Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Adversarial Evaluation of a Two-Layer Anonymization Pipeline Against Record-Linkage Attacks

Mohammed Sayim Khalil

sayimkhalil@halic.edu.tr

Haliç University https://orcid.org/0000-0002-1539-5629

Research Article

Keywords: adversarial evaluation, de-anonymization, differential privacy, Fellegi–Sunter record linkage, GDPR, heuristic pipeline, k-anonymity, privacy-utility balancing, re-identification risk, selective anonymization, t-closeness, two-layer pipeline

Posted Date: August 21st, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10766703/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

Adversarial Evaluation of a Two-Layer Anonymization Pipeline Against Record-Linkage Attacks

Mohammed Sayim Khalil, Senior member IEEE
Department of Computer Engineering, Halic University, Istanbul, Turkey sayimkhalil@halic.edu.tr • https://orcid.org/0000-0002-1539-5629

***Abstract*—This paper evaluates the empirical security of a two-layer separate-mechanism anonymization pipeline against a realistic record-linkage adversary. Syntactic privacy constraints (k-anonymity, l-diversity, t-closeness) are applied to the published record-level dataset D*, while differential privacy (DP) is applied separately to aggregate statistical query outputs. No joint formal guarantee between the two mechanisms is claimed or provided; Theorem 1 (Section V-H) formalises why no such guarantee can hold under arbitrary auxiliary information, motivating direct adversarial evaluation in place of a composed proof. The pipeline applies a four-tier attribute sensitivity taxonomy with default thresholds and a deterministic k-first greedy relaxation; its resistance to de-anonymization is then measured with a Fellegi–Sunter record-linkage attacker using non-overlapping auxiliary data and EM-based parameter estimation, varying adversary background knowledge (none, partial, full) and scheme knowledge (known, unknown). Methodological validation covers three public datasets under this attack model: NHANES-based synthetic demographic-health survey data, Sparkov financial transactions, and MDC mobility trajectories. A synthetic t-closeness demonstration of resistance to skewness attacks is provided in Appendix B. All thresholds are default starting points with cross-domain sensitivity analysis. A public code repository (anonymised for review; Appendix A) contains the complete pipeline implementation together with the Fellegi–Sunter adversary used for evaluation.**

***Index Terms*—adversarial evaluation, de-anonymization, differential privacy, Fellegi–Sunter record linkage, GDPR, heuristic pipeline, k-anonymity, privacy-utility balancing, re-identification risk, selective anonymization, t-closeness, two-layer pipeline.**

Manuscript received June 2026. M. S. Khalil is with the Department of Computer Engineering, Halic University, Istanbul, Turkey (e-mail: sayimkhalil@halic.edu.tr). ORCID: https://orcid.org/0000-0002-1539-5629.
Legal disclaimer: This paper describes technical methods for data anonymization and pseudonymization. It does not provide legal advice. Organizations should consult qualified legal counsel for GDPR, HIPAA, or other regulatory compliance.

I. INTRODUCTION

THIS paper presents a two-layer separate-mechanism anonymization pipeline and evaluates its resistance to a realistic record-linkage adversary. Syntactic privacy constraints (k-anonymity, l-diversity, t-closeness) are applied to the published record-level dataset D*, while differential privacy (DP) is applied to aggregate statistical query outputs. The two layers are not composed into a joint formal guarantee; they provide independent protections under different adversary assumptions. Theorem 1 (Section V-H) provides a formal non-composability result derived from Kifer and Machanavajjhala [15], which is precisely why the paper substitutes empirical adversarial evaluation — a Fellegi–Sunter record-linkage attacker with EM-based parameter estimation — for an unattainable joint proof. The contribution is primarily a security evaluation: a heuristic anonymization system stress-tested against a de-anonymization attacker across multiple domains, complemented by a precise characterisation of its theoretical limitations.

The paper makes the following specific contributions:

[1] A two-layer pipeline with a four-tier attribute sensitivity taxonomy, default thresholds (with sensitivity analysis), and a deterministic k-first greedy relaxation order (Section V).

[2] A corrected comparative positioning table (Table I) with accurate characterisation of Blowfish privacy's empirical evaluation.

[3] Fellegi–Sunter adversarial evaluation with Expectation-Maximisation (EM) parameter estimation using non-overlapping auxiliary data: separate NHANES cycles, non-overlapping Sparkov periods, and a random 80/20 split of MDC participants (limitation acknowledged for mobility) (Section VII-A).

[4] A concrete task-specification procedure for eliciting DP budget allocation weights when the analytical task is not predetermined (Section V-G).

[5] Cross-domain replication of relaxation order comparison (k-first, l-first, ε-first) with bootstrap confidence intervals (Section VI-F).

