AnonymEx: An Interactive Platform for Exploring and Evaluating Anonymization Techniques through Re-identification Attacks

Andrea Fieschi
fieschaa@ipvs.uni-stuttgart.de
University of Stuttgart, IPVS
Stuttgart, Germany

Christoph Stach
stachch@ipvs.uni-stuttgart.de
University of Stuttgart, IPVS
Stuttgart, Germany

Pascal Hirmer
pascal.hirmer@mercedes-benz.com
Mercedes-Benz AG
Stuttgart, Germany

Abstract

As data collection and analysis continue to expand, the need for effective and transparent anonymization becomes increasingly critical. Yet the growing diversity of anonymization techniques makes it difficult for developers to evaluate their guarantees and choose an appropriate method at design time. We present **AnonymEx**, an interactive platform that enables developers to *empirically compare anonymization techniques* using a unified metric: their susceptibility to re-identification attacks. The platform integrates (i) a *literature-grounded knowledge graph* linking anonymization techniques to documented re-identification attacks, (ii) *executable, containerized implementations* of both techniques and attacks, and (iii) an *assistant* that supports exploratory learning and helps users identify candidate techniques based on their requirements. During the demonstration, attendees explore the anonymization landscape, test techniques using provided datasets, run the associated attacks, and assess the assistant's suggestions within the *Anonymization-by-Design* workflow. *AnonymEx* thus provides a practical, transparent, and reproducible sandbox for design-time decision-making and cross-category empirical evaluation of anonymization techniques.

Keywords

anonymization, re-identification, empirical evaluation, knowledge graph, demo

1 Introduction

Selecting an appropriate anonymization technique remains a persistent challenge in privacy-preserving data sharing [6]. Although a wide range of techniques exists [1, 8, 9], their guarantees are highly heterogeneous, and most evaluation metrics are tied to specific methodological families. As a result, developers cannot meaningfully compare techniques across categories or justify their choices in a transparent, evidence-based way. A pragmatic alternative is to evaluate anonymization through its *susceptibility to re-identification attacks*, which provides a unified measure of privacy risk, hence, technique performance.

This paper adopts the perspective of the *Anonymization-by-Design* paradigm [2]. In this paradigm, anonymization is treated as a design-time decision: the choice of technique must be aligned with the data to be collected, the intended processing, and the system's functional requirements. As with other architectural decisions, selecting an anonymization technique early in the design process shapes how data will be handled, stored, and transformed throughout the system lifecycle.

However, making such a design-time choice is far from trivial. Developers must navigate a vast and fragmented landscape of anonymization approaches, each with different assumptions, guarantees, and known vulnerabilities [4]. Beyond choosing a technique, they must also assess the *quality* of that decision. We suggest to carry out this assessment by evaluating if the selected technique withstands realistic re-identification threats and if it preserves sufficient utility for the intended use case. This creates the need for practical tools that support clear evaluation.

AnonymEx addresses this need by providing a unified environment for exploring, testing, and validating anonymization options at design time. Rather than introducing new anonymity notions, the platform focuses on developer-centered, empirical assessment. It integrates (i) a literature-grounded knowledge graph linking anonymization techniques to re-identification attacks, (ii) a modular test-bed of containerized implementations of techniques and attacks, and (iii) an a chat window multiple LLMs and a local RAG model to help selecting and comparing anonymization strategies.

In Section 2 we illustrate how *AnonymEx* supports the Anonymization-by-Design workflow, followed by an explanation of its architecture and features in Section 3. Section 4 guides the readers through the demonstration scenarios that will be shown at the conference, and to conclude, a comparison between *AnonymEx* and its related systems can be found in Section 5.

2 Positioning *AnonymEx* within the Anonymization-by-Design Workflow

The Anonymization-by-Design paradigm structures privacy engineering around decisions made early in the system lifecycle, when data requirements and processing workflows are defined [2]. It distinguishes three core steps that guide developers toward selecting and justifying an anonymization strategy:

(1) **Requirement identification**: specify functional needs, data characteristics, and constraints relevant to the target system use case.
(2) **Technique selection**: identify anonymization techniques that are compatible with these requirements and capable of delivering the desired level of protection.
(3) **Assessment & validation**: empirically evaluate the shortlisted candidates to determine whether they provide sufficient privacy and utility in practice.

