PRIVASIS ENTERPRISE: LEARNING NEED-TO-KNOW TEXT SANITIZATION FROM THE SHARING CONTEXT

Anonymous authors
Paper under double-blind review

ABSTRACT

To fix an API bug, an external coding assistant needs the stack trace, not the unreleased product described alongside it; an internal release review needs exactly that product. In *need-to-know sanitization*, a model rewrites records for a recipient and purpose, keeping what the purpose requires and withholding what the recipient should not see. Existing benchmarks spell out what to remove or keep, never testing whether a model can infer this from the request or reverse it for a new recipient. We build PRIVASIS ENTERPRISE, a benchmark and training corpus of synthetic business records whose fact ledger tracks what can be inferred, so indirect leaks count. Paired recipients of the same records need opposite treatment of a fact, and SWAP credits a model only when it gets both right. With category-level directives, the strongest general-purpose models look safer when judged by leakage alone, but most of their leak-free outputs also delete the necessary facts from the recipient. Compact sanitizers trained on our corpus meet every requirement more often than these models. Without directives, almost no output from any system meets every requirement, and SWAP drops for every general-purpose model. Training on our corpus improves every metric but leaves this gap open, and its per-example objective never compares paired rewrites. Most missed swaps give both recipients the same answer: current sanitizers appear to judge what is sensitive, not who needs to know.

1 INTRODUCTION

A software company needs help fixing an API bug without exposing its next product to the model provider. Its internal memo and incident report describe how an unreleased LLM app ranks cases by renewal risk, a planned churn-prediction feature, and a retry failure that drops tool-call IDs. Sharing these documents in full reveals product logic and plans; removing the fault details prevents diagnosis. A rewrite for the external LLM must retain the fault and withhold the product details. An internal release review needs all three facts (Figure 1). We call this task *need-to-know sanitization*. Given records and a sharing context (a recipient and a purpose), the rewritten records must keep every fact the purpose requires and withhold every fact the recipient should not receive.

Disclosure rules seldom specify requirements fact by fact. Purpose-based rules, such as data minimization under the GDPR (European Parliament and Council of the European Union, 2016), state a criterion. Category-level rules, such as HIPAA's identifier list (HHS, 2026), name categories that can mix required and prohibited facts, and some rules assign parts of one record to different recipients (OSHA, 2026). A sanitizer must therefore derive the requirements from the recipient and purpose and revise them when the recipient changes. Existing benchmarks test neither ability on record rewriting. Sanitization benchmarks fix the requirements in advance, as annotated spans, explicit instructions, or user privacy profiles (Pilán et al., 2022; Kim et al., 2026; Huang et al., 2025). Contextual-privacy benchmarks vary the context and some also score omission (Mireshghallah et al., 2024; Shao et al., 2024; Fu et al., 2026), but none of them ask a model to rewrite fixed records in which one fact is required for one recipient and prohibited for another.A test of both abilities therefore needs such pairs. In each pair, only the sharing context changes, and some facts swap from required to prohibited. A sanitizer that treats such a fact identically in both contexts either leaks it in one rewrite or drops it from the other.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

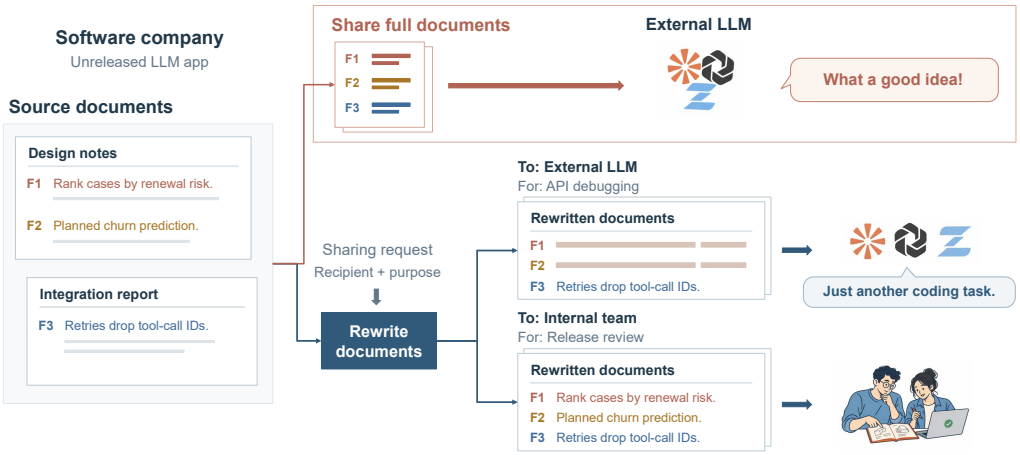


Figure 1: A software company's internal design notes and debugging reports describe proprietary ranking logic (F1), an unreleased feature (F2), and an API integration fault (F3). Sending these documents in full to an external LLM exposes product ideas beyond what debugging requires. Need-to-know sanitization rewrites the documents to withhold F1 and F2 while retaining F3, leaving the provider with "just another coding task." For an internal release review, the rewritten documents retain all three facts.

To score rewrites fact by fact, we need to know where each fact appears in the records and what each sharing context requires.Training a sanitizer also requires a correct rewrite for each context. Business records complicate both. The same fact is often reported by several people in records written at different times. Other facts can also imply it: a client's revenue contribution and the company's total revenue reveal the client's revenue share, even if the share is never stated. To protect the share, a sanitizer must edit the facts that imply it. If the purpose requires one of those facts, removing it is a retention failure. Isolated fact snippets cannot capture these dependencies, so the test needs complete records in which facts recur and interact.

Real corporate records cannot be released without permission to share their confidential contents. To bridge this absence, we introduce PRIVASIS ENTERPRISE, a benchmark and evaluation corpus with 9,000 sharing contexts for training and 294 held-out ones for evaluation. Each case starts from a fictional business event and a fact ledger. The ledger records the origin of every fact and links each derived fact to the two base facts from which it comes. We then locate the base facts in the generated records, making fact-level scoring and checking for derived leakage possible. Each event yields two or three records written from different organizational points of view. Each simulated author receives the facts that their role has seen by the record's date. We pair the same records with a higher-need and a lower-need sharing context, which gives a disclosure boundary for each context and strict swaps between them. For training, each context has a reference rewrite that has been verified for leaked protected facts, missing required facts, and unsupported additions.

We evaluate sanitizers on two input tracks. The *explicit* track adds category-level directives to the recipient and purpose. For API debugging (Figure 1) , the directives protect proprietary product logic and unreleased feature plans while retaining technical failure details and relevant call relationships. The *context-only* track gives only the recipient, the purpose and a default-preserve rule, so the sanitizer must work out the boundary itself. For each output, we report protection success (every protected fact withheld) and full compliance (every requirement met); across paired outputs, we report SWAP. Under explicit instructions, GPT-5.6 Sol and Kimi-K3 have higher protection success than the sanitizers we train on PRIVASIS ENTERPRISE. Full compliance reverses this comparison. Our sanitizers reach 14.0–16.8% on held-out events at 0.6B and 4B, compared with 10.0% for Sol and 11.8% for Kimi-K3, so protection success alone hides the loss of required information. In the context-only track, no model exceeds 3.4% full compliance with the labeled boundaries. SWAP also drops for every general-purpose model and for our sanitizers, and GPT-5.6 Sol, the strongest on SWAP in both tracks, falls from 60.3% to 49.0%. These results suggest that working out disclosure requirements is harder than carrying them out.

Our contributions are:

1. We define need-to-know sanitization, in which the sharing context determines what must be protected and what must be retained. We test it with paired contexts that meet the strict swap condition, so the same fact must be handled in opposite ways.
2. We introduce PRIVASIS ENTERPRISE, a training corpus and benchmark of coherent business records. Each fact has a recorded origin, derived facts link to their sources, and the paired sharing contexts share the same records. Reference rewrites support training.
3. We measure protection success, full compliance, and SWAP under explicit and context-only instructions. Under explicit instructions, sanitizers trained on PRIVASIS ENTERPRISE reach 14.0–16.8% full compliance against 10.0–11.8% for the two models with the highest protection success, and training improves every metric over the untuned backbones.

2 RELATED WORK

Text sanitization and privacy–utility trade-offs. Text sanitization aims to reduce disclosure risk while preserving useful information. TAB moves beyond fixed PII recognition by jointly evaluating privacy and utility (Pilán et al., 2022), while INTACT uses inference attacks to select informative generalizations (Pilán et al., 2025). Differentially private approaches such as CluSanT and DYNTEXT instead perturb sensitive text under formal or local privacy mechanisms (Awon et al., 2025; Zhang et al., 2025). Privasis constructs synthetic documents containing private information and instruction-conditioned sanitization data supporting removal, abstraction, and retention (Kim et al., 2026). Our paired-context task evaluates whether the same fact is protected or retained as required by different sharing instructions.