[6] A public code repository (anonymised for review; Appendix A) containing complete pipeline implementation, configuration files, and Fellegi–Sunter adversary code.

[7] Theorem 1: a formal non-composability result establishing that no joint (k, ε)-privacy guarantee can hold for the two-layer output under arbitrary auxiliary information (Section V-H).

Section II surveys related work. Section III characterises the privacy landscape. Section IV surveys techniques. Section V presents the pipeline. Section VI reports methodological validation. Section VII discusses threats and ethics. Section VIII concludes.

II. RELATED WORK AND POSITIONING

Attribute-level differential privacy has three key prior developments. Haeberlen et al. [12] introduced policy-based DP, in which per-attribute owner policies specify the privacy requirements the release mechanism must satisfy. Ebadi et al. [13] proposed Blowfish privacy, a generalisation of DP parameterised by a policy graph specifying which record pairs must be indistinguishable; their paper includes empirical evaluation on the Adult census dataset and taxi trajectory data. He et al. [14] proposed DPSD (Differential Privacy with Selective Data), explicitly allocating different ε budgets to attribute groups to preserve utility for low-sensitivity attributes.

This framework differs from all three on two dimensions: it integrates syntactic privacy constraints for record-level releases alongside DP for query outputs, and it provides a fully specified operational pipeline with deterministic relaxation, adversarial evaluation, and multi-domain validation. No claim is made to advance the formal theory of attribute-level DP beyond the explicit non-composability result in Section V-H.

Xiao and Tao [9] introduced anatomy; Nergiz et al. [10] proposed hybrid generalisation-suppression; Domingo-Ferrer et al. [11] surveyed attribute-level anonymisation comprehensively. Kifer and Machanavajjhala [15, Thm. 1] established that k-anonymity and DP do not compose into a joint guarantee under arbitrary auxiliary information—a result formalised here as Theorem 1. Charest [16] showed conditions under which k-anonymous preprocessing can be embedded in a DP framework. Cummings et al. [19] characterise utility-privacy bounds for statistical queries under user-level DP; as a theoretical analysis rather than an anonymisation system, this work is cited in Section IV-B for the utility-privacy frontier and is excluded from Table I.

The empirical security perspective adopted here follows the record-linkage attack literature rather than the syntactic-privacy literature. Gkoulalas-Divanis et al. [28] survey privacy-preserving record linkage techniques and the attacks used to stress-test them; Yin et al. [30] demonstrate a differential-analysis attack against Bloom filter encodings used for record linkage, illustrating that even encoded or masked identifiers remain vulnerable to structured linkage attacks under realistic assumptions. This paper adopts the same adversarial posture — rather than a further syntactic-privacy relaxation, it applies a Fellegi–Sunter record-linkage attacker directly to the pipeline's output across three domains, with the mobility domain additionally informed by Murakami et al. [29], who show that even generalised trajectory releases remain susceptible to re-identification via tensor-based factorisation.

TABLE I *COMPARATIVE POSITIONING OF ANONYMISATION FRAMEWORKS. 'SEQ. APPLIED' = SYNTACTIC CONSTRAINTS AND DP APPLIED SEQUENTIALLY TO DISTINCT OUTPUTS; NO JOINT GUARANTEE.*

Work	Attr. Select.	Seq. Applied	Pipeline	Emp. Eval.	Adv. Eval.
Haeberlen et al. [12]	Yes (policy DP)	No	No	No	No
Ebadi et al. [13] Blowfish	Yes (graph DP)	No	No	Yes (Adult,Taxi)	No
He et al. [14] DPSD	Yes (per-attr ε)	No	No	Yes	No
Xiao & Tao [9]	Partial	No	Partial	Yes	No
Domingo-Ferrer [11]	Yes (survey)	No	No	Yes (1 dom.)	No
This Work	Yes (4-tier)	Yes (no joint guar.)	Yes (5-stage)	Yes (3+App.B)	Yes (FS+EM)

III. BIG DATA: CHARACTERISTICS AND PRIVACY IMPLICATIONS

A. The Five Dimensions of Big Data

TABLE II *BIG DATA DIMENSIONS AND PRIVACY IMPLICATIONS*

Dimension	Description	Privacy Implication
Volume	Terabytes to exabytes	More unique QI combos even after generalisation; manual review infeasible
Velocity	Real-time streams (IoT)	Syntactic mechanisms are batch-only; streaming DP required – current pipeline limitation (§VII-B)
Variety	Structured, semi-structured, unstructured	Per-type anonymisation needed; cross-modal linkage creates new QI combinations
Veracity	Uncertainty and noise	Noisy records can create spurious unique QI combos, paradoxically raising re-ID risk
Value	Latent analytical worth	Utility loss from over-anonymisation creates adoption resistance; task-anchored metrics needed

B. Adversary Model

The adversary is characterised by (K, R, C, S): K ⊆ attributes known as background knowledge; R = an auxiliary dataset non-overlapping with D (different survey cycle, different synthetic

seed, or random split—limitations acknowledged in Section VII-A); C = computational budget; S ∈ {known, unknown} = knowledge of the anonymisation scheme. Background knowledge levels: A0 (K=∅), A25 (25% of QI attributes), A100 (all QI attributes), crossed with S. Privacy threats: identity disclosure, attribute disclosure, membership disclosure.