AnonymEx contributes to this workflow by making the assessment step (3) *executable, transparent, and repeatable*, while also providing indirect support for steps (1) and (2). Building on our previous demo work [3], which focused on modular exploration of anonymization techniques, the platform extends this foundation by linking techniques to their documented re-identification attacks and enabling empirical evaluation within the Anonymization-by-Design workflow. It enables users to:

EDBT '26, Tampere (Finland)

10.48786/edbt.2026.60

(i) explore the anonymization landscape through a literature-grounded knowledge graph, (ii) obtain technique recommendations based on stated requirements via the integrated assistant, and (iii) empirically test techniques by executing containerized anonymization modules together with their corresponding re-identification attacks. These capabilities allow users to understand which techniques are applicable, why they are recommended, and how they perform under realistic threat models.

Figure 1 illustrates where *AnonymEx* fits within the Anonymization-by-Design workflow by showing its role in the assessment phase and supporting design decisions.

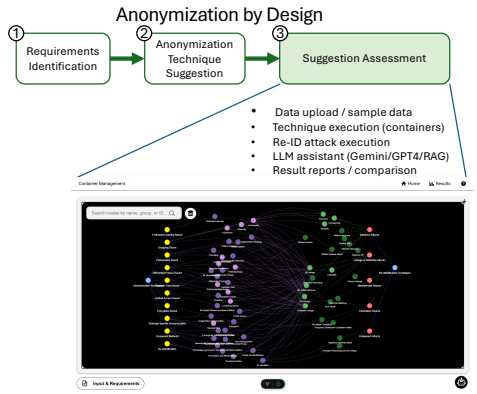

Figure 1: Anonymization-by-Design process. The figure illustrates the high-level design workflow (top) and a zoom on the assessment step where *AnonymEx* operates (bottom).

3 System Architecture and Implementation

Figure 2 provides an overview of the main components of *AnonymEx*. The platform consists of three building blocks implemented as modular, loosely coupled services:

a) **Anonymization techniques** implemented as executable Docker modules.
b) **Re-identification attacks** implemented with the same modular interface.
c) A **knowledge graph** that links techniques to documented attacks and determines which attack modules should be executed for a given technique.

These components interact through well-defined APIs and can be extended independently: researchers can add new techniques or attacks without modifying the rest of the system. The modular architecture ensures reproducibility, supports experimentation across heterogeneous technique families, and provides a clean separation between reasoning, execution, and visualization.

3.1 Knowledge Graph

The knowledge graph (Neo4j) organizes information about anonymization techniques, technique categories, and re-identification attacks. Edges connect techniques to the attacks shown in the literature to compromise them, and each edge stores provenance information referring to the corresponding publication. The schema is based on our previous work [4], allowing structured queries

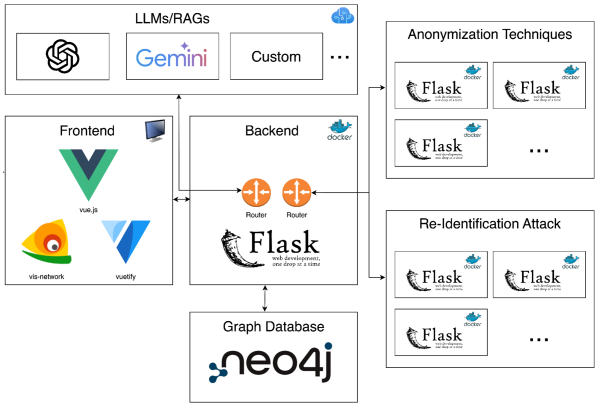

Figure 2: System architecture: the Neo4j knowledge graph encodes techniques, categories, and attack links; anonymization techniques and re-identification attacks are implemented as Docker microservices; the Flask backend orchestrates experiments; and the Vue.js frontend exposes the interactive UI and assistant.

such as: *"Which attacks are applicable to techniques in the generalization family?"* or *"Which techniques are vulnerable to linkage attacks requiring background attributes X and Y?"*

The current instance of the knowledge graph contains:

- 10 anonymization categories,
- 35 anonymization techniques,
- 24 re-identification attacks,
- 105 edges linking techniques to applicable attacks, and
- 134 literature references supporting these edges.

These edges are also directly accessible through frontend visualization as shown in Figure 4. When users click on an edge between a technique and an attack, a tooltip displays the associated literature reference, supporting the claim that the attack can be performed against the respective technique.