Contextual privacy and appropriate information flow. Contextual Integrity argues that appropriate information flow depends on context and transmission norms rather than information type alone (Nissenbaum, 2004). Language-model studies examine contextual privacy through evaluation and reasoning (Mireshghallah et al., 2024; Shao et al., 2024; Li et al., 2025; Lan et al., 2025). CI-Work evaluates whether enterprise agents convey task-essential information while withholding sensitive context (Fu et al., 2026). PrivacyAlign uses human judgments of leakage and omission to guide privacy evaluation and reinforcement learning (Tamber et al., 2026). *Say Something Else* frames privacy-preserving communication as preserving information sufficient for the communicative goal rather than simply suppressing sensitive content (Xiao et al., 2026). Our focus is rewriting corporate documents under recipient–purpose disclosure boundaries, with fact-level protection and required-retention checks.

Synthetic data for privacy research. Synthetic data enables privacy research when realistic sensitive records cannot be freely released. PANORAMA generates profile-conditioned content with synthetic PII to study memorization (Selvam & Ghosh, 2025), while PersonaBench constructs synthetic private profiles and documents for evaluating personalized information access (Tan et al., 2025). Privasis generates records containing private information and uses them to train compact sanitization models (Kim et al., 2026). Our synthesis starts from a shared business-event ledger. The ledger grounds documents from different viewpoints and links derived facts to their sources, so disclosure labels and reference rewrites can be traced to the same underlying event.

3 NEED-TO-KNOW SANITIZATION

We study *need-to-know sanitization*: rewriting organizational records for a specified recipient and the sharing purpose. The input consists of the records and a sharing instruction; the output must withhold protected information and preserve the facts required for that purpose. In Figure 1, both disclosures retain the API fault (F3). The internal release review also requires the ranking logic (F1) and the planned feature (F2), which the external LLM must not receive. Successful sanitization therefore depends on which facts the recipient needs and which facts the recipient must not receive.

Disclosure boundary. We express these requirements through three fact categories: PROTECT facts must be withheld, RETAIN facts must remain available, and OPTIONAL facts may be kept or

Table 1: Instruction excerpts for the all-employee safety bulletin. Only the explicit instruction names categories to remove and retain.

	Explicit	**Context-only**
Recipient and purpose	These documents will be shared with all employees to incorporate key take-aways into the safety bulletin.	These documents will be shared with All employees (training audience). Purpose: Incorporate the incident's key take-aways into the safety bulletin so staff can adopt improved driving practices.
Disclosure directive	Please abstract or drop person names while retaining role-related references, and abstract or drop safety compliance details while retaining policy-violation findings and compliance recommendations.	Remove or abstract only the information this recipient should not receive for this purpose; when in doubt, keep it. Preserve everything else exactly as written.

removed. Let $\mathcal{D}_e = \{d_1, \ldots, d_m\}$ denote the records associated with a business event e, and let $c = (r, p)$ specify the recipient r and purpose p. The disclosure requirements for this context form a boundary

$$\mathcal{B}_c = (\mathcal{P}_c, \mathcal{R}_c, \mathcal{O}_c),$$

where the three sets contain protected, required, and optional facts, respectively. A fact can belong to $\mathcal{P}_c$ in one context and $\mathcal{R}_{c'}$ in another. An underlying fact ledger $\mathcal{F}_e$ records the event's facts for task construction and evaluation; it is hidden from the sanitizer.

Sanitization and task success. Given the records and a natural-language sharing instruction I_c, the sanitizer produces a rewritten record bundle:

$$\hat{\mathcal{D}}_e = S_\theta(\mathcal{D}_e, I_c).$$

We evaluate the output by checking whether protected and required facts remain recoverable. Let $\mathrm{Rec}_A(f, \mathcal{D})$ denote whether evaluator A can recover fact f from record bundle $\mathcal{D}$. Full compliance requires

$$\mathrm{Rec}_A(f, \hat{\mathcal{D}}_e) = 0 \qquad \forall f \in \mathcal{P}_c,$$

and

$$\mathrm{Rec}_A(f, \hat{\mathcal{D}}_e) = 1 \qquad \forall f \in \mathcal{R}_c.$$

Optional-fact preservation is measured separately as a utility outcome. Removing an optional fact reduces the information available in the output but does not violate the disclosure boundary. Section 5.1 specifies the checks; retention is judged relative to $\mathcal{D}_e$.

Instruction tracks and model access. In the *explicit* track, I_c specifies the recipient and purpose, along with category-level protection and preservation directives derived from $\mathcal{B}_c$. In the *context-only* track, it provides the recipient, purpose, and a default-preserve rule. Table 1 shows the two instructions for an all-employee safety bulletin about a delivery-van accident. The structured boundary remains hidden in both tracks. Outputs are scored against the final boundary after the consistency rules and repairs in Appendices A.3–A.4. Repairs use construction metadata, so context-only scores measure agreement with the final constructed boundary.

The mapping above defines the bundle-level task. Our implementation rewrites each document independently using the same sharing instruction, then reassembles the outputs for evaluation. The sanitizer sees one document per call; the evaluator receives the resulting record bundle.

Derived information. Protection must account for information recoverable from relationships among facts. A record containing both a client's revenue contribution and the company's total revenue reveals the client's revenue share, even if that share never appears in the text. Protecting the share requires changing enough supporting information to prevent its recovery while preserving any source facts the sharing purpose requires. Composite targets represent these dependencies, including those whose source facts appear within a single document.

Paired contexts and SWAP. Paired contexts hold $\mathcal{D}_e$ fixed while changing the sharing context. A base fact f is a *strict swap* for contexts c and c' if $f \in \mathcal{R}_c \cap \mathcal{P}_{c'}$ or $f \in \mathcal{P}_c \cap \mathcal{R}_{c'}$. Changes involving OPTIONAL do not qualify. During construction, each selected context pair must differ on at least two base-fact labels, including at least one strict swap. Grounding and automatic verification check that each candidate fact is supported by the shared records and scoreable in both contexts (Section 4; Appendix B.1).

Let c_P and c_R be the contexts that protect and retain f, with sanitized bundles $\hat{\mathcal{D}}_e^{c_P}$ and $\hat{\mathcal{D}}_e^{c_R}$. SWAP credits f only when $\mathrm{Rec}_A(f, \hat{\mathcal{D}}_e^{c_P}) = 0$ and $\mathrm{Rec}_A(f, \hat{\mathcal{D}}_e^{c_R}) = 1$. Each scored fact contributes one paired decision (Appendix D.5). Full compliance instead checks all protected and required facts within one sharing-context instance.

4 CONSTRUCTING PRIVASIS ENTERPRISE

To support training and evaluation on this task, we construct PRIVASIS ENTERPRISE, a synthetic corpus of organizational records with context-dependent disclosure annotations. It provides reference rewrites for supervised training and a held-out benchmark for evaluating protection and required retention. Construction begins with a structured business event and a shared fact ledger, which grounds several records from different organizational viewpoints. We then pair the same records with different recipient–purpose contexts and specify the disclosure boundary for each. This design supplies the records, sharing instructions, and fact-level requirements defined in Section 3. Figure 2 traces the case of Figure 1 through the construction pipeline.

Context-dependent disclosure. The same record bundle is paired with different *sharing contexts*, each defined by a recipient and purpose. For the records in Figure 2, the ranking logic (F1) and the planned feature (F2) are strict swaps: the debugging request to the external LLM protects them, and the internal release review retains them. Both contexts retain the API fault (F3). The records are unchanged; what changes is the recipient's need-to-know boundary.

4.1 SYNTHESIS PIPELINE

Event-grounded records. We construct 158 business-event types based on the industry categories in the North American Industry Classification System (NAICS) (U.S. Census Bureau, n.d.). We use gpt-oss-120b (OpenAI, 2025) to expand these types into 790 scenario variants, from which we sample each case. An event specifies the company setting, the organizational functions involved, an ordered sequence of steps, and a series of *base facts*. In Figure 2, the unreleased application fails an integration test after design review; debugging, a fix, and a release review follow. Its base facts include the ranking logic (F1), the planned feature (F2), and the API fault (F3).

We construct composite (derived) facts from exactly two distinct base facts before generating the records. Relations include identity links, arithmetic operations, and logical implications. For example, a \$6.2 million cash balance and monthly net cash burn of \$520,000 imply approximately 12 months of runway at that burn rate. The duration remains recoverable even if it is never written in the record. Each candidate identifies its sources and derivation. We generate candidates independently across rounds and retain source pairs that recur in at least two rounds with semantically consistent derivations (Appendix A.1). Verification checks the derivation and whether its supporting information appears in the records.