C. Regulatory Context

Under GDPR Recital 26, truly anonymous data lies outside the regulation's scope; pseudonymised data (Art. 4(5)) remains within scope. The Article 29 Working Party (Opinion 05/2014) [22] established that k-anonymised data and DP outputs at moderate ε are generally pseudonymised, not truly anonymous. Data subjects retain rights (access, erasure, portability) over all pseudonymised outputs. See the first-page note for the legal disclaimer.

IV. SURVEY OF ANONYMISATION TECHNIQUES

A. Syntactic Privacy Constraints

k-Anonymity [1]: every record must be indistinguishable from at least k−1 others on quasi-identifier (QI) attributes. It is a structural dataset constraint, not a probabilistic linkage-resistance guarantee.

l-Diversity [2] requires at least l well-represented sensitive attribute values per equivalence class. *t-Closeness* [3] additionally requires that the within-class distribution be within Earth Mover's Distance t of the global distribution, addressing skewness attacks (demonstrated in Appendix B).

B. Differential Privacy

A randomised mechanism M satisfies ε-differential privacy if, for any adjacent D, D′ differing in one record and any output set S:

$$\Pr[M(D) \in S] \leq \exp(\varepsilon) \times \Pr[M(D') \in S]. \quad (1)$$

Kifer and Machanavajjhala [15, Thm. 1] established that k-anonymity and DP are separately-typed properties with no joint formal guarantee under arbitrary auxiliary information—formalised here as Theorem 1. Cummings et al. [19] characterise utility-privacy bounds for statistical queries under user-level DP. Rényi DP [6] provides tighter composition accounting.

C. Attribute-Level DP Allocation

Haeberlen et al. [12], Blowfish [13], and DPSD [14] establish the theory of attribute-level DP allocation. Section V-G applies this to the pipeline's DP layer.

D. Data Masking, Tokenisation, and Synthetic Data

Masking and tokenisation provide operational data protection without formal statistical privacy properties. Synthetic data generation [21] combined with DP-SGD [20] provides differentially private record-level releases; membership inference on generative models remains an open concern for rare subpopulations.

V. TWO-LAYER SEPARATE-MECHANISM PIPELINE

A. Design Goals and Scope

The pipeline is a practical heuristic system for applying differentiated privacy mechanisms to heterogeneous dataset attributes. Two independently characterised layers: (1) syntactic constraints on D*; (2) DP on aggregate statistical query outputs. The name "two-layer separate-mechanism" emphasises that protections operate on different outputs with no joint guarantee. The configuration space has size |M|^m ($7^{20} \approx 8\times10^{16}$ for 20 attributes); exhaustive search is infeasible. The greedy k-first relaxation is a local heuristic justified empirically by cross-domain comparison in Section VI-F.

B. Problem Formulation

Let D be a dataset with schema (A_1,…,A_m). A configuration C=(c_1,…,c_m) where each c_i ∈ M={suppress, generalise(k), l-diverse(l), t-close(t), DP(ε,δ), tokenise, generalise-low(k_min=2)}. The pipeline seeks a feasible C satisfying:

$$R(D,C,K,R_{aux}) \leq R_{max} \quad \text{AND} \quad U(D,C,task) \geq U_{min} \quad (2)$$

where R is re-identification risk (Fellegi–Sunter precision against R_aux), U is task-specific utility (Table V), and both R_max and U_min are policy-specified.

C. Attribute Classification

All attributes are classified into four tiers (Table III). T4 (non-sensitive) attributes undergo a two-stage combinatorial MI check: (i) MI(T4, T1)<θ_MI, AND (ii) MI(T4∪T2_subset, T1)<θ_MI for all T2 subsets of size ≤3. Attributes failing (ii) are reclassified to T2. No unconditional pass-through is permitted.

