

A Configurable Anonymisation Service for Semantically Annotated Data: A Case Study on REC Data

Paul Feichtenschlager^{1,*}, Christoph Fabianek¹, Fajar J. Ekaputra², Sebastian Haas³ and Gabriel Unterholzer¹

¹*OwnYourData.eu, Michael Scherz-Straße 14, 2540 Bad Vöslau, Austria*

²*Vienna University of Economics and Business, Welthandelsplatz 1, 1020 Vienna, Austria*

³*Hinterland Systems, Hörlgasse 10, 1090 Vienna, Austria*

Abstract

Due to the increasing push for data sharing in Renewable Energy Communities (RECs)—driven by evolving regulations and energy transition goals—there is a growing need for trusted, privacy-preserving data-sharing infrastructures. One of the key challenges for RECs is enabling internal and external data sharing while minimizing privacy risks and preserving data value. In this work, we introduce a generic, configurable online anonymisation service for semantically annotated data as an extension to the Semantic Overlay Architecture (SOyA). Our approach employs an automated rule-based anonymisation pipeline that reduce re-identification risks while maintaining the utility of shared data. We demonstrated how this service can support compliant, secure, and practical data-sharing practices in real-world REC scenarios.

Keywords

Data anonymisation, Renewable Energy Communities, Governance, Semantic annotation

1. Introduction

The digital transformation driven by decentralized energy resources and smart grids presents opportunities and challenges. Renewable Energy Communities (RECs) enable joint renewable energy production, consumption, and management, facilitated by smart meters capturing high-resolution household energy data. Austria mandates smart meter installation for 95% of households by 2024 [1], collecting granular consumption and production data at 15-minute intervals. Under §84 ElWOG, network operators must provide daily data by default and quarter-hourly data upon consent [2]. This raises critical privacy concerns, as smart meter data qualifies as personal data under GDPR [3], potentially revealing sensitive behavioural patterns. Hence, RECs face the challenge of sharing energy data without compromising privacy. Effective privacy-preserving mechanisms are required to protect individual identities while maintaining data utility for analysis.

Currently, several tools and services are available for data anonymisation. For example the EU-funded open-source tool Amnesia supports techniques such as masking, k-anonymity, km-anonymity, and the computation of demographic statistics [4]. The ARX Data Anonymisation Tool offers a wide range of anonymisation capabilities, including generalization, suppression, and risk analysis [5]. However, to the best of our knowledge, none of these tools offer an open, online service that enables the annotation of data with a layer that explicitly defines and enforces its anonymisation requirements. Therefore, we developed a freely accessible online service, enabling users to experiment with anonymisation features prior to integrating them into their workflows. To ensure transparency, community engagement, and continuous improvement, we released the entire service as open-source software, encouraging feedback and contributions. Additionally, the implementation emphasizes open standards and semantic technologies, aligning with best practices to enhance interoperability and adoption. This approach lowers the entry

NeXt-generation Data Governance workshop 2025, September 03, 2025, Vienna, AT

*Corresponding author.

✉ paul@ownyourdata.eu (P. Feichtenschlager); christoph@ownyourdata.eu (C. Fabianek); fajar.ekaputra@wu.ac.at (F. J. Ekaputra); seb@hinterland.systems (S. Haas); gabriel@ownyourdata.eu (G. Unterholzer)

ORCID 0009-0002-4410-8796 (C. Fabianek); 0000-0001-7116-9338 (F. J. Ekaputra)



© 2022 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

barrier, allowing the service to be effortlessly deployed both publicly and privately, facilitated by an MIT license and the availability of a ready-to-use Docker image. The service configuration can be performed via a YAML file, without necessitating prior expertise in software engineering.

This paper introduces a practical solution addressing what we consider—and others have also highlighted [6, 7]—as one of the key open challenges in this regulatory and technical landscape: the privacy-preserving sharing of high-resolution PV production data within energy communities. We present the design and implementation of an anonymisation service initially motivated by the requirements for anonymisation explicitly stated in the EU Data Governance Act (DGA) [8], exemplified through its integration into a DGA-compliant data intermediary to automatically anonymise REC data prior to sharing. While this intermediary use case highlights the immediate relevance and practical necessity of anonymisation under the DGA, our service aims to be a versatile, general-purpose tool, lowering barriers to adoption by enabling organisations such as data intermediaries and data altruism organisations to exchange data securely, without compromising individual privacy. Our contribution lies in demonstrating how automated, rule-based anonymisation pipelines can reduce re-identification risks while maintaining data utility, thus enabling compliant and trustworthy data sharing in real-world Renewable Energy Communities (RECs) scenarios.

2. Background and Related Work

This section outlines the background and related work relevant to this work. First, the foundations of SOyA, as a core component of our service are explained. Afterwards, this sections gives an overview of the most prominent anonymisation techniques.

