

Statistical Disclosure Control: Moving Forward

Journal of Official Statistics

2025, Vol. 41(3) 820–826

© The Author(s) 2025

Article reuse guidelines:

sagepub.com/journals-permissions

DOI: 10.1177/0282423X241312023

journals.sagepub.com/home/jof



Josep Domingo-Ferrer¹ , David Sánchez¹ and Krishnamurthy Muralidhar²

Keywords

privacy, anonymity, statistical disclosure control

1. Introduction

The very first privacy notions in official statistics can be found in the randomized response mechanism (RR; Warner 1965), which protects query answers by reporting true answers with probability p and false answers with probability $1-p$, and the *Dalenius desideratum* (Dalenius 1977), which states that nothing about any specific individual should be learnable from a publicly released statistical database that cannot be learned without access to the database. These foundations ushered in the statistical disclosure control (SDC) discipline (Adam and Worthmann 1989; Hundepool et al. 2012; Willenborg and De Waal 2012), whose goal is to provide methods for data anonymization.

As the needs for privacy-preserving data sharing increased, it was soon realized that (i) Dalenius-like privacy notions were too stringent to produce protected data with reasonable analytical utility, and (ii) non-interactive data releases were much more useful and needed than interactive query answering. Hence, the first practical approaches to anonymization focused on the release of statistical databases and were utility-first, where anonymization parameters are iteratively tried until

¹Department of Computer Science and Mathematics, CYBERCAT–Center for Cybersecurity Research of Catalonia, Universitat Rovira i Virgili, Tarragona, Catalonia, Spain

²Department of Marketing and Supply Chain Management, Price College of Business, University of Oklahoma, Norman, OK, USA

Corresponding author:

Josep Domingo-Ferrer, Department of Computer Science and Mathematics, CYBERCAT–Center for Cybersecurity Research of Catalonia, Universitat Rovira i Virgili, Av. Països Catalans 26, Tarragona, Catalonia 43007, Spain.

Email: josep.domingo@urv.cat



Creative Commons Non Commercial CC BY-NC: This article is distributed under the terms of the Creative Commons Attribution-NonCommercial 4.0 License (<https://creativecommons.org/licenses/by-nc/4.0/>) which permits non-commercial use, reproduction and distribution of the work without further permission provided the original work is attributed as specified on the SAGE and Open Access pages (<https://us.sagepub.com/en-us/nam/open-access-at-sage>).

sufficient analytical utility is empirically achieved while reducing the observed disclosure risk below a certain threshold. Examples of these heuristic mechanisms are data suppression, sampling, noise addition, microaggregation, data swapping, or value generalization.

To overcome the limitations of the empirical nature of classical SDC methods, the computer science community approached the problem in the late 90s from the privacy-first angle by introducing the notion of privacy model, by which an *ex ante* privacy condition is enforced using one or several SDC methods. Privacy models have the advantage of setting the privacy level by design, without having to empirically evaluate it. The first privacy model was *k*-anonymity (Samarati and Sweeney 1998), which states that the probability of re-identifying a record in a protected database by leveraging available background knowledge on the individual to whom the record corresponds should be at most $1/k$. *k*-Anonymity was followed by a number of extensions, such as *l*-diversity (Machanavajjhala et al. 2007), *t*-closeness (Li et al. 2010), and others, which mitigate the disclosure risk of confidential values. All privacy models in the *k*-anonymity family were designed to anonymize microdata sets, that is, data sets formed by records that correspond to individual respondents.

In 2006, a new privacy model named ϵ -differential privacy (DP) was proposed in the cryptographic community, which aimed to bring to data sharing a privacy guarantee based on indistinguishability by reformulating the old Dalenius desideratum in an even more stringent manner: “the risk to one’s privacy [...] should not substantially increase as a result of participating in a statistical database” (Dwork 2006).

DP was originally proposed for interactive statistical queries to a database, where a randomized query function (that returns the query answer plus some noise) satisfies ϵ -DP if, for any two databases that differ in one record, the respective query answers do not differ by more than an exponential factor of ϵ (called the budget). This ensures the presence or absence of any single record does not significantly affect query answers. The smaller ϵ , the higher the protection, with $\epsilon \leq 1$ being considered “safe” (Dwork 2011). The noise to be added to the answer to enforce a certain ϵ depends on the global sensitivity of the query to the presence or absence of any single record. One can see that the enforcing mechanism employed by DP is very similar to the RR forerunner, as both produce randomized answers to interactive queries with a predefined probability.