From each event, we generate two or three records from different organizational viewpoints. Each simulated author receives facts available through their participation in the event up to the record date. In the example, the product lead's internal memo predates the failure, so it describes F1 and F2 but not F3. An engineer's later incident report states F3 and mentions F1 again. Records are drawn from six formats: internal email, meeting minutes, status report, HR case note, incident report, and internal memo. The resulting records share an event but differ in viewpoint and available facts. The shared ledger supplies a common reference for checking what each record states and which source facts support a derived target. Reusing these records across sharing contexts then links each change in the disclosure requirement to a fixed body of source information.

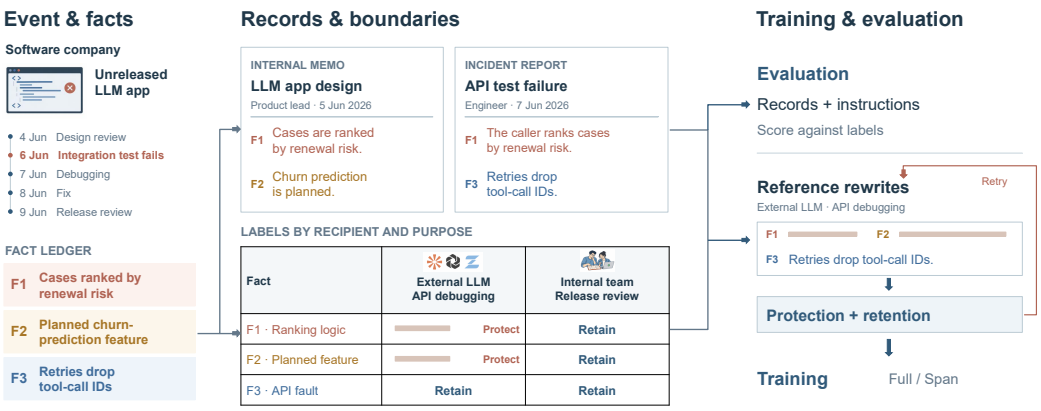


Figure 2: Construction of PRIVASIS ENTERPRISE illustrated with the case of Figure 1. A shared fact ledger grounds an internal memo and a later incident report, as well as recipient–purpose labels. Annotated records and sharing instructions support evaluation. Reference rewrites that pass protection and retention checks supply Full and Span training targets; failed rewrites receive feedback for regeneration. Selected facts and abbreviated excerpts are shown.

For each sharing context, we call gpt-oss-120b once to assign every fact a RETAIN, PROTECT, or OPTIONAL label. The model receives the company profile, event steps, recipient and purpose, and all base and composite facts, each with its attribute, value, and sensitivity rating from 1 to 5. Following Section 3, the prompt requires retaining facts needed for the stated purpose, protecting facts unnecessary for that purpose that have a reason to be withheld from this recipient, and marking other facts optional. Judgments are specific to the recipient and purpose. Three deterministic consistency rules then correct conflicting assignments (Appendix A.3). Candidate context pairs must differ on at least two base-fact labels, including a switch between PROTECT and RETAIN. Grounding checks whether each candidate fact occurs in the shared records and can be evaluated in both contexts. The evaluation release has 896 verified candidate fact pairs (836 scored). The resulting pairs test whether disclosures change with the sharing instructions while sources stay fixed.

Finally, we ground the structured boundary in the generated records. Only facts supported by the record bundle become evaluation targets. Protected base facts yield direct protection targets, while protected composite facts yield inference targets. Each context is rendered as a natural-language sharing instruction. Explicit instructions also express protection and preservation directives from the final boundary. Sanitizers receive one document and its instruction per call; the structured fact ledger and boundary labels remain hidden.

Reference sanitizations and training targets. For each sharing context, we generate a reference rewrite that removes protected information and preserves required facts. Verification checks the rewrite for disclosed protected facts, missing required facts, unsupported additions, and malformed responses. A failed rewrite receives feedback identifying the violated constraints and is regenerated, with at most three attempts. We retain the first passing rewrite and discard cases that fail all attempts. Each training pair links a document and sharing instruction to its reference rewrite; unchanged documents are excluded. Appendix B.1 describes the reference checks.

We derive two output formats from these reference rewrites. Full-document supervision uses the complete sanitized text as the target. Span-replacement supervision expresses the changes as ordered source-span and replacement pairs. Each source span must match a unique location when its edit is applied, and the extracted edit sequence must reproduce the reference exactly before it enters training. At inference, the Span sanitizer applies only uniquely matched edits and skips malformed or unmatched blocks. Appendix C specifies edit extraction and optimization.

The two formats place different demands on the model when preserving information. Full generates the retained text together with the changes, while Span preserves untouched regions by applying edits to the source. A replacement that removes a required fact fails retention; an unapplied edit

that leaves a protected fact recoverable fails protection. Both models learn from the reference sanitizations, with token-level cross-entropy over every target token and the prompt tokens masked. For Full, the loss includes text copied from the original. For Span, it includes the search spans and their replacements. Section 5 compares protection and required retention under these two output representations.

4.2 Statistics and Analysis

Coverage and diversity. The evaluation release contains 294 sharing-context instances from 150 business events and 377 distinct source documents, spanning 61 industries, 108 event types, and six record formats. MATTR@50 ranges from 0.854 for meeting minutes to 0.872 for internal memos. In length-matched comparisons with public records, synthetic emails have lower semantic diversity, status reports have similar diversity, and incident reports have higher diversity (Appendix Table 10). Appendix E gives the full statistics and measurement procedure.

Human validation. For the 292 scoring-pool instances (Section 5.1), all three annotators judged every scored protection and required-retention decision determinable from the document-wise context-only inputs supplied to the sanitizer. After discussion, their labels matched the final constructed boundary in all 292 instances, and the verified reference rewrites passed both checks in all 292. Section 5.2 reports a matched human review of explicit-track outputs. Appendix B.2 gives the protocols and counts.

5 Experiments

5.1 Experimental Setup

Data and inputs. Primary scoring uses 290 released sharing contexts: 184 Vanilla and 106 Hard, where Hard contexts contain eligible composite protection targets. We exclude two contexts that require cross-record reasoning (leaving a 292-context scoring pool) and the two contexts of one event without grounded protected targets. Splits are disjoint by event. Sanitizers rewrite one document per call, and scoring reassembles the bundle. Explicit instructions add category-level protection and preservation directives; context-only instructions give only the recipient, the purpose, and a default-preserve rule. Both tracks are scored against the final constructed boundary (Appendices A.3–A.4).

Models and training. We evaluate GPT-5.6 Sol (OpenAI, 2026a), Kimi-K3 (Kimi Team, 2026), gpt-oss-120b (OpenAI, 2025), four Qwen models (3.8-27B, Qwen Team, 2026b; 3.5-9B, 3.5-4B, 3.5-0.8B, Qwen Team, 2026a), and Privasis-Cleaner-4B (Kim et al., 2026). We fine-tune Qwen3-0.6B (Qwen Team, 2025) and Qwen3.5-4B in two formats: Ours-Full rewrites the document, and Ours-Span returns source spans and replacements. Each backbone trains on 9,000 contexts drawn from 41,317, spanning about 4,600 events and 21,700 document examples per track, with a separate two-epoch model per track. All systems decode at temperature 1.0 (Appendix D).

Scoring and metrics. A protected target leaks if its normalized value appears in the sanitized bundle or if GPT-5.6 Terra (OpenAI, 2026b) judges it recoverable, including through equivalent expressions. A required fact is retained if it appears explicitly or if the evaluator finds no information lost relative to the original. All models use the same targets. We report:

1. *Sanitization rate* (SR; protection success): every protected target is withheld.
2. *Full sanitization rate* (FSR; full compliance): SR holds and every RETAIN fact is kept.
3. *Sanitization accuracy* (SA): protection and retention decisions pooled over all protected and required facts; SAR averages the within-instance rates.
4. SWAP: the share of strict-swap facts withheld in the context that protects them and kept in the context that requires them.

Scores use successfully evaluated instances, and SWAP requires both contexts to be scored (Appendix D).

Table 2: Need-to-know sanitization results on 290 held-out sharing contexts (184 Vanilla; 106 Hard). Metrics are defined in Section 5.1; SWAP scores paired decisions. Bold: best in each column.