2.1. Semantic Overlay Architecture (SOyA)

SOyA (Semantic Overlay Architecture) is an open-source framework designed to facilitate the authoring, validation, and transformation of semantically annotated data models [9]. It addresses the need for flexible yet rigorous data modeling in decentralized and privacy-sensitive environments—such as those found in energy data ecosystems. SOyA introduces a modular architecture based on two primary concepts: Bases, which define the core structure of a dataset, and Overlays, which semantically enrich the base by specifying annotations, classifications, constraints, or transformations. These overlays enable the separation of concerns by allowing different semantic layers (e.g., validation, encoding, personal data classification) to be composed as needed.

SOyA (data) structure are usually authored in YAML for ease of use and are ultimately translated into an ontology, allowing compatibility with Semantic Web technologies and Linked Data tools. To this end, we automatically generated an ontology in JSON-LD and Turtle serialization for each registered SOyA structure¹. SOyA also supports the handling of conventional flat JSON by converting it into JSON-LD compliant to the ontology, enabling the subsequent application of semantic web operations and queries. By incorporating these user-friendly features, SOyA aims to lower the barrier to entry, facilitating broader adoption and easier integration of semantic web technologies into diverse workflows.

For the use case of renewable energy communities, SOyA enables the explicit modeling of personal data attributes (e.g., energy usage, timestamps) and facilitates the annotation of privacy-relevant elements. This functionality ensures that data intermediaries can apply standardized anonymisation or minimization logic before external sharing, in line with GDPR and Data Governance Act requirements.

2.2. Data Anonymisation

As data volumes grow and data mining technologies become more prevalent, concerns regarding data protection have also intensified. Personal data, in particular, is highly sensitive and demands

¹For the AnonymizationDemo example used throughout this paper, the SOyA structure (YAML) and its respective TTL and JSON-LD serializations can be accessed in the following links: YAML <https://soya.ownyourdata.eu/AnonymisationDemo/yaml>, TTL <https://soya.ownyourdata.eu/AnonymisationDemo/ttl>, and JSON-LD: <https://soya.ownyourdata.eu/AnonymisationDemo>

careful handling [10]. Anonymisation aims to protect the privacy of personal data while preserving the structural integrity and utility of the underlying data [11]. Fung et al. [12] identify three key objectives in the anonymisation process: adherence to privacy constraints, preservation of data utility, and retention of data truthfulness. To achieve this, direct identifiers and quasi-identifiers - attributes that can be used to reidentify individuals alone or in combination - must be removed or transformed appropriately [13].

One of the primary threats that anonymisation techniques seek to address is the linking attack, in which two datasets are combined using shared quasi-identifiers to reveal sensitive information. For example, voter registration data can be linked to medical records using attributes such as ZIP code, age, and gender, potentially making individuals' medical histories identifiable [14]. To mitigate this risk, Samarati and Sweeney introduced the concept of k -anonymity, which requires that each record in a dataset be indistinguishable from at least $k - 1$ other records with respect to a set of quasi-identifiers.

Data Anonymisation Techniques. To address the challenges inherent in the anonymisation process, various approaches have been proposed in the literature. We introduce two of the most prominent techniques in the following.

- **Generalization** involves replacing specific attribute values with more abstract or less precise representations, thereby grouping multiple records under a generalized value [11]. This technique is applicable for attributes that exhibit hierarchical or multi-level structures—such as geographic locations or for numerical attributes, where values can be grouped into ranges or buckets [10].
- **Randomization** adds a random salt to all values of an attribute. It is suitable only for attributes with cardinal measurement scales, where value distances are defined [16].

Data Governance. The governance of personal data within the European Union is shaped by a comprehensive legal framework that seeks to balance data utility with fundamental rights to privacy and data protection. Central to this framework is the General Data Protection Regulation (GDPR) [3], which has been in force since 2018 and applies to all organizations processing personal data of EU citizens. The GDPR establishes key principles such as data minimization, purpose limitation, and accountability, and it grants data subjects enforceable rights including access, rectification, and erasure [17, 18].

An important GDPR mechanism is *anonymisation*, the irreversible removal of personal identifiers making data subjects unidentifiable. Such anonymised data is exempt from GDPR, enabling secondary uses like research, analytics, or data sharing without additional constraints [19], making anonymisation crucial for compliant data utilization [17]. Complementing GDPR, the *Data Governance Act* (DGA) introduces further provisions promoting secure and trustworthy reuse of data—particularly from public bodies, data shared altruistically, and exchanges via neutral intermediaries. The DGA outlines registration, obligations, and neutrality rules for data intermediaries, extending regulatory oversight from individual rights to ecosystem-level accountability and trust.