TABLE III *ATTRIBUTE SENSITIVITY TAXONOMY WITH DEFAULT THRESHOLDS*

Tier	Type	Criterion	Threshold†	Mechanism(s)
T1	Direct ID	HIPAA 18-id. or ISO 29101; or ≥99% unique	—	Suppress / tokenise
T2	Quasi-ID	MI(A_i,T1)>θ_MI; or >1% k-viol. at k_QI	θ_MI=0.15†; k_QI=5†	gen.(k), l-div./t-close
T3	Sensitive	GDPR Art. 9; HIPAA PHI; policy	Reg. lists	l-div./t-close; DP agg.
T4	Non-Sens.	Not T1/T2/T3; passes 2-stage MI check	θ_MI=0.15†	gen.-low(k_min=2)

† Default values, not validated standards. θ_MI=0.15 nats is motivated by MI values (0.1–0.3 nats) for canonical QIs in prior work. k_QI=5 aligns with HIPAA Expert Determination practice. Sensitivity analysis over θ_MI∈{0.05,...,0.30} in Section VI-F.

D. Release Profiles

TABLE IV *RELEASE PROFILES AND POLICY THRESHOLDS*

Profile	Use Case	k	l/t	DP ε_total	R_max	U_min
A	Open data	k≥10	l≥5 or t≤0.15	ε≤1.0	5%	Table V
B	Trusted partner+DUA	k≥5	l≥3 or t≤0.20	ε≤2.0	10%	Table V
C	Internal (RBAC)	T1: tokenise	Access ctrl.	N/A	15%	N/A

E. Five-Stage Pipeline with Deterministic Relaxation

```
Algorithm 1 Two-Layer Separate-Mechanism Pipeline
  Input: D, profile P, adversary (K, R_aux, C, S),
         R_max, U_min, k_min=2, l_min=2,
         eps_max=5.0, max_relax=5
  Output: D* or ESCALATE

  1. CLASSIFY: Assign tier per Table III.
     T4: 2-stage comb. MI check;
     reclassify T4→T2 if check fails.
  2. QI-DISCOVER: For T2 subsets S, |S|≤K_max:
     v(S) = unique-QI-combo count.
     If v(S)>0.01n: flag S.
     Greedy fwd. sel. if |T2|>10.
  3. PROFILE: Load C_P from Table IV.
     k_curr=k_P; l_curr=l_P;
     eps_curr=eps_P; relax=0.
  4. Apply and Evaluate (structured loop):
     feasible ← false
     while relax ≤ max_relax and not feasible:
       APPLY: T1 suppress/tokenise;
              T2 gen.(k_curr);
              T3 l-div.(l_curr)/t-close(t_curr);
              T4 gen.-low(k_min);
              DP on agg. query outputs only.
       R ← FS_precision(D*, R_aux, Alg. 2)
       U ← domain_utility(D*)
       if R≤R_max AND U≥U_min:
         feasible ← true
       elif k_curr > k_min:
         k_curr -= 1; relax += 1
       elif l_curr > l_min:
         l_curr -= 1; relax += 1
       elif eps_curr < eps_max:
         eps_curr += 0.5; relax += 1
       else: break
     return D* if feasible else ESCALATE
```

If the pipeline returns ESCALATE, the data release must not proceed; a human expert must review the configuration.

Subgroup privacy warning: The pipeline does not provide fairness guarantees. Uniform DP noise and syntactic generalisation may disproportionately affect minority subgroups. Subgroup-specific privacy-utility audits are recommended before deployment.

F. Layer Characterisation and Interaction

Syntactic constraints are properties of D* and protect against equivalence-class-based matching by passive adversaries. DP is applied to aggregate statistical query outputs and provides a probabilistic bound (ε-DP) regardless of adversary background knowledge. The acknowledged interaction: syntactic generalisation reduces attribute precision in D*, changing the global sensitivity Δf of subsequent DP queries. Δf is computed over D* (post-generalisation), which is conservative—generalisation reduces value ranges, so $\Delta f_D^* \leq \Delta f_D$. No formal composition between the layers is claimed.

G. DP Budget Allocation with Task Elicitation

Budget allocation follows He et al. [14] (DPSD): $\varepsilon_3=\varepsilon_total\times w3$, $\varepsilon_4=\varepsilon_total\times(1-w3)$, where $w3=v3/(v3+v4)$. When the analytical task is predetermined, v3 and v4 are counted from the declared query specification. When not predetermined, a structured stakeholder questionnaire is used (template in repository). Default w3=0.5 if questionnaire unavailable; sensitivity analysis over $w3\in\{0.3,0.5,0.7\}$ is reported.

H. Theorem 1: Non-Composability

Theorem 1 *(Non-composability, adapted from Kifer and Machanavajjhala [15, Thm. 1]).*

Let D be a dataset. Suppose mechanism M first transforms D into D* satisfying k-anonymity and then answers a query q with an ε-DP mechanism. For any given $k \geq 1$ and $\varepsilon \geq 0$, there exists an adversary with auxiliary information that can violate a (ε',δ)-DP guarantee for the combined output (D*, q(D*)) for any finite ε', $\delta<1$, provided the adversary has knowledge of the anonymisation scheme.

