

Measuring Privacy: Critical Reflections and Directions for a Metrology-Based Approach

Eulalia Balestrieri¹, Iliaria Amelia Caggiano², Francesco Picariello³, Ioan Tudosa¹

¹ *University of Sannio, Benevento, Italy, {balestrieri, ioan.tudosa}@unisannio.it*

² *University Suor Orsola Benincasa, Naples, Italy, ilaria.caggiano@unisob.na.it*

³ *Universitas Mercatorum, Rome, Italy, francesco.picariello@unimercatorum.it*

Abstract – Privacy measurement in digital systems lacks a standardised metrological framework to ensure reliable and comparable assessments. A metrological approach ensures that privacy measurements are reliable, reproducible, and comparable over time and across different contexts. In this work most common privacy metrics, including k-anonymity, ℓ -diversity, t-closeness, differential privacy, and mutual information, are critically evaluated, identifying their strengths and limitations from a metrological perspective. Initial directions and open challenges toward a metrology-based approach to privacy measurement are outlined, too.

I. INTRODUCTION

The increasing reliance on personal data in modern digital systems has made privacy protection a central concern across multiple domains. Although regulatory frameworks have established strong principles for data protection, the quantitative assessment of privacy remains an open challenge. Numerous privacy metrics have been proposed, aimed at attempting to quantify the level of protection afforded to individuals against data disclosure, but they often lack the foundational properties expected of reliable measurement tools. The term "metric" is often used loosely, without a consistent reference to measurement theory.

Metrology provides a well-established framework for evaluating the quality of metrics. Key concepts to do that are included in the International Vocabulary of Metrology (VIM) [1] and the Guide to the Expression of Uncertainty in Measurement (GUM) [2]. Although developed for physical measurements, they are increasingly being adapted to digital and informational domains. Efforts by institutions such as PTB (Physikalisch-Technische Bundesanstalt) through its "Metrology for Digital Transformation" initiatives (e.g., D-SI metadata model and European Metrology Cloud projects) [3], and NIST (National Institute of Standards and Technology), through its extension of legal metrology principles such as the "chain of trust" to digital measuring systems, are beginning to explore the intersection of metrology and digital trust [4]. Privacy is governed by legal frameworks

like the European General Data Protection Regulation (GDPR) [5], which emphasise principles such as data minimisation, purpose limitation, and accountability. However, these regulations do not prescribe specific quantitative metrics, further highlighting the need for rigorous and standardised measurement approaches. This paper investigates the applicability of metrological principles to the most commonly used privacy metrics, aiming to identify gaps, promote rigour, and suggest directions for future standardisation. The paper is structured as follows. First, a brief overview of existing privacy regulations and standards is provided, followed by a description of the most commonly used privacy metrics. Next, the relationship between these privacy metrics and the metrological approach is examined, with particular attention to the challenges and requirements for building a metrology-based framework to measure privacy. Finally, conclusions are drawn and directions for future work are outlined.

II. EXISTING PRIVACY REGULATIONS AND STANDARDISATION

Modern data protection regulations emphasise accountability, that is, the obligation of data controllers and processors to demonstrate compliance with legal and ethical obligations. This entails not only the implementation of appropriate technical and organisational measures but also the ability to prove their effectiveness. Common themes concern emphasis on proactive measures and continuous assessment, encouragement (or requirement) of risk-based evaluations, focus on documentation, and demonstrable safeguards.

However, these requirements often lack concrete guidance on how to measure or verify the effectiveness of privacy measures.

Regulatory frameworks and standardisation efforts worldwide have established the legal and procedural basis for privacy protection, but often without clearly defined, quantitative metrics. This lack of metrological rigour hinders the objective assessment and comparability of privacy-preserving techniques.