The knowledge graph serves three purposes within the platform: (i) it populates the interactive visualization in the frontend, (ii) it determines which re-identification attacks can be executed against a selected technique, and (iii) it helps users gain an overview of available anonymization techniques [4] and their weaknesses, and potential privacy risks in practical deployments.

3.2 Executable Modules

Each anonymization technique and re-identification attack is implemented as a standalone Docker microservice exposing a standardized REST API. Technique modules accept datasets and technique-specific parameters and return anonymized outputs. The designated attack modules take these outputs and attempt targeted re-identification, generating structured reports that quantitatively summarize attack success rates, required adversarial background knowledge, and recovered attributes.

This modular setup extends the patterns introduced in our previous demo work [3]. It ensures isolation, i.e., a re-identification attack crashing does not affect other modules, and makes it easy to integrate community-contributed implementations. Because all modules adhere to the same interface, anonymization pipelines across different technique families can be replicated and compared simply by adding new containers.

3.3 Backend and Frontend

A Flask-based backend orchestrates end-to-end experiments: it resolves which attacks correspond to a selected technique via knowledge-graph queries, launches the required containers, manages data flow between modules, and aggregates results for presentation. The backend also records execution metadata to support reproducibility and debugging.

The frontend, implemented in Vue.js and illustrated in Figure 3, provides three coordinated views: (i) the knowledge graph exploration panel, (ii) the experiment runner for techniques and attacks testing, and (iii) the assistant interface. Updates provided via WebSockets allow users to follow experiment progress live as results are generated.

3.4 Assistant Stack

The platform includes an assistant presented as a chat interface through which users can interact with different language models. Users can choose to connect the interface to ChatGPT, Gemini, or a local RAG system built on Llama 3. The local model operates over a curated knowledge base of approximately 3,500 anonymization-related papers, enabling suggestions that are grounded in a controlled and transparent corpus rather than external sources. While the assistant does not perform automated reasoning over system components, it helps users effectively explore complex anonymization concepts, understand technique properties, and obtain literature-backed guidance during the design-time decision making process.

The assistant offers a single chat-based interface in which users can describe their use case or ask questions about anonymization concepts. Based on the connected model, the assistant suggests relevant anonymization techniques and points to literature that discusses their properties, assumptions, and known vulnerabilities. Users can then take the suggested techniques and evaluate them within the platform by running the corresponding anonymization modules and re-identification attacks. This interaction supports the design-time workflow by helping users identify plausible candidates and understand their trade-offs before performing empirical validation.

4 Demonstration Scenarios

During the demo session, visitors can engage with *AnonymEx* through three guided scenarios. Each scenario highlights a different aspect of the platform, showcasing how users can explore anonymization options, empirically test techniques, and validate design-time decisions.

4.1 Scenario 1: Testing Anonymization Techniques

The central scenario demonstrates the core functionality of *AnonymEx*: executing anonymization techniques and empirically evaluating their robustness against documented re-identification attacks. Visitors may upload their own dataset or select one of the platform's example datasets. They can then choose one or more anonymization techniques to apply. Each selected technique is executed inside its dedicated Docker module, producing an anonymized version of the input dataset.

Based on the knowledge graph, the system automatically determines which re-identification attacks are applicable and launches the corresponding attack modules. The platform returns a structured evaluation that includes:

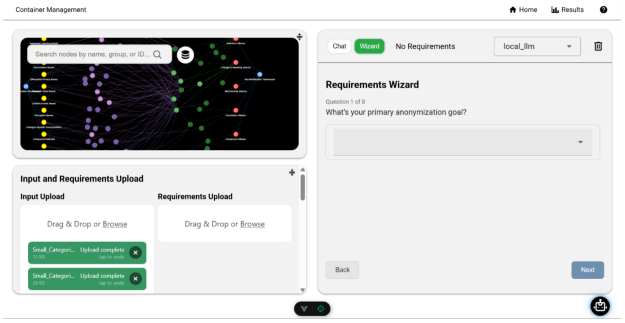

Figure 3: Main application interface integrating the knowledge graph exploration panel, assistant interface, and experiment dashboard.

- quantitative attack success rates,
- required adversarial background knowledge per attack,
- recovered attributes or records, and
- utility indicators for common analytical tasks (e.g., aggregation accuracy, classification performance).