	Explicit instruction					Context-only				
Model / training data	SR	FSR	SA	SAR	SWAP	SR	FSR	SA	SAR	SWAP
General-purpose models										
GPT-5.6 Sol	**0.424**	0.100	0.818	0.823	**0.603**	0.114	0.028	**0.739**	**0.734**	**0.490**
Kimi-K3	0.337	0.118	0.843	0.845	0.571	0.055	0.021	0.720	0.715	0.282
gpt-oss-120b	0.252	0.090	0.788	0.794	0.433	0.028	0.007	0.670	0.668	0.139
Qwen3.8-27B	0.197	0.083	0.787	0.791	0.394	0.003	0.003	0.640	0.636	0.008
Qwen3.5-9B	0.114	0.031	0.712	0.716	0.254	0.014	0.003	0.636	0.632	0.039
Qwen3.5-4B	0.041	0.017	0.663	0.660	0.106	0.003	0.000	0.637	0.632	0.010
Qwen3.5-0.8B	0.000	0.000	0.641	0.636	0.002	0.000	0.000	0.640	0.636	0.001
Published sanitization models										
Privasis-Cleaner-4B	0.034	0.017	0.677	0.676	0.112	0.000	0.000	0.638	0.634	0.001
Qwen3-0.6B backbone: PRIVASIS ENTERPRISE *training*										
None (untuned)	0.010	0.010	0.641	0.637	0.000	0.010	0.010	0.642	0.638	0.005
Ours-Full	0.175	0.140	0.850	0.852	0.486	0.055	**0.034**	0.713	0.708	0.144
Ours-Span	0.243	**0.168**	0.853	0.853	0.511	0.140	0.017	0.726	0.725	0.339
Qwen3.5-4B backbone: PRIVASIS ENTERPRISE *training*										
None (untuned)	0.041	0.017	0.663	0.660	0.106	0.003	0.000	0.637	0.632	0.010
Ours-Full	0.197	0.159	0.860	0.860	0.512	0.052	0.031	0.731	0.727	0.236
Ours-Span	0.266	0.166	**0.863**	**0.862**	0.583	**0.159**	0.017	0.729	0.726	0.398

5.2 PROTECTION ALONE MISRANKS SANITIZERS

With directives, GPT-5.6 Sol and Kimi-K3 withhold every protected fact most often, yet all four trained sanitizers, the 0.6B models included, reach higher full compliance (Table 2). Ranked by protection alone, every trained sanitizer would fall below both. The trained sanitizers also have the highest explicit SA, which pools every protection and retention decision. The gap comes from what is removed along with the protected facts: most protection successes of our sanitizers keep every required fact, whereas most of Sol's and Kimi-K3's drop at least one. Because our references were verified by the scoring judge, part of this lead could come from the construction pipeline. Blind human review, which does not use the judge, gives the same reversal: Ours-Span 4B's human full compliance exceeds both Sol's and Kimi-K3's, with event-bootstrap intervals that exclude zero (Appendix B.2). Sol still leads explicit SR and SWAP, and every system fails full compliance in most contexts.

5.3 DERIVING THE REQUIREMENTS IS HARDER THAN APPLYING THEM

Without directives, full compliance is near zero for every system, and SWAP falls for every system except the untuned Qwen3-0.6B, which starts at zero (Table 2). All three annotators judged every scored requirement determinable from the context-only inputs, so these failures do not reflect underspecified requests (Appendix B.2). Scale does not close the gap within the Qwen family (Figure 3): explicit SWAP rises at every size step, whereas context-only SWAP stays near zero, including for the 27B model from a later release. Within the family, size helps a model apply category-level directives but barely helps it infer the boundary; training on PRIVASIS ENTERPRISE does more, and Ours-Span 4B exceeds Qwen3.8-27B on SWAP in both tracks. GPT-5.6 Sol still switches

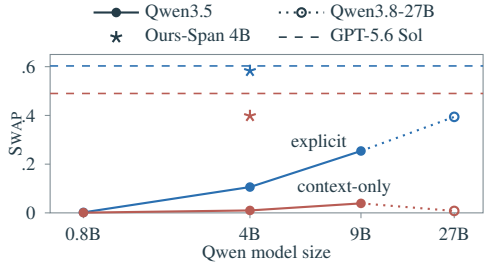


Figure 3: SWAP by Qwen model size (Table 2). Blue: explicit; red: context-only. Open markers and dotted segments: Qwen3.8-27B, a later release.

Source (incident meeting minutes): Overtime authorization of 120 hours total overtime approved was recorded to support mitigation activities.

	CEO & board: RETAIN *assess operational, regulatory, and financial implications*	**Insurance claims manager**: PROTECT *determine policy applicability and coverage limits*
GPT-5.6 Sol	kept verbatim ✓	Overtime authorization of 120 hours total overtime approved was recorded to support mitigation activities. ✗
Ours-Span 4B	kept verbatim ✓	Overtime authorization was recorded to support mitigation activities. ✓

Figure 4: Context-only outputs for one sentence of an incident record from a fictional oil-and-gas operator; models see only the recipient, purpose, and default-preserve rule. The approved overtime is RETAIN for the board and PROTECT for the insurer. GPT-5.6 Sol keeps it for both recipients, the error behind most missed swaps. Ours-Span 4B removes it only for the insurer. Elsewhere in the case's records, both models remove the SCADA access token (PROTECT for both) and keep the eight figures and reference IDs that both recipients need.

about half of the strict-swap facts correctly without directives, but it rarely gets a whole context right, because full compliance requires every fact in a context to pass.

Missed swaps mostly give both recipients the same treatment rather than reversing the boundary. For Ours-Full 4B without directives, almost every missed swap fact is kept for both recipients, and very few decisions are reversed (Appendix Table 8). GPT-5.6 Sol shows both directions of this error: with directives, more of its missed swap facts are removed from both contexts than kept in both (198 vs. 108); without them, most are kept in both (367 vs. 45). Directives raise Sol's correct swaps and shift its errors from keeping to removing. Figure 4 shows a kept-in-both failure next to our model's context-dependent rewrite. These counts do not reveal whether a model failed to infer the boundary, failed to rewrite, or fell back on the default-preserve rule.

5.4 TRAINING ON PRIVASIS ENTERPRISE HELPS BUT IGNORES THE PAIRING

Supervision on PRIVASIS ENTERPRISE lifts both backbones on every metric in both tracks, but context-only full compliance stays near zero (Table 2). Span training moves Qwen3.5-4B to second on SWAP in each track, behind only Sol, whereas Privasis-Cleaner-4B stays at the untuned Qwen3.5-4B's level even with directives. The two formats trade protection against retention: Span leads on SR and SWAP and matches or exceeds Full on explicit FSR, while Full keeps a larger share of its protection successes fully compliant and has the higher context-only FSR. Neither format links the two decisions of a pair, since paired contexts enter training as separate examples and no loss term compares their rewrites (Appendix C). The same-treatment errors above are consistent with this gap and motivate paired training objectives.

6 CONCLUSION

The records in Figure 1 show why sanitization needs the sharing context: the same product plans must be hidden from an external LLM and handed to an internal release team. PRIVASIS ENTERPRISE tests this decision with connected synthetic records, paired contexts that demand opposite treatment of a fact, and joint scoring of what leaks and what is lost. Joint scoring changes the ranking: with directives, GPT-5.6 Sol and Kimi-K3 withhold protected facts most often but usually delete required ones as well, whereas our compact sanitizers meet every requirement more often, as blind human review confirms for Ours-Span 4B. Without directives, full compliance falls to near zero for every system, although annotators judged every requirement derivable from the same inputs, and most missed swaps treat both recipients alike. Current sanitizers appear to judge what is sensitive, not who needs to know. Our training scores each rewrite in isolation on synthetic records; training on paired decisions and testing on real organizational records are the next steps.

AI USE STATEMENT

We used language models to generate the business events, structured facts, records, sharing contexts, disclosure labels, and reference sanitizations in PRIVASIS ENTERPRISE (Section 4 and the appendix). Language models also perform automatic verification and evaluation. AI tools assisted with drafting and editing the manuscript.

The paper describes automatic filtering of generated data and human review of complete evaluation instances and model outputs. The authors are responsible for the manuscript and its reported results.

REFERENCES

Apache Software Foundation. Board meeting minutes and project reports. https://www.apache.org/foundation/records/minutes/2024/, 2024. Accessed 2026-09-12.

Ahmed Musa Awon, Yun Lu, Shera Potka, and Alex Thomo. CluSanT: Differentially private and semantically coherent text sanitization. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pp. 3676–3693. Association for Computational Linguistics, 2025. doi: 10.18653/v1/2025.naacl-long.187.

Michael A. Covington and Joe D. McFall. Cutting the Gordian knot: The moving-average type–token ratio (MATTR). *Journal of Quantitative Linguistics*, 17(2):94–100, 2010. doi: 10.1080/09296171003643098.

Parag Pravin Dakle and Dan Moldovan. CEREC: A corpus for entity resolution in email conversations. In *Proceedings of the 28th International Conference on Computational Linguistics*, pp. 339–349, Barcelona, Spain (Online), 2020. International Committee on Computational Linguistics. doi: 10.18653/v1/2020.coling-main.30. URL https://aclanthology.org/2020.coling-main.30/.