The European Union's GDPR (Regulation (EU) 2016/679) [5] is among the most influential data protection

regulations. While it emphasises principles such as data minimisation, purpose limitation, and accountability, it does not prescribe specific measurement methods. However, Recital 78 encourages the use of privacy-enhancing technologies and data protection by design, indirectly requiring accessible levels of privacy.

ISO/IEC 29100 [6] provides a high-level privacy framework but lacks operational metrics. ISO/IEC 27701 [7] extends ISO 27001 with controls for Privacy Information Management Systems (PIMS), introducing risk-based approaches. Nevertheless, privacy risk is typically evaluated qualitatively or semi-quantitatively, based on likelihood and impact, without standardised units of measure or uncertainty.

The NIST Privacy Framework [8] adopts a structured, flexible, and voluntary model to help organisations identify and manage privacy risks. Although it includes control categories like “Minimize Data Collection and Retention” or “Manage Data Processing,” it does not define universal or standardised privacy metrics. However, related NIST documents (e.g., NISTIR 8062 [9]) discuss the concept of privacy engineering and encourage the development of “measurable privacy goals”, leaving implementation to specific domains.

The International Telecommunication Union has published several technical standards (e.g., ITU-T X.1149 [10]) on privacy risk management and X.1205 [11] on cybersecurity, which include references to impact assessments and, in some cases, probabilities of leakage. However, they fall short of providing reproducible and standardised measurement protocols.

Overall, while these frameworks recognise the importance of assessing privacy risk, few provide formalised, metrology-compliant metrics, nor do they establish traceable, uncertainty-aware measurement practices.

III. EXISTING PRIVACY METRICS

Over the past two decades, a variety of privacy metrics have been developed across academic and applied settings. There have been so many different and complex metrics, making the choice among them very challenging, that rather than choosing, it is often preferred to propose and create new metrics. Consequently, the comparison among different privacy systems, in terms of the assured degree of privacy to the users, as well as the level of protection assured by privacy-enhancing technologies, becomes impossible to carry out.

There is not the same approach in defining the privacy metric and its requirements and functioning. Not all proposed metrics have a mathematical base [12]. Multiple areas are of concern regarding assured privacy levels, as communication systems, databases, location services, smart metering, social networks, etc. [12]. Therefore, many privacy metrics are context-dependent and often lack comparability and interpretability [12].

The most common and used metrics for privacy are k-

anonymity, ℓ -diversity, t-closeness, Differential Privacy (DP) and entropy-based metrics.

k-anonymity ensures that each individual in the dataset is indistinguishable from at least k-1 other individuals with respect to Quasi-Identifying attributes (QI) (e.g., age, gender). The dataset is transformed (generalisation, suppression) so that each combination of QI appears at least k times. It does not protect against attacks on homogeneous sensitive attributes within a group [12],[13],[14],[15].

ℓ -diversity extends k-anonymity by introducing the concept of diversity in the sensitive attributes within each anonymous group. Each anonymous group must have at least ℓ different sensitive attribute values (e.g., different disease categories). It can be bypassed if the sensitive values are semantically similar (e.g., "severe" vs. "very severe"). Although it provides a level of heterogeneity, the definition of diversity can be subjective [12],[13],[15].

t-Closeness requires each anonymous group to have a distribution of sensitive attributes that is "close" (within a threshold t) to the distribution of the entire dataset. It uses statistical distance metrics (e.g., Earth Mover's Distance) to quantify proximity. It requires more complex calculation, dependent on the choice of threshold t [12],[15].

DP defines a formal limit on the amount of information that can be inferred from an output, regardless of prior knowledge. It adds calibrated noise (e.g., Laplacian noise) to statistical query responses to ensure that the presence or absence of an individual has a limited impact [12],[15]. The scale of the noise added to outputs is controlled by the parameter ϵ ; a smaller ϵ requires injecting more noise, making it harder to infer whether any individual's data was included. A larger ϵ allows less noise, improving utility but reducing privacy protection.