Proof sketch. Kifer and Machanavajjhala show that k-anonymity does not bound information leakage about an individual's presence when the adversary has arbitrary auxiliary information: the equivalence class structure together with external knowledge enables membership inference with probability 1, irrespective of k. Since ε-DP requires that membership inference be bounded by $\exp(\varepsilon)$, the two properties are incompatible as a joint guarantee. ■

Corollary 1 *(Separate-mechanism operation).* The pipeline avoids this impossibility by never claiming a joint guarantee. The syntactic layer protects against equivalence-class-based matching under bounded auxiliary information. The DP layer provides a probabilistic bound for query outputs only, under arbitrary auxiliary information.

VI. METHODOLOGICAL VALIDATION

A. Datasets, Baselines, and Framing

Results are framed as methodological validation on public datasets, not as real-world privacy risk estimates. Three datasets: (1) NHANES-based synthetic demographic-health survey data (250,000 records; NHANES 2017–2018 public files; no clinical EHR); (2) Sparkov financial transactions (12M

records, MIT License); (3) MDC mobility trajectories (6.8M records, CC-BY 4.0)—auxiliary data is a random 80/20 participant split (not independent; worst-case estimates).

Six baselines: B1 (no anonymisation); B2 (uniform k=5); B3 (attribute-matched k=5); B4 (record-level DP at equivalent ε_total); B5 (DP on queries only, no syntactic protection); B6 (record-level DP then k-anonymise). Statistical significance: 100 bootstrap resamples, Bonferroni correction ($\alpha_corr = 0.05/12 \approx 0.0042$).

TABLE V *CASE STUDY CONFIGURATIONS. OD-R = PEARSON CORRELATION OF OD FLOW MATRIX. * MDC USES 80/20 SPLIT (WORST-CASE).*

Case Study	Records	Attrs.	U_min (A/B)	R_max	w3
Dem.-Health	250K	42	RMSE≤0.030/0.040	5/10%	0.60
Financial	12M	12	AUC≥0.900/0.880	5/10%	0.70
Mobility*	6.8M	8	OD-r≥0.920/0.900	5/10%	0.50

B. Demographic-Health Survey

Utility metrics: smoking-status prevalence RMSE and BMI-category distribution RMSE. Profile B (R_max=10%, RMSE ≤0.040). Non-overlapping auxiliary: NHANES 2015–2016 cycle. Fellegi–Sunter parameters estimated via EM [27]. Zero relaxation steps required.

TABLE VI *DEMOGRAPHIC-HEALTH RESULTS (PROFILE B, K=5, L=3, E=1.0). FS = FELLEGI–SUNTER PRECISION.*

Method	FS Prec.(%)	95% CI	RMSE	r	p vs. B3
B1 No anon.	14.3%	[12.8,15.9]	0.000	1.000	—
B2 Unif. k=5	4.1%	[3.3,5.0]	0.031	0.961	—
B3 Attr-match	3.8%	[3.0,4.7]	0.022	0.971	—
B4 Rec. DP	6.9%	[5.9,7.9]	0.041	0.943	—
B5 DP queries only	13.1%	[11.8,14.5]	0.015	0.984	—
B6 DP then k=5	5.2%	[4.3,6.2]	0.035	0.951	—
Proposed	1.8%	[1.3,2.5]	0.019	0.914	<0.0001

The proposed method achieves significantly lower FS precision than B3 (1.8% vs. 3.8%, p<0.0001, Bonferroni-corrected). B5 (DP on queries only) shows that omitting syntactic record-level protection produces unacceptably high FS precision (13.1%), confirming the necessity of the syntactic layer even when DP is applied to query outputs.

C. Financial Transactions

Utility metric: LightGBM fraud detection AUC (5-fold CV). Profile A. Non-overlapping auxiliary: Sparkov 2023 (different simulation period). Selective DP concentrates noise on the most temporally distinctive 20% of transactions.

TABLE VII *FINANCIAL TRANSACTION RESULTS (PROFILE A). FS = FELLEGI–SUNTER PRECISION.*

Method	FS Prec.(%)	95% CI	AUC	Deg.	p vs. B3
B1 No anon.	90.0%	[88.1,91.8]	0.943	—	—
B2 Unif. k=5	31.4%	[29.7,33.2]	0.901	−0.042	—
B3 Attr-match	9.1%	[8.0,10.3]	0.904	−0.039	—
B4 Rec. DP	8.7%	[7.5,9.9]	0.887	−0.056	—
B5 DP queries only	78.3%	[76.1,80.5]	0.938	−0.005	—
B6 DP then k=5	10.2%	[9.0,11.5]	0.893	−0.050	—
Proposed	4.8%	[3.9,5.8]	0.911	−0.032	<0.0001

D. Mobility Trajectories (Worst-Case Auxiliary)

Limitation: R_aux is a random 80/20 MDC participant split—not independent data. Results represent a worst-case adversary estimate and are not directly comparable to health and finance.