Thanks to its neat privacy guarantee, DP has been rapidly adopted by the research community, to the point that previous approaches tend to be regarded as obsolete. In fact, DP has become the *de facto* privacy standard in most data intensive areas, such as data analytics, statistics, or machine learning (Wood et al. 2018). However, while mild noise may suffice to enforce DP on aggregated statistical queries (e.g., averages), its stringent privacy guarantee requires a very large noise to be added for identity queries returning the contents of a specific record. In those cases, the noise needed to enforce any safe-enough ϵ is so large that the DP-protected outcomes become nearly random and, therefore, analytically useless. The latter is precisely the case of microdata releases, which are more in demand than

interactive queries (Domingo-Ferrer et al. 2021). The research community has devoted enormous efforts during the last eighteen years trying to fit (or bend) DP to those settings, with the arduous aim of reconciling the DP privacy guarantee with sufficient data utility.

2. Challenges

A shortcoming of the utility-first approach traditionally embraced by the SDC community, especially in official statistics, is that it requires (possibly several rounds of) empirical disclosure risk assessment to make sure the protected data are safe enough. Such an assessment is expensive, and it requires selecting plausible attack scenarios and parameters. In recent years, there has been a tendency to seek *ex ante* privacy guarantees. Whereas official statisticians have mostly opted for synthetic data generation (Burnett-Isaacs et al. 2022), computer scientists have favored privacy models, especially DP. We next sketch the challenges associated with those *ex ante* approaches.

Synthetic data are artificial data preserving some of the statistical features of the original data. Their main attractive is that, being artificial, it would seem they can be released without privacy concerns. However, they are not without issues. If they are generated with the only requirement that certain statistics be exactly preserved, their utility is limited to those statistics, whereas other statistics or subdomain analyses are not preserved. Then one might wonder why not just publish the statistics to be preserved rather than a synthetic data set. On the other hand, if synthetic data are generated by fitting a statistical or machine learning (ML) model to the original data and then sampling that model, the issue of potential overfitting appears. Indeed, sampling an overfitted model is likely to yield synthetic data extremely close to the original: while this is good for utility, it is very bad for privacy (reidentification and attribute disclosure are likely). Overfitting is especially worrisome when using complex ML models, such as deep neural networks, whose millions of parameters may encode a snapshot of the confidential original data they are trained on.

The deployment of privacy models has mostly focused on DP, whereas the *k*-anonymity family of models have so far enjoyed less popularity. Yet, *k*-anonymity models modify the input microdata, from which consistent statistical outputs can be computed; in contrast, DP adds noise independently to each statistical output, which is likely to result in inconsistencies among outputs (Citro et al. 2022). A relevant case study in official statistics is the application of DP to the 2020 U.S. Census. This application has faced operational and scientific challenges (Garfinkel et al. 2018), which have resulted in significant delays in data release, the use of relaxed DP with budgets as large as $\epsilon = 39.907$ (far above the $\epsilon \leq 1$ recommended by the DP inventors), and very substantial inaccuracies in the DP-protected data (Kenny et al. 2021). Other DP-deployments by the big companies (Apple, Google, Microsoft, Meta, LinkedIn) have also resorted to large ϵ budgets (whose privacy guarantee is ineffectual) or to questionable composition assumptions, in an attempt to preserve utility.

Regarding the combination of synthetic data and DP, it amounts to training the synthetic data generation models on the original data in a differentially private manner. Since training is an iterative procedure, the sequential composition property of DP applies, which means that the budget used increases with the number of iterations (Abadi et al. 2016) and hence the privacy level decreases. Despite this budget increase, the impact of DP-training on the utility of the trained model remains significant (Domingo-Ferrer et al. 2021), which means that the utility of the synthetic data obtained by sampling the model can be seriously affected. In fact, in Blanco-Justicia et al. (2022) it is shown that anti-overfitting techniques applied during ML model training may yield a better privacy-utility trade-off than DP.

3. Directions

From the challenges above, it becomes clear that empirical disclosure risk assessment remains as unavoidable for synthetic or DP data as it was under the utility-first traditional approach to SDC. DP can only provide enough ex ante privacy to dispense with ex post disclosure risk assessment when the budget ϵ is really small (say below 1), but in this case utility is likely to be too low for most practical applications. When the budget is large, DP is little more than mild classical noise addition, and hence disclosure risk assessment is still needed. Future directions include:

- Ex ante privacy that does not require disclosure risk assessment remains a very desirable goal. Privacy models are the way to go but, given the difficulties of deploying DP in practice, exploring the other large family of privacy models (k -anonymity and its extensions) seems a good option. The k -anonymity family can provide reasonable utility in data releases and data collection with ex ante privacy guarantees that are easy to understand in terms of limiting identity disclosure (upper bound on the re-identification probability) and attribute disclosure (lower bounds on the diversity of confidential attributes). Further, k -anonymity and variants of it remain the best privacy models to anonymize some data types for which perturbative approaches would cause severe utility damage, such as stream data (Xiao and Tao 2007) or georeferenced data (Cremer et al. 2024).
- Coming up with realistic disclosure risk assessment in practical scenarios is necessary. Potential re-identification/disclosure attacks are sometimes invoked as an argument to justify the use of stronger SDC methods and privacy models. However, no successful attacks against properly anonymized data (e.g., those released by national statistical institutes) have been reported, which contradicts the pessimism in Ohm (2009). Much publicized re-identification attacks were mounted against data that had undergone mere pseudonymization or suppression of identifiers—the reidentification of the Governor of Massachusetts (Sweeney 1997), the AOL attack (Barbaro and Zeller 2006), and the Netflix Prize Dataset attack (Narayanan and Shmatikov 2008)—or against poorly anonymized data (De Montjoye et al.

2015; Sánchez et al. 2016). Also, the much discussed reconstruction of data compatible with the released statistical outputs does not in general imply re-identification (Muralidhar and Domingo-Ferrer 2023).

- When using ML models in data protection, especially to generate synthetic data, overfitting must be avoided. Overfitted models may memorize the entire original data and yield synthetic data extremely close to the original. One option is to use anti-overfitting methods. An alternative is to use only simple models for synthetic data generation, such as statistical models or basic ML methods (decision trees, random forests, etc.). The bottom line is that models used for synthetic data generation should not be too accurate if we do not want the resulting synthetic data to be too close to the original data.
- Non-perturbative SDC methods (such as data suppression, generalization, or sampling, see Hundepool et al. (2012)) are probably a better choice than perturbative methods in settings where perturbed data (e.g., DP data obtained via noise addition) might be unacceptable or invalidate analytical conclusions (e.g., in the medical domain or official statistics). Non-perturbative methods keep protected data truthful, and they are sufficient to enforce formal k -anonymity privacy.
- Alternative data sharing mechanisms that are privacy- and utility-preserving-by-design also deserve attention. Let us mention federated learning, anonymous channels, or secure multiparty computation as an alternative to sharing anonymized data.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This research was funded by the European Commission (project H2020-871042 “SoBigData++”), the Government of Catalonia (ICREA Acadèmia Prizes to J. Domingo-Ferrer and to D. Sánchez and grant 2021SGR-00115), MCIN/AEI/10.13039/501100011033 and “ERDF A way of making Europe” under grant PID2021-123637NB-I00 “CURLING,” and the EU’s NextGenerationEU/PRTR via INCIBE (project “HERMES” and INCIBE-URV Cybersecurity Chair).

ORCID iD

Josep Domingo-Ferrer  <https://orcid.org/0000-0001-7213-4962>

References

- Abadi, M., A. Chu, L. Goodfellow, et al. 2016. “Deep Learning with Differential Privacy.” Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, New York, NY, October 24–28.
- Adam, N. R., and J. C. Worthmann. 1989. “Security-Control Methods for Statistical Databases: A Comparative Study.” *ACM Computing Surveys* 21 (4): 515–56. DOI: <https://doi.org/10.1145/76894.76895>.