Entropy-based metrics quantify indeterminacy about private information (e.g., Shannon entropy, conditional entropy); higher entropy generally implies higher privacy. They are sensitive to the granularity of data representation and may overestimate privacy if the attacker's knowledge is not uniform. Mutual Information (MI) is a special case of entropy-based metrics that measures information leakage. It quantifies the reduction in indeterminacy (entropy) of one variable due to knowledge of another. Lower MI indicates better privacy. It requires accurate estimation of joint probability distributions, which is challenging for large or continuous datasets, is sensitive to assumptions about data distribution and computationally intensive for high-dimensional data [12],[15]. It is worth noting that k-anonymity, ℓ -diversity, t-closeness and DP are in fact privacy models or mechanisms, originally conceived as techniques to protect data. MI, instead, is an evaluation metrics that quantify information leakage.

IV. METROLOGY AND PRIVACY METRICS

Measuring privacy is complex because privacy is a

multidimensional concept and often context-dependent. A metrological approach for privacy metrics ensures that privacy measurements are reliable, reproducible, and comparable over time and across different contexts. When metrics are based on metrological principles, other people or systems can replicate the measurements and obtain similar results, which is essential for verifying and validating privacy levels. Standardising metrics allows comparing the results of different privacy solutions or technologies, making it easier to choose the best strategy or technology. Accurate and standardised measurements help companies, regulators, and developers make informed decisions and correctly assess privacy risks and protections. Moreover, a metrological method enables monitoring how metrics change over time, useful for evaluating the effectiveness of implemented measures and for continuous improvement.

A metrological compliant privacy metric is required to have a well-defined measurand (the quantity being measured) and a meaningful estimate of uncertainty.

A valid measurement also requires repeatability (the ability to obtain consistent results under the same conditions) and reproducibility (the ability to replicate results across different settings and analysts). In the following subsection, the measurands of the most common privacy metrics and their repeatability and reproducibility are analysed from a metrological point of view.

A. Privacy Metrics measurands

Table 1 provides concise definitions of these measurands, summarising the quantities that each metric is intended to measure according to current practice in the literature. k-anonymity is easy for people to interpret, while DP and MI require a technical background. DP and MI are theoretically precise, while k-anonymity, ℓ -diversity, and t-closeness rely on contextual definitions of QI and sensitive attributes, which can vary.

For k-anonymity and ℓ -diversity, the measurand is not a continuous quantity but a cardinality, and the scale is discrete.

For t-closeness, the measurand is a distance between distributions and depends on the chosen distance function; the unit is conventional (dimensionless real value).

For DP, ϵ is unitless (log of the odds ratio). It is a dimensionless quantity on a ratio scale.

For MI, the amount of shared information is measured in bits.

The interpretation of metric values differs across methods. For k-anonymity and ℓ -diversity, higher values correspond to greater privacy. Conversely, t-closeness achieves stronger privacy with lower values.

In practice, the choice of parameter values depends on the dataset and the desired privacy level.

The k-anonymity threshold is most commonly set at 3 or 5 [13]. Values above 5 are uncommon, and values exceeding 15 are rare [14]. However, in other cases, it can be found

that considering values of 10 or lower are insufficient to ensure adequate data protection [15]. It is worth noting that the selection of k must carefully balance privacy protection with data utility, adapting to the specific context and objectives of the study.

Similar to k-anonymity, for ℓ -diversity, there is no ideal value for ℓ , although it is typically less than or equal to k and greater than 1 [15],[16].

Values of t are typically between 0 and 1 [17],[18]. Lower values of t (for example, 0.1 or 0.05) mean that the distributions within each group closely match the overall distribution, resulting in stronger privacy [17],[18],[19]. Higher values of t (for example, 0.2 or 0.3) permit greater variation, which can improve data utility but reduce privacy [17],[18]. The suitable value of t depends on the specific dataset, the sensitivity of the attribute, and the desired privacy level. It is typically determined through experimentation, aiming to balance data privacy with usefulness [17],[18].