Utility metric: OD-r. Profile A. One relaxation step required: initial k=10 produced OD-r=0.906<0.920; k relaxed to 9, yielding OD-r=0.940. Note: resulting FS precision (11.0%) exceeds Profile A R_max=5%; Profile B (R_max=10%) would be adopted in practice for mobility data.

TABLE VIII *MOBILITY RESULTS (WORST-CASE 80/20 AUXILIARY SPLIT). OD-R = PEARSON CORRELATION.*

Method	FS Prec.(%)	95% CI	OD-r	p vs. B3
B1 No anon.	73.0%	[71.2,74.8]	1.000	—
B2 Grid sup.	28.4%	[26.8,30.0]	0.891	—
B3 Attr-match k=5	19.2%	[17.8,20.7]	0.921	—
B4 Spatial DP	14.8%	[13.5,16.2]	0.927	—
B5 DP queries only	64.1%	[61.9,66.3]	0.972	—
B6 DP then k=5	17.3%	[15.9,18.8]	0.918	—
Proposed (k=9)	11.0%	[9.8,12.3]	0.940	0.0002

E. t-Closeness Validation

A synthetic skew-augmented demonstration of t-closeness is provided in Appendix B. t-Closeness with t=0.15 reduces skewness attack success from 18.7% (l-diversity only) to 6.3% on the synthetic benchmark. These results are from a synthetic case; real-world validation is required before deployment.

F. Relaxation Order: Cross-Domain Comparison

The k-first order is compared against l-first and ε-first across all three domains using an infeasible starting configuration. Due to the MDC auxiliary data limitation, the cross-domain conclusion rests primarily on health and finance.

TABLE IX *RELAXATION ORDER COMPARISON. * MOBILITY USES WORST-CASE 80/20 AUXILIARY; NOT DIRECTLY COMPARABLE.*

Domain	Order	Steps	FS Prec.(%)	95% CI
Health	k-first	5	3.6%	[2.8,4.5]
Health	l-first	3	5.1%	[4.2,6.1]
Health	ε-first	2	4.8%	[3.9,5.8]
Finance	k-first	3	3.2%	[2.5,4.0]
Finance	l-first	2	5.4%	[4.5,6.4]
Finance	ε-first	2	4.9%	[4.0,5.9]
Mobility*	k-first	1	11.0%	[9.8,12.3]
Mobility*	l-first	1	11.0%	[9.8,12.3]
Mobility*	ε-first	1	11.0%	[9.8,12.3]

Across health and finance (domains with non-overlapping auxiliary data), k-first achieves lower FS precision at the feasible point than l-first ($p<0.01$, Bonferroni-corrected). The k-first preference is an empirical observation, not a theoretical guarantee.

G. Sensitivity Analysis

Sensitivity over $k\in\{2,5,10,15,20\}$, $l\in\{2,3,5\}$, $t\in\{0.10,0.15,0.20\}$, $\varepsilon\in\{0.5,1.0,2.0,5.0\}$, $\theta_MI\in\{0.05,...,0.30\}$, $w3\in\{0.3,0.5,0.7\}$. θ_MI sensitivity: varying from 0.05 to 0.30 changes T2 attribute count by ±2–4 per domain with re-identification risk effects ±0.3–0.6%—confirming θ_MI is not a critical parameter. w3 sensitivity: fraud AUC varies ±0.008; FS precision ±0.5%.

H. Runtime Measurements

TABLE X *PIPELINE WALL-CLOCK RUNTIMES. SINGLE-THREAD PYTHON, INTEL XEON GOLD 6248, 256 GB RAM.*

Stage	Health (250K)	Finance (12M)	Mobility (6.8M)
1. Classify + MI	2.1 s	18.4 s	11.2 s
2. QI-Discover	4.8 s	112 s (greedy)	6.7 s
3. Profile load	<0.1 s	<0.1 s	<0.1 s
4. Apply	1.3 s	8.4 s	5.1 s
5. FS Evaluation	18 s	314 s	89 s
Total (0 relax)	26.2 s	453 s (7.6 min)	112 s (1.9 min)

FS evaluation dominates runtime. For the 12M financial dataset, R_aux is capped at 50,000 records with locality-sensitive hashing pre-filtering, reducing pair evaluations from 6×10^11 to 6×10^8.

VII. DISCUSSION, THREATS, AND ETHICAL CONSIDERATIONS

A. Adversarial Evaluation

Re-identification risk is measured as Fellegi–Sunter precision with EM-based parameter estimation [27] against non-overlapping auxiliary datasets: NHANES 2015–2016 for health, Sparkov 2023 for finance. For mobility, a random 80/20 MDC participant split is used (limitation: worst-case estimates not comparable to other domains).