This scenario shows how cross-category comparisons can be achieved through the susceptibility to re-identification attacks.

Beyond this core scenario, the next two scenarios offer complementary perspectives on the uses of *AnonymEx* that can foster open discussion with visitors of this demonstration.

4.2 Scenario 2: Assistant-Guided Exploration in the Anonymization Sandbox

The second scenario demonstrates how the platform can be used as a learning and exploration tool during the earlier stages of design-time. Visitors interact with the assistant through the chat interface, where they can describe their use case or ask conceptual questions about anonymization. Depending on the selected model (ChatGPT, Gemini, or the local Llama 3 RAG system), the assistant suggests relevant anonymization techniques.

Using these suggestions as a starting point, users can experiment directly with the anonymization techniques available on the platform. The demo environment includes several dummy datasets that allow attendees to try out techniques without providing their own data. Users can anonymize these datasets, observe how each technique transforms the data, and then apply the corresponding re-identification attacks to assess their privacy protection in practice—mirroring the workflow from Scenario 1.

4.3 Scenario 3: Knowledge Graph Exploration

In this scenario, visitors explore the anonymization landscape through the interactive knowledge graph shown in Figure 4. The KG serves both a conceptual and a functional role: it provides a high-level overview of how anonymization techniques and re-identification attacks interrelate *and* it is used by the platform to determine which attack modules to run for a selected technique.

The visualization allows filtering techniques by category, abstraction level, or specific methodological properties, and inspecting edges that link techniques to the re-identification attacks documented in the literature. Clicking on an edge reveals the corresponding publications where the connection is explained, giving users immediate access to a source and deeper context.

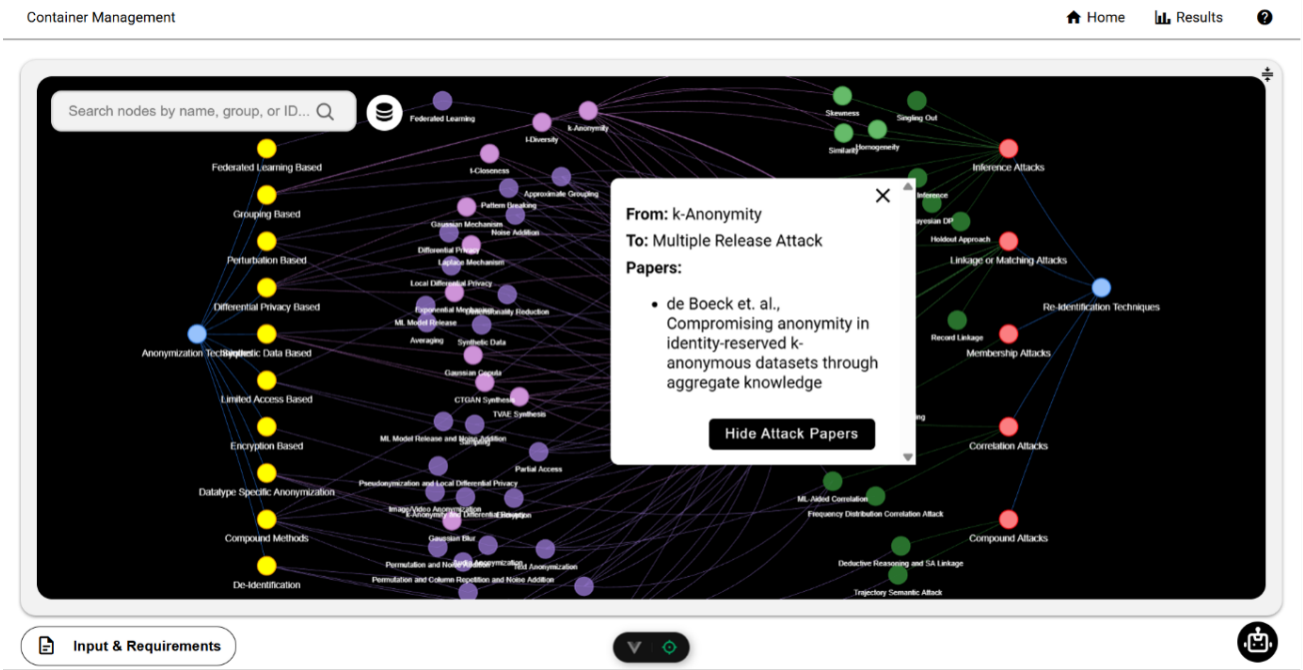

Figure 4: Interactive KG view: nodes represent techniques, categories, and attacks; edges encode literature-backed evidence linking specific attacks to specific techniques.