Concerning DP, smaller values of ϵ offer greater privacy protection but may reduce data accuracy. Larger values of ϵ may improve accuracy but reduce privacy protection. The choice of ϵ remains a debated issue in the community. While theoretical works generally recommend small values ($\epsilon < 1$) to guarantee strong privacy, in practice, larger values are often used. Values in the range 0–5 are considered conservative and provide robust protection in most contexts [20]. However, higher values ($5 < \epsilon < 20$) may still be acceptable in some practical settings [20].

For MI, small values (>0 , but low) express a small dependency, so a minimal risk of privacy leak. Large values instead express strong correlation and so a greater risk of inferring sensitive data.

Table 1. Metrics measurands

Metric	Measurand
k-anonymity	The minimum equivalence class size (number of records with identical QI after generalisation/suppression).
ℓ -diversity	The minimum number of distinct sensitive attribute values in any equivalence class.
t-closeness	The maximum statistical distance between the sensitive attribute distribution in any equivalence class and the global distribution.
Differential Privacy	The maximum change in the output distribution of a randomised algorithm when any one individual's data is changed, expressed by the parameter ϵ .
Mutual Information	The amount of information that observable attributes reveal about sensitive data.

Most privacy metrics are defined in terms of parameters chosen a priori, k in k-anonymity, ℓ in ℓ -diversity, t in t-

closeness, and ϵ in DP. These parameters are fixed by design and do not directly reflect the actual, realised protection of user data in a given dataset. For DP, ϵ represents a theoretical bound on leakage, not an observable quantity. Similarly, k , ℓ , and t are design thresholds rather than measurements of effective risk.

MI quantifies information leakage but can vary depending on dataset realisation and noise mechanisms, making it partially suitable for statistical treatment.

Existing privacy metrics are largely based on design parameters that do not directly quantify actual information leakage. They are often context-dependent, rely on abstract or discrete scales, and can have non-intuitive interpretations. By defining an observable measurand, privacy can be assessed through a direct, consistent, and quantitative indicator, enabling rigorous evaluation, reproducibility, and meaningful comparison across datasets and methods.

B. Repeatability and Reproducibility of Privacy Metrics

Existing privacy metrics were not originally conceived with requirements concerning repeatability and reproducibility in mind, which raises significant challenges for their adoption as standardised measures.

k -anonymity, ℓ -diversity, and t -closeness are formally well-defined metrics whose reported values are fully repeatable when applied to the same dataset. Being deterministic metrics, they are reproducible for definition [18],[21],[22],[23]. This doesn't mean that if the same k , ℓ , or t is achieved, there are no variations in the resulting dataset in the case of applying randomisation or heuristic tie-breaking, common in greedy or optimisation-based anonymisation methods [18].

Beyond algorithmic factors, the practical utility of these metrics as indicators of privacy is limited by external conditions. The actual risk of re-identification depends heavily on the adversary's background knowledge, access to external datasets, and data characteristics such as skewed distributions of sensitive attributes or attribute correlations [18],[21],[22],[23]. Empirical measures of re-identification risk can vary significantly across datasets, implementations, or even repeated applications, limiting the metrics' reliability as indicators for real-world privacy protection [18],[21],[22],[23].

DP introduces a formal guarantee based on the privacy parameter ϵ , ensuring that the influence of any single individual on the output is strictly bounded. However, different mechanisms calibrated to the same ϵ can produce substantially different empirical outcomes in practice, due to variations in function sensitivity, noise distribution, or implementation details. While ϵ defines a theoretical upper bound on privacy loss, it does not directly quantify the actual risk or information leakage in a specific dataset or execution. Moreover, the intrinsic randomness of DP mechanisms means that repeated executions on the same dataset may yield different outputs, affecting the

operational reproducibility of reported results. Consequently, transparent and standardised evaluation protocols are essential to assess the effective privacy and ensure comparability and reproducibility of empirical measurements across studies [24].