TABLE XI *ADVERSARIAL EVALUATION (PROPOSED METHOD). FS = FELLEGI–SUNTER PRECISION. * MOBILITY: WORST-CASE 80/20 SPLIT.*

Domain	Adversary	Scheme	FS Prec.(%)	95% CI	FMR(%)
Health	A0	Known	1.9%	[1.2,2.8]	0.93
Health	A25	Known	6.4%	[5.1,7.8]	0.95
Health	A100	Known	14.2%	[12.4,16.1]	0.91
Health	A100	Unknown	9.8%	[8.3,11.4]	0.96
Finance	A100	Known	17.3%	[15.4,19.3]	0.90
Mobility*	A100	Known	28.4%	[26.0,30.9]	0.94

Unknown-scheme adversaries achieve lower FS precision than known-scheme, confirming known-scheme provides the conservative bound. Residual risk at A100 for mobility (28.4%) reflects both the inherent re-identifiability of trajectory data [8], [29] and the worst-case auxiliary assumption.

B. Pipeline Limitations

[8] Greedy relaxation has no global optimality guarantee.

[9] Pipeline is batch-only; streaming data (high-velocity dimension) is not supported.

[10] T4 combinatorial MI check covers subsets of size ≤3; higher-order interactions are not quantified.

[11] EM-based Fellegi–Sunter assumes conditional attribute independence.

[12] w3 elicitation depends on stakeholder questionnaire quality.

[13] Mobility auxiliary data (80/20 split) is not independent; results are worst-case estimates.

C. Regulatory Considerations

Pipeline outputs are likely pseudonymised under GDPR for Profile B/C releases ($k=5$, $\varepsilon\leq 2.0$) and likely for most Profile A releases. Whether a specific release qualifies as truly anonymous under GDPR Recital 26 is a context-specific legal determination that this paper cannot make. Organizations should seek independent legal counsel. Data subjects likely retain GDPR rights over all pseudonymised outputs.

D. Ethical Considerations

All datasets are public (NHANES, Sparkov, MDC); no real patient or financial data was used. Uniform DP noise disproportionately distorts query outputs for minority subgroups (disparate privacy [19]); subgroup-aware DP budget allocation is a direction for future work. All configuration decisions must be traceable to documented risk assessments.

VIII. CONCLUSION

This paper has presented a two-layer separate-mechanism anonymization pipeline and evaluated its resistance to a

Fellegi–Sunter record-linkage adversary: syntactic constraints on D* and differential privacy on aggregate query outputs, applied with no joint formal guarantee. Theorem 1 formalises why such a guarantee is impossible under arbitrary auxiliary information, and the adversarial evaluation in Section VII-A supplies the empirical security evidence that the theorem shows cannot be obtained analytically. The pipeline uses a greedy k-first relaxation order supported by cross-domain empirical comparison. A public code repository provides full implementation, including the record-linkage adversary used for evaluation.

Residual limitations—greedy local search, batch-only operation, limited T4 combinatorial coverage, worst-case mobility auxiliary, and disparate privacy effects on minority subgroups—define primary future work: Pareto-optimal multi-objective search, streaming DP adaptation, higher-order MI discovery, and subgroup-aware budget allocation.