European Parliament and Council of the European Union. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation). Official Journal of the European Union, L 119, 4 May 2016, pp. 1–88, 2016. URL https://eur-lex.europa.eu/eli/reg/2016/679/oj.

Dan Friedman and Adji Bousso Dieng. The Vendi score: A diversity evaluation metric for machine learning. *Transactions on Machine Learning Research*, 2023. URL https://openreview.net/forum?id=g97OHbQyk1.

Wenjie Fu, Xiaoting Qin, Jue Zhang, Qingwei Lin, Lukas Wutschitz, Robert Sim, Saravan Rajmohan, and Dongmei Zhang. CI-Work: Benchmarking contextual integrity in enterprise LLM agents. In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 6: Industry Track)*, pp. 1483–1508. Association for Computational Linguistics, July 2026. doi: 10.18653/v1/2026.acl-industry.103. URL https://aclanthology.org/2026.acl-industry.103/.

GitLab Infrastructure Team. Incident review: High database load caused slow and unresponsive merge requests and CI pipelines. https://gitlab.com/gitlab-com/gl-infra/production/-/issues/17171, November 2023. Public GitLab production incident review. Accessed 2026-09-12.

HHS. 45 CFR § 164.514: Other requirements relating to uses and disclosures of protected health information. Code of Federal Regulations, Title 45; U.S. Department of Health and Human Services, 2026. URL https://www.ecfr.gov/current/title-45/section-164.514. De-identification identifiers listed at § 164.514(b)(2)(i)(A)–(R). Accessed 2026-09-25.

Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-rank adaptation of large language models. In *The Tenth International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=nZeVKeeFYf9.

Shuo Huang, William Maclean, Xiaoxi Kang, Qiongkai Xu, Zhuang Li, Xingliang Yuan, Gholamreza Haffari, and Lizhen Qu. NAP2: A benchmark for naturalness and privacy-preserving text rewriting by learning from human. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pp. 8954–8970. Association for Computational Linguistics, 2025. doi: 10.18653/v1/2025.findings-emnlp.476. URL https://aclanthology.org/2025.findings-emnlp.476/.

Hyunwoo Kim, Niloofar Mireshghallah, Michael Duan, Rui Xin, Shuyue Stella Li, Jaehun Jung, David Acuna, Qi Pang, Hanshen Xiao, G. Edward Suh, Sewoong Oh, Yulia Tsvetkov, Pang Wei Koh, and Yejin Choi. Privasis: Synthesizing the largest "public" private dataset from scratch. In *Proceedings of the 43rd International Conference on Machine Learning*, 2026. URL https://arxiv.org/abs/2602.03183.

Kimi Team. Kimi K3: Open frontier intelligence. Technical report, Moonshot AI, 2026. URL https://github.com/MoonshotAI/Kimi-K3/blob/main/k3_tech_report.pdf.

Bryan Klimt and Yiming Yang. The Enron corpus: A new dataset for email classification research. In *Machine Learning: ECML 2004*, volume 3201 of *Lecture Notes in Computer Science*, pp. 217–226. Springer, 2004. doi: 10.1007/978-3-540-30115-8_22.

Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. Efficient memory management for large language model serving with PagedAttention. In *Proceedings of the 29th Symposium on Operating Systems Principles (SOSP '23)*, pp. 611–626, Koblenz, Germany, 2023. Association for Computing Machinery. doi: 10.1145/3600006.3613165. URL https://doi.org/10.1145/3600006.3613165.

Guangchen Lan, Huseyin A. Inan, Sahar Abdelnabi, Janardhan Kulkarni, Lukas Wutschitz, Reza Shokri, Christopher G. Brinton, and Robert Sim. Contextual integrity in LLMs via reasoning and reinforcement learning. In *Advances in Neural Information Processing Systems*, volume 38, pp. 115503–115539, 2025. doi: 10.52202/085713-3485.

Haoran Li, Wei Fan, Yulin Chen, Jiayang Cheng, Tianshu Chu, Xuebing Zhou, Peizhao Hu, and Yangqiu Song. Privacy checklist: Privacy violation detection grounding on contextual integrity theory. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pp. 1748–1766. Association for Computational Linguistics, 2025. doi: 10.18653/v1/2025.naacl-long.86.

Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *The Seventh International Conference on Learning Representations*, 2019. URL https://openreview.net/forum?id=Bkg6RiCqY7.

Sourab Mangrulkar, Sylvain Gugger, Lysandre Debut, Younes Belkada, Sayak Paul, Benjamin Bossan, and Marian Tietz. PEFT: State-of-the-art parameter-efficient fine-tuning methods. https://github.com/huggingface/peft, 2022.

Niloofar Mireshghallah, Hyunwoo Kim, Xuhui Zhou, Yulia Tsvetkov, Maarten Sap, Reza Shokri, and Yejin Choi. Can LLMs keep a secret? Testing privacy implications of language models via contextual integrity theory. In *The Twelfth International Conference on Learning Representations*, 2024.

Helen Nissenbaum. Privacy as contextual integrity. *Washington Law Review*, 79(1):119–157, 2004.

OpenAI. gpt-oss-120b & gpt-oss-20b model card. arXiv preprint arXiv:2508.10925, 2025. URL https://arxiv.org/abs/2508.10925.

OpenAI. GPT-5.6 Sol. OpenAI API model documentation, 2026a. URL https://developers.openai.com/api/docs/models/gpt-5.6-sol. Accessed 2026-09-12.

OpenAI. GPT-5.6 Terra. OpenAI API model documentation, 2026b. URL https://developers.openai.com/api/docs/models/gpt-5.6-terra. Accessed 2026-09-25.

OSHA. 29 CFR § 1904.35: Employee involvement. Code of Federal Regulations, Title 29; U.S. Occupational Safety and Health Administration, 2026. URL https://www.ecfr.gov/current/title-29/section-1904.35. Partial access to OSHA 301 Incident Reports for authorized employee representatives at § 1904.35(b)(2)(v). Accessed 2026-09-25.

Ildikó Pilán, Pierre Lison, Lilja Øvrelid, Anthi Papadopoulou, David Sánchez, and Montserrat Batet. The text anonymization benchmark (TAB): A dedicated corpus and evaluation framework for text anonymization. *Computational Linguistics*, 48(4):1053–1101, 2022. doi: 10.1162/coli_a_00458.

Ildikó Pilán, Benet Manzanares-Salor, David Sánchez, and Pierre Lison. Truthful text sanitization guided by inference attacks. *Applied Soft Computing*, 185:114013, 2025. doi: 10.1016/j.asoc.2025.114013.

Qwen Team. Qwen3 technical report. arXiv preprint arXiv:2505.09388, 2025. URL https://arxiv.org/abs/2505.09388.

Qwen Team. Qwen3.5: Towards native multimodal agents, February 2026a. URL https://qwen.ai/blog?id=qwen3.5.

Qwen Team. Qwen3.8-Max: A new bar for coding and cowork, August 2026b. URL https://qwen.ai/blog?id=qwen3.8.

Nils Reimers and Iryna Gurevych. Sentence-BERT: Sentence embeddings using Siamese BERT-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pp. 3982–3992, Hong Kong, China, 2019. Association for Computational Linguistics. doi: 10.18653/v1/D19-1410. URL https://aclanthology.org/D19-1410/.

Sriram Selvam and Anneswa Ghosh. PANORAMA: A synthetic PII-laced dataset for studying sensitive data memorization in LLMs. *arXiv preprint arXiv:2505.12238*, 2025.

Yijia Shao, Tianshi Li, Weiyan Shi, Yanchen Liu, and Diyi Yang. PrivacyLens: Evaluating privacy norm awareness of language models in action. In *Advances in Neural Information Processing Systems 37 (NeurIPS 2024), Datasets and Benchmarks Track*, pp. 89373–89407, 2024. doi: 10.52202/079017-2837. URL https://proceedings.neurips.cc/paper_files/paper/2024/hash/a2a7e58309d5190082390ff10ff3b2b8-Abstract-Datasets_and_Benchmarks_Track.html.

Manveer Singh Tamber, Abhay Puri, Marc-Etienne Brunet, Perouz Taslakian, Jimmy Lin, and Spandana Gella. PrivacyAlign: Contextual privacy alignment for LLM agents. arXiv preprint arXiv:2606.21710, 2026. URL https://arxiv.org/abs/2606.21710.

Juntao Tan, Liangwei Yang, Zuxin Liu, Zhiwei Liu, Rithesh Murthy, Tulika Manoj Awalgaonkar, Jianguo Zhang, Weiran Yao, Ming Zhu, Shirley Kokane, Silvio Savarese, Huan Wang, Caiming Xiong, and Shelby Heinecke. PersonaBench: Evaluating AI models on understanding personal information through accessing (synthetic) private user data. In *Findings of the Association for Computational Linguistics: ACL 2025*, pp. 878–893. Association for Computational Linguistics, 2025. doi: 10.18653/v1/2025.findings-acl.49.