Empirical estimation of Mutual Information (MI) can exhibit high variability due to factors such as sample size, underlying data distribution, and the choice of estimator. Methods like k -nearest neighbours (kNN) or neural-based approaches such as MINE may yield different values even on the same dataset, depending on parameter settings [25],[26]. Unlike the theoretical MI, which is well-defined mathematically, empirical estimates carry intrinsic uncertainty that must be quantified and reported. Standardised protocols for MI estimation, including consistent data preprocessing, estimator selection, and uncertainty evaluation, are therefore essential to ensure reproducible and comparable measurements across studies and datasets.

Overall, these limitations highlight that, despite their solid theoretical foundations, existing privacy metrics exhibit substantial variability in empirical estimation and limited metrological reproducibility, and therefore cannot be considered fully consistent or universally comparable measures of real-world privacy risk.

V. TOWARD A METROLOGY-BASED FRAMEWORK

Privacy should be defined as a physical informational quantity, that is, something that can be observed, simulated, or calculated in a replicable way. Traditional metrics such as k -anonymity, ℓ -diversity, t -closeness, DP, or MI are relative definitions: they describe properties of data with respect to a dataset or algorithm, but they do not rely on a universally accepted physical or informational unit. Each of these metrics measures a different aspect of privacy: k -anonymity expresses indistinguishability between records; ℓ -diversity ensures protection against inference of sensitive attributes; t -closeness preserves the statistical distribution of attributes; DP captures the probability of inference for a single record; MI quantifies the amount of information revealed. However, there is no way to combine them into a single, consistent number.

Different metrics capture different dimensions of privacy, but without a common unit or scale, it is not possible to rigorously compare one solution with another. If the value of a metric is just a parameter set upstream (for instance, k , ℓ , t , or ϵ), it makes no sense to speak about uncertainty regarding it. Uncertainty only makes sense with respect to quantities that are actually measured or observable, that is, with respect to the effective privacy that a system guarantees. Rather than referring to fixed parameters, it would be more useful and effective to measure observed privacy or actual risk. All systems should therefore be evaluated with the same observable quantity and the same protocol. Comparisons between systems must be based on

observable measurement $\pm$ uncertainty, and this requires a metric that is an observable quantity on real data, reflecting the actual risk or the actual amount of information revealed.

The measurement must represent how effectively a system protects user data, regardless of the nominal parameter chosen. In order to be scientifically rigorous, such a metric must be: observable on real data (measurable); comparable across different systems; and quantifiable with uncertainty. Building this metric by combining the existing ones can have some advantages: it leverages results that are already consolidated and recognised by the community, it is easier to be accepted by standardisation bodies such as ISO or NIST, and it covers various attack scenarios. On the other hand, each metric has different semantics, making them difficult to compare. Adding uncertainty becomes complicated, and there is the risk of creating a “patchwork index” without true metrological consistency.

By inventing a new metric from scratch, defined metrologically, it could be designed from the outset as an observable measurand with associated uncertainty. This would make it easier to build a common scale and would allow direct comparisons between different systems and technologies. However, such a new metric would require significant effort in theoretical modelling and experimental validation, and it could be perceived as too new, struggling to become standardised, unless the community is convinced to converge toward it.

A combination of approaches may be appropriate to integrate different measurement perspectives into a composite indicator; multiple risk dimensions can be joined into a single observable measure. A standard measure of effective privacy could be derived from a weighted or aggregated combination of existing metrics, after transforming them into observable and normalised quantities. Such a composite measure would provide several advantages: different types of attacks or inferences are covered; the risk of overestimating protection based on a single metric is reduced; heterogeneous systems can be compared with a single index.