REFERENCES

1. [1] L. Sweeney, "k-anonymity: A model for protecting privacy," Int. J. Uncertainty Fuzziness Knowl.-Based Syst., vol. 10, no. 5, pp. 557–570, 2002.
2. [2] A. Machanavajjhala, D. Kifer, J. Gehrke, and M. Venkitasubramaniam, "l-diversity: Privacy beyond k-anonymity," ACM Trans. Knowl. Discovery Data, vol. 1, no. 1, p. 3, 2007.
3. [3] N. Li, T. Li, and S. Venkatasubramanian, "t-closeness: Privacy beyond k-anonymity and l-diversity," in Proc. IEEE ICDE, 2007, pp. 106–115.
4. [4] C. Dwork, "Differential privacy," in Proc. ICALP, LNCS vol. 4052, 2006, pp. 1–12.
5. [5] C. Dwork, F. McSherry, K. Nissim, and A. Smith, "Calibrating noise to sensitivity in private data analysis," in Proc. TCC, LNCS vol. 3876, 2006, pp. 265–284.
6. [6] I. Mironov, "Rényi differential privacy," in Proc. IEEE CSF, 2017, pp. 263–275.
7. [7] A. Narayanan and V. Shmatikov, "Robust de-anonymization of large sparse datasets," in Proc. IEEE S&P, 2008, pp. 111–125.
8. [8] Y.-A. de Montjoye, C. A. Hidalgo, M. Verleysen, and V. D. Blondel, "Unique in the crowd: The privacy bounds of human mobility," Sci. Rep., vol. 3, p. 1376, 2013.
9. [9] X. Xiao and Y. Tao, "Anatomy: Simple and effective privacy preservation," in Proc. VLDB, 2006, pp. 139–150.
10. [10] M. E. Nergiz, M. Atzori, and C. Clifton, "Hiding the presence of individuals from shared databases," in Proc. ACM SIGMOD, 2009, pp. 665–678.
11. [11] J. Domingo-Ferrer, D. Sánchez, and J. Soria-Comas, "Database anonymization," Synth. Lectures Inf. Security Privacy Trust, vol. 8, no. 1, pp. 1–136, 2016.
12. [12] A. Haeberlen, B. C. Pierce, and A. Narayan, "Differential privacy under fire," in Proc. USENIX Security, 2011.
13. [13] H. Ebadi, D. Sands, and G. Schneider, "Differential privacy: Now it's getting personal," in Proc. ACM POPL, 2015, pp. 69–81. [Blowfish privacy]
14. [14] X. He, A. Machanavajjhala, and B. Ding, "Blowfish privacy: Tuning privacy-utility trade-offs using policies," in Proc. ACM SIGMOD, 2014, pp. 1447–1458. [DPSD]
15. [15] D. Kifer and A. Machanavajjhala, "No free lunch in data privacy," in Proc. ACM SIGMOD, 2011, pp. 193–204.
16. [16] A.-S. Charest, "How can we analyze differentially private synthetic datasets?" J. Privacy Confidentiality, vol. 2, no. 2, 2010.
17. [17] I. P. Fellegi and A. B. Sunter, "A theory for record linkage," J. Amer. Statist. Assoc., vol. 64, no. 328, pp. 1183–1210, 1969.
18. [18] M. Abadi et al., "Deep learning with differential privacy," in Proc. ACM CCS, 2016, pp. 308–318.
19. [19] R. Cummings, S. Krehbiel, N. Melikechi, and A. Roth, "Mean estimation with user-level privacy under data heterogeneity," in Proc. NeurIPS, 2021.
20. [20] M. Jagielski, J. Ullman, and A. Oprea, "Auditing differentially private machine learning," in Proc. NeurIPS, 2020.
21. [21] L. Xu, M. Skoularidou, A. Cuesta-Infante, and K. Veeramachaneni, "Modeling tabular data using conditional GAN," in Proc. NeurIPS, 2019, pp. 7335–7345.
22. [22] Article 29 Data Protection Working Party, "Opinion 05/2014 on anonymisation techniques," WP216, European Commission, 2014.
23. [23] European Parliament, "General data protection regulation (EU) 2016/679," Official J. Eur. Union, 2016.
24. [24] U.S. Dep. Health Human Services, "Guidance regarding methods for de-identification of protected health information under HIPAA," Office for Civil Rights, 2012.
25. [25] P. Samarati and L. Sweeney, "Protecting privacy when disclosing information: k-anonymity," Tech. Rep., SRI International, 1998.
26. [26] J. Wieringa et al., "Data analytics in a privacy-concerned world," J. Business Res., vol. 122, pp. 915–925, 2021.
27. [27] W. E. Winkler, "Using the EM algorithm for weight computation in the Fellegi-Sunter model of record linkage," in Proc. Sect. Survey Res. Methods, Amer. Statist. Assoc., 1988, pp. 667–671.
28. [28] A. Gkoulalas-Divanis, D. Vatsalan, D. Karapiperis, and M. Kantarcioglu, "Modern privacy-preserving record linkage techniques: An overview," IEEE Trans. Inf. Forensics Security, vol. 16, pp. 4966–4987, 2021.
29. [29] T. Murakami, A. Kanemura, and H. Hino, "Group sparsity tensor factorization for re-identification of open mobility traces," IEEE Trans. Inf. Forensics Security, vol. 12, no. 3, pp. 689–704, 2017.
30. [30] W. Yin, L. Yuan, Y. Ren, W. Meng, D. Wang, and Q. Wang, "Differential cryptanalysis of Bloom filters for privacy-preserving record linkage," IEEE Trans. Inf. Forensics Security, vol. 19, pp. 7935–7948, 2024.

AI-Assisted Editing Disclosure: AI-assisted tools were used for copy editing purposes to improve readability and grammar. All intellectual content, analysis, framing, and conclusions are the original work of the authors. The authors take full responsibility for the integrity of the manuscript.

Mohammed Sayim Khalil received the Ph.D. degree from Universiti Teknologi Malaysia (UTM). He is currently an Associate Professor in the Department of Computer Engineering, Halic University, Istanbul, Turkey. His research interests include big data analytics, data privacy, selective anonymization, and the ethics of artificial intelligence. He is a senior member of IEEE. (https://orcid.org/0000-0002-1539-5629)