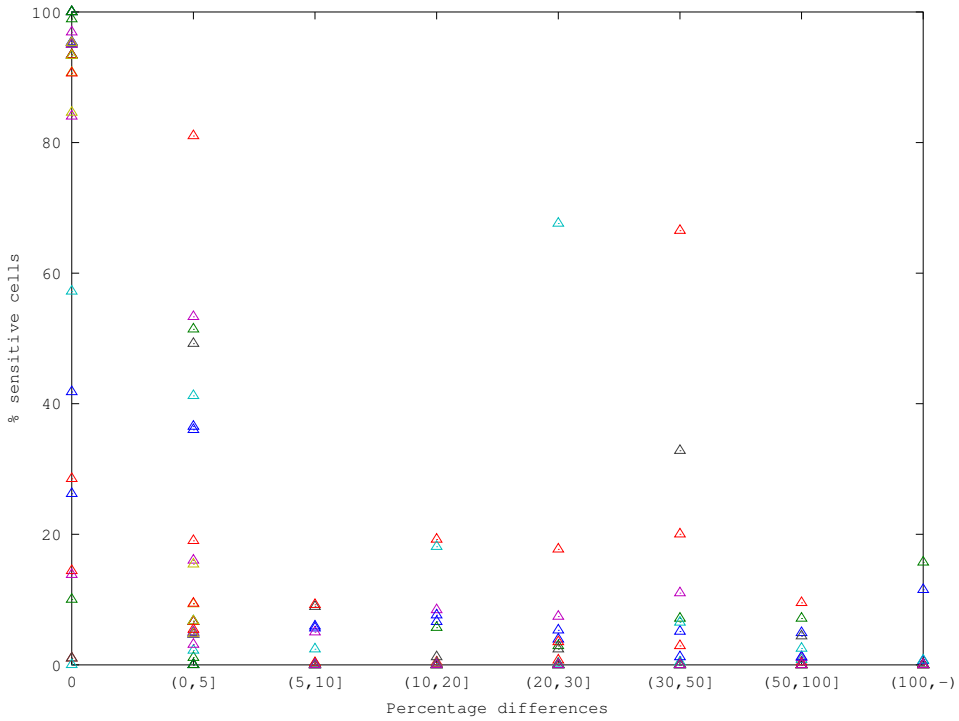

CTA is not safe for scenario B3

ℓ_2 is safer than ℓ_1

Results for scenario C

ℓ_1

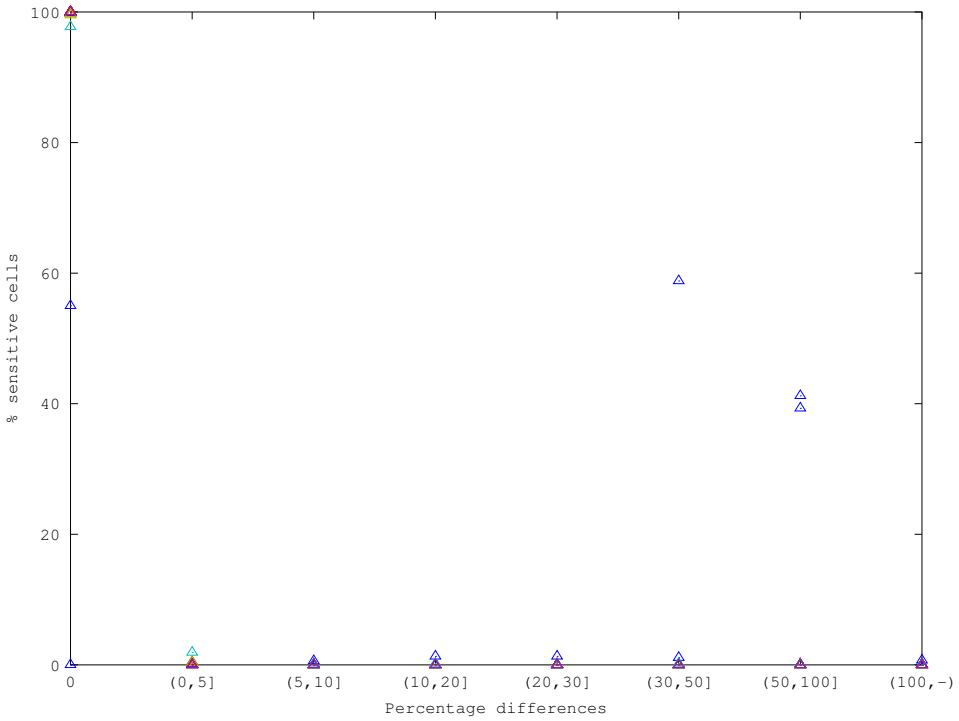

ℓ_2

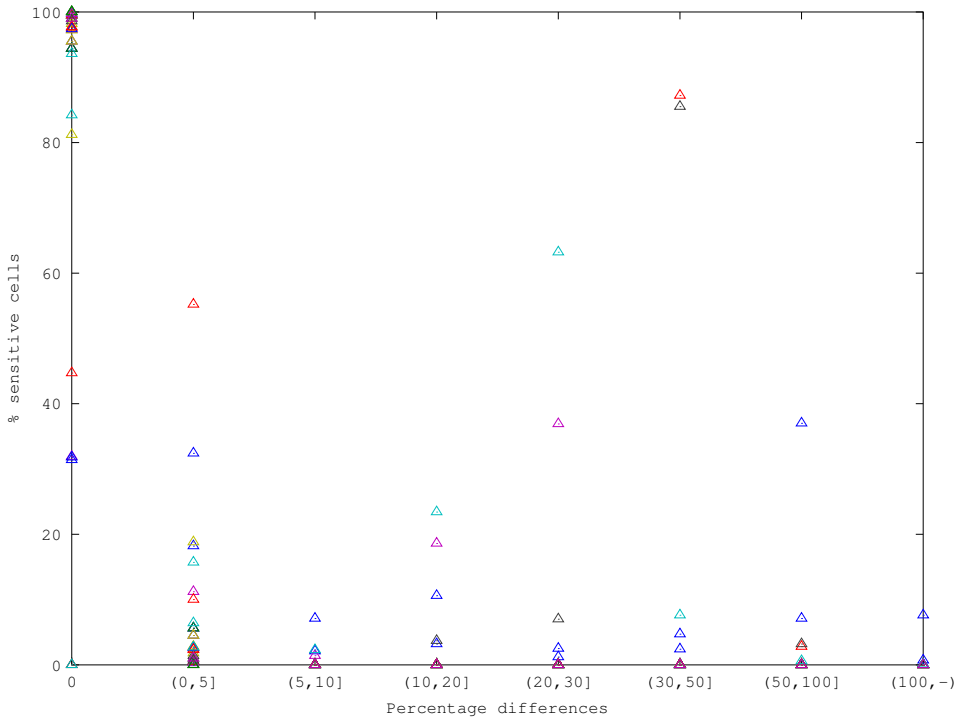

CTA is not safe for scenario C

ℓ_2 is safer than ℓ_1

Conclusions

- Scenarios B1 and B2 are safe. The estimate $\hat{a}_i$ was never equal to the true cell value a_i, and the distribution is not concentrated on the left intervals.

Conclusions

- Scenarios B1 and B2 are safe. The estimate $\hat{a}_i$ was never equal to the true cell value a_i, and the distribution is not concentrated on the left intervals.
- For scenarios B3 and C the attacker was able to re-compute in almost 100% of the cases the original values a.

Conclusions

- Scenarios B1 and B2 are safe. The estimate $\hat{a}_i$ was never equal to the true cell value a_i, and the distribution is not concentrated on the left intervals.
- For scenarios B3 and C the attacker was able to re-compute in almost 100% of the cases the original values a.
- ℓ_2 seems to reduce more than $\ell_1$the disclosure risk: the distribution is more left-skewed for ℓ_2 in scenarios B1 and B2.

Conclusions

- Scenarios B1 and B2 are safe. The estimate $\hat{a}_i$ was never equal to the true cell value a_i, and the distribution is not concentrated on the left intervals.
- For scenarios B3 and C the attacker was able to re-compute in almost 100% of the cases the original values a.
- ℓ_2 seems to reduce more than $\ell_1$the disclosure risk: the distribution is more left-skewed for ℓ_2 in scenarios B1 and B2.
- Conclusion: CTA is safe, unless the attacker has good information about the protection levels, protection senses, set of sensitive cells, and lower and upper bounds. However the knowledge of such big amount of information by the attacker may be a strong assumption.

Thanks for your attention!