Nevertheless, several common metrological challenges remain. First, the definition of the measurand: current metrics do not directly measure the observable phenomenon of interest, such as the probability of re-identification or the extent of information loss. It is therefore necessary to clearly define what is being observed and in what units it is measured. Second, the problem of common units of measurement: each metric currently uses different scales, such as number of records, probabilities, distances, or bits of information. To compare different systems, a normalised scale is required, for example, an index ranging from 0 to 1. Third, the problem of associated uncertainty: traditional privacy metrics do not provide a way to estimate uncertainty. A metrological approach instead requires the identification of uncertainty

sources and repeatable procedures to quantify the reliability of the measurement. Fourth, the issue of traceability: standardised test datasets and attack models must be created to achieve traceability. Finally, universal applicability: each metric has been designed for a specific context (tabular data, query releases, information-theoretic channels), while a privacy metric must be adaptable to different scenarios while still remaining interpretable.

Traditional privacy metrics, therefore, cannot be directly used as measurands in a metrological sense. By defining observable quantities that reflect real privacy risk and combining them into a standardised metric, we can achieve a measurement-based, uncertainty-aware framework for evaluating privacy systems. This approach would bridge the gap between theoretical guarantees and practical, quantitative assessment. Attempting to stitch together metrics born in different eras and contexts, such as k-anonymity for relational databases, DP for streaming or query systems, and MI for communication channels, risks creating an inelegant compromise that is difficult to standardise. Instead, a composite metric of effective privacy, grounded in observable risk and accompanied by an uncertainty estimate, could provide a quantitative, comparable, and reproducible measure of privacy protection. Such an approach could facilitate the integration of privacy metrics into existing regulatory frameworks and standards, enabling the benchmarking of privacy-preserving systems according to a measurable and standardised criterion, evidence-based compliance assessment for data protection regulations such as GDPR or ISO/IEC standards, and consistent evaluation across heterogeneous systems and datasets, improving transparency and accountability. By bridging the gap between theoretical guarantees and empirically measured privacy, the combined metric could serve as a normative tool, guiding both technical design and regulatory verification of privacy-preserving technologies. Initial steps for an interdisciplinary effort to bring metrological thinking into privacy evaluation may include defining theoretical measurands aligned with privacy goals, creating protocols for uncertainty evaluation in privacy computations, and aligning privacy measurement methods with international metrology standards such as the GUM and the VIM.

VI. CONCLUSIONS AND FUTURE WORK

Quantifying privacy remains a complex challenge that requires a solid measurement foundation. This paper outlines the essential requirements for metrologically compliant privacy metrics and critically examines the limitations of current approaches. Future work will focus on empirical studies and case analyses, paving the way for proposals aimed at standardisation and integration with existing measurement science frameworks. In particular, observable quantities for each base metric, k-anonymity, ℓ -diversity, t-closeness, DP, and MI, will be investigated

to identify which observables most accurately capture effective privacy under realistic threat models. Various strategies for combining these observables into a single index will be explored, including weighted schemes, conservative minima, and multi-criteria optimisation. Subsequently, procedures to estimate the uncertainty of the composite metric will be developed, considering how dataset characteristics, mechanism randomness, and attacker models influence uncertainty estimates. By applying a metrology-based framework, privacy can be transformed into a comparable and reproducible quantity, enabling both practitioners and regulators to quantify, compare, and enforce privacy in a consistent and scientifically grounded manner.