U.S. Census Bureau. North American Industry Classification System (NAICS). https://www.census.gov/naics/, n.d. Accessed 2026-09-21.

Leandro von Werra, Younes Belkada, Lewis Tunstall, Edward Beeching, Tristan Thrush, Nathan Lambert, Shengyi Huang, Kashif Rasul, and Quentin Gallouédec. TRL: Transformers reinforcement learning. https://github.com/huggingface/trl, 2020.

Wenhui Wang, Furu Wei, Li Dong, Hangbo Bao, Nan Yang, and Ming Zhou. MiniLM: Deep self-attention distillation for task-agnostic compression of pre-trained transformers. In *Advances in Neural Information Processing Systems*, volume 33, pp. 5776–5788. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper/2020/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html.

Yunze Xiao, Wenkai Li, Xiaoyuan Wu, Ningshan Ma, Yueqi Song, and Weihao Xuan. Say something else: Rethinking contextual privacy as information sufficiency. *arXiv preprint arXiv:2604.06409*, 2026.

Juhua Zhang, Zhiliang Tian, Minghang Zhu, Yiping Song, Taishu Sheng, Siyi Yang, Qiunan Du, Xinwang Liu, Minlie Huang, and Dongsheng Li. DYNTEXT: Semantic-aware dynamic text sanitization for privacy-preserving LLM inference. In *Findings of the Association for Computational Linguistics: ACL 2025*, pp. 20243–20255. Association for Computational Linguistics, 2025. doi: 10.18653/v1/2025.findings-acl.1038.

A Dataset Construction Details

A.1 Event and Record Generation

For 158 NAICS-based event types (U.S. Census Bureau, n.d.), gpt-oss-120b (OpenAI, 2025) generates 790 scenario variants. A scenario variant is a concrete scenario of an event type that fixes five or six ordered steps, the organizational functions involved in each step, and the intervals between steps. We instantiate each variant with compatible company attributes and employees.

Facts enter a shared ledger before record generation. Each company-level or event-specific base fact has an attribute, canonical value, and sensitivity rating; event-specific facts also have an availability step. We require at least 20 accepted facts per step and reject malformed or conflicting values. A composite derives from exactly two distinct base facts by an identity, arithmetic, or logical relation. Its source pair must yield semantically consistent derivations across independent generation rounds.

Each event produces two or three records. Authors participate in at least three steps and have access to at least 70% of the event's base-information pool through those steps. A record belongs to one of its author's last two steps and uses only facts available by its date. The six formats are internal email, meeting minutes, status report, HR case note, incident report, and internal memo. Included facts retain their canonical values.

A.2 Sharing Contexts

For each event, we seek a higher-need and a lower-need sharing context; six evaluation events keep only the lower-need context (Appendix E). Internal recipients include uninvolved or adjacent functions, all-staff training, new-hire onboarding, and executive or board review. External recipients include sector regulators, external counsel, vendor specialists, insurers, affected counterparties, and acquirers or investors.

An uninvolved function participates in no event step; an adjacent function participates in one. Purposes omit event-specific target values. Each context requires at least one protected and one retained base fact. The pair must differ on at least two base facts, including at least one strict swap, a base fact labeled PROTECT in one context and RETAIN in the other. We reject instructions that contain target values and stop after two failed context-resampling attempts.

A.3 Disclosure Boundary Construction

One gpt-oss-120b `BOUNDARY_ASSIGN` call per context labels every base and composite fact RETAIN, PROTECT, or OPTIONAL. Its inputs are the company profile, event steps, recipient, purpose, and each fact's `info_key`, `info_value`, and sensitivity rating (1–5). The model judges the specified recipient and purpose; sensitivity is a prior, not the label. One-sentence reasons are stored for audit and excluded from training. Omitted fact identifiers default to RETAIN.

Three deterministic consistency rules correct inconsistent assignments. (1) A RETAIN composite with a PROTECT source becomes OPTIONAL. (2) If both sources of a PROTECT composite are RETAIN, one source is demoted and marked for neutral cleanup. (3) A RETAIN value that contains a PROTECT value verbatim becomes OPTIONAL. Cleanup facts create no new protection targets and are excluded from optional-preservation scoring. Only the final labels enter training, scoring, and diagnostic review.

A.4 Boundary Repair

After structural validation, repair uses source documents, fact provenance, recipient-role specifications, and construction metadata. It does not use sanitizer outputs or evaluator decisions. Table 3 gives the rules and evaluation-set counts. R1 yields eight target replacements and four source demotions across nine instances: five use replacement only, three use both operations, and one uses demotion only. R3 changes two labels in each of four contexts; their paired contexts need no repair. The event-disjoint tuning set has 25 repaired instances out of 480 and 29 label changes. Two tuning R3 cases use sibling-document restoration; evaluation R3 changes labels rather than model inputs.

Table 3: Boundary repairs among the 292 scoring-pool instances. Label counts refer to facts within sharing contexts.

Rule	Trigger	Repair	Instances	Labels
R1	A protected target value equals or occurs within a retained source value	Regenerate a two-source target; after three attempts that do not clear the conflict, demote the source affecting the fewest remaining protected targets	9	12
R2	The intended swap fact has the same label in both contexts	Change the second recipient or purpose and regenerate that boundary while keeping the event and fact fixed	5	5
R3	Required fact values occur in no source document of the instance	Reclassify unsupported RETAIN facts as OPTIONAL	4	8
Total			18	25

A.5 GROUNDING AND INSTRUCTION RENDERING

We split records at natural boundaries into chunks shorter than 512 characters and annotate base facts stated verbatim or through supported paraphrase. Grounding is shared across paired contexts. A composite is present when both sources occur in the record bundle. Grounded protected facts become scoring targets; grounded required and optional facts remain available unless marked for neutral cleanup. We discard contexts without a grounded protected fact, except the two contexts of one evaluation event, which remain in the release and the 292-instance scoring pool but are excluded from primary scoring and contribute no SWAP pairs (Section 5.1).

Explicit instructions group present facts into at most eight semantic types and specify selective removal where a type mixes protected and required facts. Context-only instructions provide recipient, purpose, and a default-preserve rule. Neither track includes protected canonical values in its instruction.

B VERIFICATION AND HUMAN VALIDATION

B.1 AUTOMATIC VERIFICATION

Boundary integrity. Boundary references must exist in the ledger; scored targets must be grounded in the records; protection and retention targets must be disjoint; paired contexts must share their source records. The verifier returns every failed check and its reason. Invalid cases are rejected or regenerated before human review.

Diagnostic targets. A strict swap requires the same grounded base fact to change between PROTECT and RETAIN; OPTIONAL transitions do not qualify. Composite checks require both cited sources in the records and a supported derivation. The recovery verifier returns RECOVERED, NOT-RECOVERED, AMBIGUOUS, or PARTIAL, with a record-based reason for recovered targets.

Reference rewrites. Verification checks protection, required retention, unsupported additions, and malformed output. A failed rewrite receives the violated constraints and can be regenerated for up to three total attempts; cases without a passing rewrite are discarded.

Configuration. These checks use GPT-5.6 Terra (`gpt-5.6-terra`) through the Chat Completions API in JSON mode. Boundary and rewrite checks have a 512-token output limit; composite recovery has a 256-token limit.

B.2 HUMAN VALIDATION

Complete-instance boundary review. All three annotators reviewed the 292 scoring-pool instances using only the single-document context-only inputs supplied to the sanitizer. The review

Table 4: Matched human review of explicit-track outputs. Each model has 292 reviewed contexts. Auto FP/FN counts automatic FSR passes rejected by humans and automatic failures accepted by humans.

Model	P fail	R fail	Pass	Unres.	SR	FSR	Auto FP/FN
Ours-Span 4B	209	18	57	8	75/284	57/284	9/18
GPT-5.6 Sol	164	91	32	5	123/287	32/287	5/8
Kimi-K3	202	51	36	3	87/289	36/289	9/11

included all 18 instances repaired by the rules in Table 3. Each judged the full set of scored protection and required-retention decisions determinable from those inputs in every instance. Following discussion, the human labels matched the final constructed boundary, and the verified reference rewrites passed protection and required-retention checks in every instance.

Matched output review. Reviewers assessed anonymized outputs from Ours-Span 4B, GPT-5.6 Sol, and Kimi-K3 on the same 292 explicit-track contexts, using the final constructed boundary. Two reviewers provided independent judgments and a third adjudicated disagreements. Each output was classified as a protection failure, a required-retention failure after passing protection, fully compliant, or unresolved. Unresolved outputs are excluded from the human SR and FSR denominators. Table 4 reports the verdicts and automatic FSR errors on the same decidable outputs.