- Barbaro, M., and T. Zeller. 2006. "A Face is Exposed for AOL Searcher No. 4417749." *The New York Times*, August 9.
- Blanco-Justicia, A., D. Sánchez, J. Domingo-Ferrer, and K. Muralidhar. 2022. "A Critical Review on the Use (and Misuse) of Differential Privacy in Machine Learning." *ACM Computing Surveys* 55 (8): 1–16. DOI: <https://doi.org/10.1145/3547139>.
- Burnett-Isaacs, K., C. Girard, K. Sallier, et al. 2022. *Synthetic Data for Official Statistics: A Starter Guide*. United Nations Economic Commission for Europe.
- Citro, C., R. Prevost, and I. Youngs. 2022. "U.S. Political and Statistical Geography and Census Small-Area Statistics: A Primer." Working Draft 2.1, Georgetown University. <https://mdi.georgetown.edu/census-geographies-project/>
- Cremer, S., L. Jehmlich, and R. Lenz. 2024. "Masking Georeferenced Health Data-An Analysis Taking the Example of Partially Synthetic Data on Sleep Disorder." In *International Conference on Privacy in Statistical Databases (PSD 2024)*, edited by J. Domingo-Ferrer and M. Önen. New York: Springer.
- Dalenius, T. 1977. "Towards a Methodology for Statistical Disclosure Control." *Statistik Tidskrift* 15: 429–44.
- De Montjoye, Y.-A., L. Radaelli, V. K. Singh, and A. S. Pentland. 2015. "Unique in the Shopping Mall: On the Reidentifiability of Credit Card Metadata." *Science* 347 (6221): 536–9. DOI: <https://doi.org/10.1126/science.1256297>.
- Domingo-Ferrer, J., D. Sánchez, and A. Blanco-Justicia. 2021. "The Limits of Differential Privacy (and Its Misuse in Data Release and Machine Learning)." *Communications of the ACM* 64 (7): 33–5. DOI: <https://doi.org/10.1145/3433638>.
- Dwork, C. 2006. "Differential Privacy." Proceedings of the 33rd International Colloquium on Automata, Languages and Programming-ICALP 2006, Venice, Italy, July 10–14.
- Dwork, C. 2011. "A Firm Foundation for Private Data Analysis." *Communications of the ACM* 54 (1): 86–95. DOI: <https://doi.org/10.1145/1866739.1866758>.
- Garfinkel, S. L., J. M. Abowd, and S. Powazek. 2018. "Issues Encountered Deploying Differential Privacy." Proceedings of the 2018 Workshop on Privacy in the Electronic Society, Toronto, ON, Canada, October 15.
- Hundepool, A., J. Domingo-Ferrer, L. Franconi, et al. 2012. *Statistical Disclosure Control*. Hoboken, NJ: John Wiley & Sons.
- Kenny, C. T., S. Kuriwaki, C. McCartan, E. T. R. Rosenman, T. Simko, and K. Imai. 2021. "The Use of Differential Privacy for Census Data and Its Impact on Redistricting: The Case of the 2020 US Census." *Science Advances* 7 (1): eabk3283. DOI: <https://doi.org/10.1126/sciadv.abk3283>.
- Li, N., T. Li, and S. Venkatasubramanian. 2010. "Closeness: A New Privacy Measure for Data Publishing." *IEEE Transactions on Knowledge and Data Engineering* 22 (7): 943–56. DOI: <https://doi.org/10.1109/TKDE.2009.139>.
- Machanavajjhala, A., J. Gehrke, D. Kiefer, and M. Venkitasubramanian. 2007. "L-Diversity: Privacy Beyond k -Anonymity." *ACM Transactions on Knowledge Discovery from Data* 1 (1): 3. DOI: <https://doi.org/10.1145/1217299.1217302>.
- Muralidhar, K., and J. Domingo-Ferrer. 2023. "Database Reconstruction is Not So Easy and is Different from Reidentification." *Journal of Official Statistics* 39 (3): 381–98. DOI: <https://doi.org/10.2478/jos-2023-0017>.
- Narayanan, A., and V. Shmatikov. 2008. "Robust De-Anonymization of Large Sparse Datasets." 2008 IEEE Symposium on Security and Privacy (SP 2008), Oakland, CA, USA, May 18–22.
- Ohm, P. 2009. "Broken Promises of Privacy: Responding to the Surprising Failure of Anonymization." *UCLA Law Review* 57: 1701.

- Samarati, P., and L. Sweeney. 1998. "Protecting Privacy When Disclosing Information: k -Anonymity and Its Enforcement Through Generalization and Suppression." Technical Report, SRI International.
- Sánchez, D., S. Martínez, and J. Domingo-Ferrer. 2016. "Comment on 'Unique in the Shopping Mall: On the Reidentifiability of Credit Card Metadata.'" *Science* 351 (6279): 1274. DOI: <https://doi.org/10.1126/science.aad9295>.
- Sweeney, L. 1997. "Weaving Technology and Policy Together to Maintain Confidentiality." *Journal of Law, Medicine and Ethics* 25: 98–110. DOI: <https://doi.org/10.1111/j.1748-720X.1997.tb01885.x>.
- Warner, S. L. 1965. "Randomized Response: A Survey Technique for Eliminating Evasive Answer Bias." *Journal of the American Statistical Association* 60 (309): 63–9. DOI: <https://doi.org/10.1080/01621459.1965.10480775>.
- Willenborg, L., and T. De Waal. 2012. *Elements of Statistical Disclosure Control*. New York: Springer Science & Business Media.
- Wood, A., M. Altman, A. Bembenek, et al. 2018. "Differential Privacy: A Primer for a Non-Technical Audience." *Vanderbilt Journal of Entertainment and Technology Law* 21: 209.
- Xiao, X., and Y. Tao. 2007. "M-Invariance: Towards Privacy Preserving Re-Publication of Dynamic Datasets." Proceedings of the 2007 ACM SIGMOD International Conference on Management of Data, Beijing, China, June 11–14.

Received: November 22, 2024

Accepted: December 19, 2024