REFERENCES

- [1] JCGM 200:2012, *International Vocabulary of Metrology*, 2012. doi:10.59161/JCGM200-2012
- [2] JCGM 100:2008(E), *Evaluation of Measurement Data – Guide to the Expression of Uncertainty*, 2008. doi:10.59161/JCGM100-2008E
- [3] F.Thiel et al., “A Digital Quality Infrastructure for Europe: The European Metrology Cloud,” *PTB-Mitteilungen*, vol.127, no.4, 2017, doi:10.7795/310.20170404
- [4] NIST, “Legal Metrology Meets the Digital Age,” 2025. <https://www.nist.gov/news-events/news/2025/07/legal-metrology-meets-digital-age>
- [5] European Union, *Regulation (EU) 2016/679 – GDPR*, 2016.
- [6] ISO/IEC 29100:2011, *Information Technology – Security Techniques – Privacy Framework*, ISO, 2011.
- [7] ISO/IEC 27701:2019, *Extension to ISO/IEC 27001 and ISO/IEC 27002 for Privacy Information Management*, ISO, 2019.
- [8] NIST, “NIST Privacy Framework: A Tool for Improving Privacy through Enterprise Risk Management”, 2020.
- [9] NIST, “An Introduction to Privacy Engineering and Risk Management in Federal Systems,” *NISTIR 8062*, 2017.
- [10] ITU-T, “Risk Management and Risk Assessment Techniques in Information Security Management,” *ITU-T X.1149*, 2011.
- [11] ITU-T, “Overview of Cybersecurity,” *ITU-T X.1205*, 2008.
- [12] I.Wagner, D.Eckhoff, “Technical Privacy Metrics: A Systematic Survey,” *ACM Comput. Surv.*, vol.51, no.3, 2018, doi:10.1145/3168389
- [13] T.Benschop, M.Welch, “Statistical Disclosure Control for Microdata: A Practice Guide”, 2025. <https://sdcpractice.readthedocs.io/en/latest/>
- [14] K.El Emam, F.K.Dankar, “Protecting Privacy Using k-Anonymity,” *J. Am. Med. Inform. Assoc.*, vol.15, no.5, 2008, doi:10.1197/jamia.M2716
- [15] R. Aufschläger et al., “Anonymization Procedures for Tabular Data,” *Information*, vol.14, no.9, 2023, doi:10.3390/info14090487
- [16] Utrecht University, *Data Privacy Handbook*, May 2025. <https://utrechtuniversity.github.io/dataprivacyhandbook/disclaimer.html>
- [17] D.Roy, S.K.Jena, “Determining t in t-Closeness Using Multiple Sensitive Attributes,” *Int. J. Comput. Appl.*, vol.70, 2013, doi:10.5120/12179-8291
- [18] N.Li et al., “t-Closeness: Privacy Beyond k-Anonymity and l-Diversity,” *ICDE*, 2007, pp.106–115, doi:10.1109/ICDE.2007.367856
- [19] H.Liang, H.Yuan, “On the Complexity of t-Closeness Anonymization and Related Problems,” *Data Structures and Algorithms*, 2013. <https://arxiv.org/abs/1301.17511>
- [20] J.Near, D.Darais, “Differential Privacy: Future Work & Open Challenges,” *NIST Cybersecurity Insights Blog*, 2022. <https://www.nist.gov/blogs/cybersecurity-insights/differential-privacy-future-work-open-challenges>
- [21] J.Domingo-Ferrer, V.Torra, “A Critique of k-Anonymity,” *ARES*, 2008, pp.990–993, doi:10.1109/ARES.2008.97
- [22] K.El Emam et al., “Re-identification Attacks on Health Data,” *PLoS ONE*, vol.6, no.12, 2011, doi:10.1371/journal.pone.0028071
- [23] S. Li et al., “Reidentification Risk in Panel Data,” *Inf. Syst. Res.*, vol.34, no.3, 2022, doi:10.1287/isre.2022.1169
- [24] NIST, *Special Publication 800-226: Guidelines for Evaluating Differential Privacy Guarantees*, 2025. doi: 10.6028/NIST.SP.800-226
- [25] M. I. Belghazi et al., “Mutual Information Neural Estimation,” *ICML*, 2018, arXiv:1801.04062
- [26] P. Zhao, L. Lai, “Analysis of KNN Information Estimators,” *IEEE Trans. Inf. Theory*, vol.66, no.6, 2020, doi:10.1109/TIT.2019.2945041