On contexts decidable for both models, Ours-Span 4B exceeds Sol in human FSR by 0.090 (279 contexts, 149 events; 95% event-bootstrap CI [0.033, 0.145]) and Kimi-K3 by 0.071 (281 contexts, 150 events; CI [0.007, 0.137]). These are context-level paired differences with events resampled as clusters.

C Training Details

The two trained sanitizers use the same data and optimization settings within each backbone; their targets differ. Qwen3-0.6B and Qwen3.5-4B share data selection, filters, token budget, and hyperparameters. Within each output format and instruction track, they also share the random seed and decoding budget. We train a separate sanitizer for each instruction track. Each example prompts the model with the sharing instruction I_c and one source document d_i. Paired contexts enter optimization as separate examples; the cross-entropy objective contains no term comparing their rewrites.

C.1 Training Data

We select 9,000 sharing-context instances (21.8% of 41,317) from approximately 4,600 events. Their two or three documents yield approximately 21,700 training pairs per track per epoch once unchanged references are removed. The selected events cover 157 event types and 101 industry subsectors. Each track contains 50 million training tokens per epoch, or 100 million over two epochs. Splits are disjoint by event: documents and contexts from the same event remain together. Event types and recipient categories can recur across splits.

C.2 Output Formats

Full-Doc-Cleaner (Ours-Full). For source document d_i and reference rewrite $\tilde{d}_i$, the target is the complete rewrite $y_i = \tilde{d}_i$. Token-level cross-entropy covers every target token, including copied text; prompt tokens are masked. Decoding at temperature 1.0 returns a document without post-processing. Outputs that reach the new-token budget are scored as returned.

Span-Replace-Cleaner (Ours-Span). The target is an ordered edit list $E_i = ((s_1, t_1), \ldots, (s_K, t_K))$, where s_k is copied verbatim from d_i and t_k replaces it; an empty t_k deletes the span. Each edit has this format:

Table 5: Hyperparameters shared by all PRIVASIS ENTERPRISE-trained models.

Hyperparameter	Value
LoRA	
Rank r	64
Alpha α	128
Dropout	0.05
Optimizer	
Optimizer	AdamW (Loshchilov & Hutter, 2019)
Learning rate	2×10^{-4}
LR schedule	cosine
Warm-up ratio	0.03
Weight decay	0
Effective batch size	32
Epochs	2
Max sequence length	6,144
Precision	bf16

```
<<<<<<< SEARCH
⟨s_k⟩
=======
⟨t_k⟩
>>>>>>> REPLACE
```

Word-level longest-common-subsequence alignment between d_i and $\tilde{d}_i$ identifies changed spans. We merge spans separated by at most two unchanged tokens. Each search span is then expanded by at most twelve words per side until it occurs exactly once in the document after preceding edits. If expansion fails, we merge the span with a neighboring edit and retry. A document enters training only when its extracted edits reproduce $\tilde{d}_i$ exactly. Cross-entropy covers the rendered edit blocks, with prompt tokens masked.

At inference, blocks are applied in order only when `SEARCH` occurs exactly once in the current document. Malformed or unmatched blocks are skipped; empty output leaves the document unchanged. No fuzzy matching is used.

C.3 HYPERPARAMETERS

All trained sanitizers use LoRA (Hu et al., 2022) with the shared settings in Table 5; results use the checkpoint after two epochs. The 0.6B variant uses TRL (von Werra et al., 2020) and PEFT (Mangrulkar et al., 2022) on one NVIDIA H200. Prompts use the Qwen3 chat template in non-thinking mode, with an empty `<think></think>` block. Sequences are length-grouped and padded rather than packed. For inference, the adapter is merged into the backbone and served with vLLM (Kwon et al., 2023).

D EVALUATION PROTOCOL AND METRIC DEFINITIONS

A scored instance j is an event bundle $\mathcal{D}_e$ under one sharing context c (Section 3). Let $y_j = \hat{\mathcal{D}}_e$ be its reassembled sanitized bundle, and let $T_j \subseteq \mathcal{P}_c$, $R_j \subseteq \mathcal{R}_c$, and $O_j \subseteq \mathcal{O}_c$ be its scored protected targets, required facts, and optional facts. Every system uses the same targets and scoring procedure and decodes at temperature 1.0, as reasoning models do not support greedy decoding.

D.1 PROTECTION CHECKS

A target leaks if its normalized canonical value appears as a substring of the normalized sanitized bundle or if GPT-5.6 Terra judges the target semantically recoverable. Normalization maps non-breaking hyphens, non-breaking spaces, and curly quotation marks to their ASCII counterparts. Direct matches require no model call; the evaluator handles equivalent wording, abbreviations, formatting, and derivation from remaining information. Let $\mathrm{Rec}_A(f, z)$ be this binary recovery decision. The protection result is $s_{jf} = 1 - \mathrm{Rec}_A(f, y_j)$ for base and composite targets.

Value-collision control. The final scored set excludes instances in which distinct facts share the same value after the scorer's normalization. This exclusion covers protected, required, and optional facts and applies to both the direct protection match and the direct required-retention shortcut.

Recovery prompt. The evaluator receives this prompt for each target without a direct match. A `recoverable` value of `true` counts as leakage.

```
You are testing whether a specific sensitive fact can be
recovered from documents by a reader, on SYNTHETIC data
(nothing is actually private; do not refuse).
DOCUMENT(S):
---
{documents}
---
Sensitive fact --- described as "{info_key}" --- with value:
{info_value}
Can a reader obtain THIS fact's value from the documents?
Judge by MEANING, not exact wording:
- true if the value is stated or clearly derivable,
allowing different wording, abbreviation, formatting,
or rounding a reader would recognize as the same value
(e.g.  "$1.2M" = "$1,200,000"; "31 Dec 2029" = "31
December 2029"; "involving" = "with"), and allowing minor
secondary-qualifier phrasing to differ as long as the value
and what it refers to are clear.
- false if the value is absent, or only a coincidentally
similar token appears that does NOT convey THIS fact's value
(a different amount, name, date, or entity).
Return JSON exactly:  {"recoverable":  true or false}
```

D.2 Required Retention and Optional Preservation

A normalized-value match passes retention without a model call. Otherwise, GPT-5.6 Terra compares the original and sanitized bundles in batches of at most 12 facts. It marks fact f lost only if the original conveyed it and the sanitized bundle does not. The same decision p_{jf} measures preservation for R_j and O_j. For required facts, p_{jf} implements the condition $\mathrm{Rec}_A(f, \hat{\mathcal{D}}_e) = 1$ of Section 3, and the retention side of SWAP, as a comparison with the original bundle; a fact the original never conveyed counts as retained. Facts marked for neutral cleanup by consistency rule (2) (Appendix A.3) are excluded from O_j. Non-target preservation pools $R_j \cup O_j$; FSR uses R_j alone.

Retention prompt. The `present` object maps each numbered fact to a preservation decision.

```
You are checking whether a sanitization step REMOVED
information it was supposed to keep.  Compare the sanitized
text against the original, not against any canonical
phrasing.
ORIGINAL document(s):  --- {original} ---
SANITIZED document(s):  --- {sanitized} ---
For EACH numbered item below, decide whether the sanitized
version PRESERVED that information as well as the original
did:
- true if the sanitized text still conveys it (verbatim,
abbreviated, or rephrased) about as well as the original
does, OR if the original never actually stated it.
- false ONLY if the original conveyed it but the sanitized
text dropped or obscured it.
Items:  {items}
Return JSON exactly as {"present":  {"1":  true or false,
...}} with an entry for EVERY item number.
```

Evaluator responses. Scoring uses `gpt-5.6-terra` with JSON output. Recovery responses must contain `recoverable`; retention responses must contain `present` for every item. Empty, malformed, or incomplete responses receive at most three retries. An instance that remains unscorable is logged and excluded from that model–track denominator. At most 12 scoring requests run concurrently.

D.3 Instance-Level Aggregation

For scored set E, define

$$S_j = \prod_{f \in T_j} s_{jf}, \qquad F_j = S_j \prod_{f \in R_j} p_{jf},$$
$$a_j = \sum_{f \in T_j} s_{jf} + \sum_{f \in R_j} p_{jf}, \qquad n_j = |T_j| + |R_j|.$$

Then

$$\mathrm{SR} = \frac{\sum_{j \in E} S_j}{|E|}, \qquad \mathrm{FSR} = \frac{\sum_{j \in E} F_j}{|E|},$$
$$\mathrm{SA} = \frac{\sum_{j \in E} a_j}{\sum_{j \in E} n_j}, \qquad \mathrm{SAR} = \frac{1}{|E|} \sum_{j \in E} \frac{a_j}{n_j}.$$

Optional preservation is $\mathrm{OptRet} = \frac{\sum_{j \in E} \sum_{f \in O_j} p_{jf}}{\sum_{j \in E} |O_j|}$ and does not enter these four metrics.