5 Related Systems

Several existing systems support the evaluation of anonymization or privacy-preserving techniques, but most focus on a single methodological family. For example, DPBench [5] supports exploration of differential privacy mechanisms, while ARX [7] provides a comprehensive environment for grouping-based techniques. Other tools offer implementations and utility metrics, but do not cover the full spectrum of anonymization approaches.

As discussed in Section 2, Step 3 of Anonymization-by-Design requires assessment under realistic threat models, which existing systems only partially support. They rarely provide a unified overview of techniques and documented re-identification vulnerabilities, and cross-category comparisons remain infeasible due to incompatible evaluation metrics [2]. Consequently, current tools fall short of supporting Step 3.

AnonymEx addresses these gaps by combining (i) a literature-grounded knowledge graph linking techniques and re-identification attacks, and (ii) executable modules that enable attack-based comparison across heterogeneous techniques. Rather than replacing prior systems, *AnonymEx* complements them with a practical, design-oriented perspective for transparent and reproducible cross-category assessment.

6 Conclusion and Future Work

We presented *AnonymEx*, a platform that integrates attack-based evaluation into the Anonymization-by-Design workflow. By combining a literature-grounded knowledge graph with executable anonymization and re-identification modules, it provides a practical sandbox for understanding and comparing techniques. Future work includes expanding the graph and module library, adding automated evaluation, and conducting user studies to assess design-time support.

Acknowledgments

Special thanks go to our students, Prajakta Deshpande, Richard Below, and Daniel Peys, for assisting with the implementation of parts of the *AnonymEx* application.

References

[1] Yves-Alexandre De Montjoye, César A Hidalgo, Michel Verleysen, and Vincent D Blondel. 2013. Unique in the crowd: The privacy bounds of human mobility. *Scientific Reports* 3, 1 (2013), 1376.

[2] Andrea Fieschi, Pascal Hirmer, Sachin Agrawal, Christoph Stach, and Bernhard Mitschang. 2024. HySAAD–A Hybrid Selection Approach for Anonymization by Design in the Automotive Domain. In *2024 25th IEEE International Conference on Mobile Data Management (MDM)*. IEEE, 203–210.

[3] Andrea Fieschi, Pascal Hirmer, and Christoph Stach. 2025. Discovering Suitable Anonymization Techniques: A Privacy Toolbox for Data Experts. In *Datenbanksysteme für Business, Technologie und Web (BTW 2025)*. Gesellschaft für Informatik.

[4] Andrea Fieschi, Christoph Stach, Pascal Hirmer, and Bernhard Mitschang. 2025. Navigating the Anonymization Landscape: An Ontological Support for Developers to Make Informed Privacy Decisions. In *Information Systems Security and Privacy. ICISSP 2025*. Springer. To appear.

[5] Michael Hay, Ashwin Machanavajjhala, Gerome Miklau, Yan Chen, and Dan Zhang. 2016. Principled evaluation of differentially private algorithms using dpbench. In *Proceedings of the 2016 International Conference on Management of Data*. 139–154.

[6] Abdul Majeed and Sungchang Lee. 2020. Anonymization techniques for privacy preserving data publishing: A comprehensive survey. *IEEE Access* 9 (2020), 8512–8545.

[7] Fabian Prasser, Florian Kohlmayer, Harald Berger, and Klaus A. Kuhn. 2020. Flexible data anonymization using ARX—Current status and challenges ahead. *Software - Practice and Experience* 50, 7 (2020), 1277–1304. https://www.scopus.com/inward/record.uri?eid=2-s2.0-85081026799&doi=10.1002%2fspe.2812&partnerID=40&md5=30161bec20eb30529aefb955b5822b72

[8] Luc Rocher, Julien M Hendrickx, and Yves-Alexandre De Montjoye. 2019. Estimating the success of re-identifications in incomplete datasets using generative models. *Nature Communications* 10, 1 (2019), 3069.

[9] Latanya Sweeney. 1997. Weaving technology and policy together to maintain confidentiality. *The Journal of Law, Medicine & Ethics* 25, 2-3 (1997), 98–110.