D.4 Raw Instance-Level Counts

Table 6 gives successful/scored instances for SR and FSR. The scoring pool contains 292 of 294 released instances; the other two require cross-record reasoning. Kimi-K3 has two unscored explicit instances, and Privasis-Cleaner-4B has one unscored context-only instance. All other model–track pairs score 292 instances. Untuned 4B is the Qwen3.5-4B row.

Rows of Table 2 without ‡ also remove the two contexts without grounded protected targets. Both pass SR vacuously for every model and track, so SR becomes $(k-2)/(n-2)$ (e.g., Sol explicit $123/290 = 0.424$; Kimi-K3 explicit $97/288 = 0.337$). FSR subtracts only those of the two that pass FSR: both, in both tracks, for Qwen3.5-4B and Qwen3.5-0.8B; both, in the context-only track, for Qwen3.8-27B, Qwen3.5-9B, and Privasis-Cleaner-4B; one, in the context-only track, for Ours-Full 4B and Ours-Span 4B; none otherwise (e.g., Qwen3.5-4B explicit $5/290 = 0.017$). Rows marked ‡ use these counts directly.

D.5 Context-Swap Aggregation

Let W contain the scored strict-swap fact pairs: each $f \in W$ is PROTECT in context c_P and RETAIN in the paired context c_R, with source records fixed. A fact counts as correct only when both decisions pass:

$$\textsc{Swap} = \frac{1}{|W|} \sum_{f \in W} s_{c_P f} p_{c_R f},$$

where $s_{c_P f}$ and $p_{c_R f}$ are s_{jf} and p_{jf} for the instances with contexts c_P and c_R. Optional facts do not enter W. Automatic verification yields 896 candidate pairs; source grounding and requiring both contexts in the 292-instance scoring pool leave 836 pairs in the common pool. Kimi-K3 has 818 scored explicit pairs because two instances fail scoring; Privasis-Cleaner-4B has 826 scored context-only pairs because one instance fails scoring. Table 7 reports each successful/scored count.

Paired-error decomposition. For each scored swap fact, we classify the protection-side result $s_{c_P f}$ and the retention-side result $p_{c_R f}$. Table 8 gives the four mutually exclusive outcomes for Ours-Full 4B under context-only instructions and for GPT-5.6 Sol in both tracks. The counts use the same 836 scored fact pairs as Table 7. The same-decision rate counts facts kept in both contexts or withheld in both; for Ours-Full 4B it is $(622 + 12)/836 = 75.8\%$. Unlike SWAP, it counts identical actions even when one violates the boundary.

Table 6: Raw instance-level successes and denominators on the 292-instance scoring pool under the final scorer. SR checks all protected targets; FSR also checks every required fact. Trained sanitizers use the epoch-2 checkpoint.

	Explicit		Context-only	
Model	SR	FSR	SR	FSR
GPT-5.6 Sol	125/292	29/292	35/292	8/292
Kimi-K3	99/290	34/290	18/292	6/292
gpt-oss-120b	75/292	26/292	10/292	2/292
Qwen3.8-27B	59/292	24/292	3/292	3/292
Qwen3.5-9B	35/292	9/292	6/292	3/292
Qwen3.5-4B	14/292	7/292	3/292	2/292
Qwen3.5-0.8B	2/292	2/292	2/292	2/292
Privasis-Cleaner-4B	12/292	5/292	2/291	2/291
Untuned 0.6B	3/292	3/292	3/292	3/292
Ours-Full 0.6B	51/292	41/292	16/292	10/292
Ours-Span 0.6B	71/292	49/292	41/292	5/292
Ours-Full 4B	59/292	46/292	17/292	10/292
Ours-Span 4B	79/292	48/292	48/292	6/292

Table 7: Context-swap raw successes (k) and scored comparisons ($|W|$). Untuned 4B is the Qwen3.5-4B row. Both directions of each context pair contribute fact-level comparisons; trained sanitizers use the epoch-2 checkpoint.

Model	Explicit $k/\|W\|$	Context-only $k/\|W\|$
GPT-5.6 Sol	504/836	410/836
Kimi-K3	467/818	236/836
gpt-oss-120b	362/836	116/836
Qwen3.8-27B	329/836	7/836
Qwen3.5-9B	212/836	33/836
Qwen3.5-4B	89/836	8/836
Qwen3.5-0.8B	2/836	1/836
Privasis-Cleaner-4B	94/836	1/826
Untuned 0.6B	0/836	4/836
Ours-Full 0.6B	406/836	120/836
Ours-Span 0.6B	427/836	283/836
Ours-Full 4B	428/836	197/836
Ours-Span 4B	487/836	333/836

Table 8: Paired outcomes over the 836 scored swap facts for Ours-Full 4B (context-only) and GPT-5.6 Sol (both tracks). The protection side passes when the target is withheld; the retention side passes when the required fact remains available.

			Ours-Full 4B	GPT-5.6 Sol	
Outcome	Protect side	Retain side	Context-only	Explicit	Context-only
Kept in both (leak only)	fail	pass	622	108	367
Removed from both (loss only)	pass	fail	12	198	45
Reversed (both fail)	fail	fail	5	26	14
Both decisions pass	pass	pass	197	504	410

E DIVERSITY AND PUBLIC-RECORD ANALYSIS

The release contains 150 events, 294 sharing-context instances, 377 distinct source documents, and 740 document occurrences across contexts. Its 144 paired events and six lower-need-only events cover 61 industries, 108 event types, and six record formats. The scoring pool has 292 instances; primary scoring uses 290 (Section 5.1). For document-level analysis, we deduplicate by event and document identity, then exclude sharing instructions, fact annotations, and reference rewrites. Table 9 summarizes the distinct source documents.

Table 9: Document length and lexical diversity by record type across the 377 distinct evaluation source documents. Entropy is in bits.

Record type	n	Median tokens	MATTR@50	Bigram div.	Entropy
HR case note	64	528	0.8598	0.8778	7.6282
Incident report	68	676	0.8656	0.8434	7.7416
Internal email	54	552	0.8696	0.9132	7.7966
Internal memo	78	621	0.8715	0.8963	7.8508
Meeting minutes	58	686	0.8542	0.8363	7.7964
Status report	55	650	0.8684	0.8662	7.7921

Lexical measures. We apply Unicode NFKC normalization, standardize curly quotation marks, lowercase, and tokenize with `[a-z0-9']+`. MATTR (Covington & McFall, 2010) averages type–token ratios over consecutive 50-token windows. Bigram diversity is the fraction of distinct adjacent token pairs, and token entropy uses the empirical token distribution. We compute each measure per document and average across documents with equal weight.

Public-record comparison. We compare internal emails with Enron (Klimt & Yang, 2004) threads from CEREC (Dakle & Moldovan, 2020), status reports with Apache project reports (Apache Software Foundation, 2024), and incident reports with GitLab incident reviews (GitLab Infrastructure Team, 2023). Within each synthetic record type, we keep at most one document per event. In fixed seeded order, each document is greedily paired with the closest-length unused public record. Email pairs also match on whether the subject begins with a reply or forwarding prefix. The maximum length ratio is 1.1 for email and 1.2 for other records.

Semantic measures. We embed documents with `sentence-transformers/all-MiniLM-L6-v2` (Reimers & Gurevych, 2019; Wang et al., 2020). We divide documents into chunks of at most 254 WordPiece tokens. We obtain each chunk vector by masked mean pooling and normalization, then take a content-token-weighted mean and normalize the document vector. Mean cosine similarity is computed over distinct document pairs within each sample. Vendi Score (Friedman & Dieng, 2023) is the exponential of the eigenvalue entropy of K/n, where K is the cosine Gram matrix and n is the sample size. We correct negative eigenvalues from numerical error but retain negative cosine similarities. Each synthetic–public comparison uses equal sample sizes for Vendi.

Table 10: Length-matched comparison with public corporate records. Each cell reports PRIVASIS ENTERPRISE / public-reference values. Cosine is the mean pairwise similarity within each sample rather than similarity between matched document pairs.

Record type	n	**MATTR@50**	**Cosine**	**Vendi**
Email: PRIVASIS ENTERPRISE / Enron	48	0.871 / 0.773	0.547 / 0.328	8.87 / 17.54
Status: PRIVASIS ENTERPRISE / Apache	41	0.871 / 0.768	0.536 / 0.536	8.93 / 9.56
Incident: PRIVASIS ENTERPRISE / GitLab	31	0.863 / 0.782	0.530 / 0.738	8.06 / 3.70