Together, the GDPR and DGA form a layered governance model that simultaneously protects individuals and facilitates responsible data innovation within the EU.

3. The Data Anonymiser Service

In this work, we developed a service that applies anonymisation techniques to datasets containing personal data. The primary objective was to create a generic and configurable solution that ensures compliance with the General Data Protection Regulation (GDPR). To reduce the entry barrier for users, the service is designed to be easily configurable through SOyA using a simple YAML file.

The degree of anonymisation dynamically adapts to the size of the dataset and the number of attributes to be anonymised. This ensures that a sufficient level of anonymisation is applied while preserving the overall structure and utility of the original data set as much as possible.

The anonymisation service is provided as a free-to-use, web-based portal², which comprises of a lightweight single-page application built with Rails offers a responsive graphical user interface (GUI) that abstracts all protocol details, so that domain experts without programming experience can execute the same workflows that advanced users may invoke programmatically through the OpenAPI endpoint³. The front-end and back-end components are packaged as Docker images and are continuously deployed on an OwnYourData Kubernetes cluster; nightly builds are published under an MIT licence to foster reproducibility and community contributions. The open-source code is available in the accompanying GitHub repositories⁴.

To meet GDPR and DGA accountability requirements, every request is logged in an append-only audit ledger, while raw input files are stored only for the duration of processing and deleted immediately afterwards. Organisations with stricter data-sovereignty needs can deploy the service on-premises via the provided Docker images, achieving functional parity with the managed Software-as-a-Service instance but retaining full control over network boundaries and compliance monitoring. This dual provisioning model—public SaaS and self-hosted container images—minimises the entry barrier for exploratory usage while ensuring that production environments can satisfy sector-specific governance constraints.

3.1. Data Anonymisation Workflow

The service accepts requests containing both the data to be anonymised and a URL pointing to a valid anonymisation configuration. The process begins by fetching the configuration from the provided URL. Once validated, an anonymiser object is created for each attribute based on the configuration. This is enabled through an anonymiser interface, which is implemented by various anonymisation strategies described in Section 3.2.




Figure 1: Workflow of anonymisation Service

To apply anonymisation, the input data is first converted into an attribute-oriented schema, and for each attribute, the service aggregates values across all instances, while allowing attributes to remain blank for instances where no value is provided. After the transformation, anonymisers are applied to the aggregated attribute values. The number of values per attribute remains unchanged throughout the process. If an instance lacks a value for a specific attribute, that attribute remains empty in the anonymised dataset. Finally, the anonymised attributes are transformed back to an instance oriented data structure (might require defining specific fields in SOyA, e.g. min/max), to produce the final output.

3.2. Data Anonymisation Implementation

The anonymisation service is designed to allow the seamless integration of custom anonymisation implementations. Each anonymiser must implement a service interface that accepts a list of attribute values along with the total attribute count, and returns a corresponding list of anonymised values. This work introduces two anonymiser implementations: (a) **Generalisation** and (b) **Randomisation**, each tailored to different data types and anonymisation strategies.

Both anonymisation techniques require the definition of a specific number of groups. To ensure a sufficient level of anonymity, the desired number of generalization groups g is calculated based on the dataset size k and the number of anonymised attributes n . The objective is to achieve a 99% probability that no instance in the dataset remains unique after anonymisation.

²<https://anonymiser.ownyourdata.eu>

³<https://anonymizer.go-data.at>

⁴<https://github.com/OwnYourData/anonymisation-service>

To determine the number of groups under this requirement, an approach by Jiang et al. [20] to quantify re-identification risk, was adapted. Assuming that the attributes in the input data are statistically independent and that all groups are of equal size, a formula was derived to guarantee that the probability of any individual being uniquely identifiable is less than 1%. Given k as the number of individuals and n as the number of anonymised attributes, the number of required groups g is computed as:

$$g = \frac{1}{\sqrt[n]{1 - \sqrt[k]{1 - \sqrt[k]{0.99}}}} \quad (1)$$

Generalization. In generalization, attribute values are grouped into categories, and the corresponding group labels are written in the anonymised dataset. The number of groups g is defined in Equation 1.

Values are assigned to buckets in such a way that each group contains the same number of values. Outliers are not assigned to separate groups, avoiding easy re-identification. Instances are first sorted by value and then assigned to groups on the basis of their position. Each group is then labeled on the basis of the values it contains.

Generalization also supports object-type attributes, assuming a hierarchical structure is defined in the configuration. The algorithm traverses the hierarchy from the most specific level upward, reducing each data point to a generalized value. A valid generalization is accepted when:

- The number of resulting groups is less than or equal to the calculated number of groups g .
- Each group contains at least $\frac{k}{2g}$ elements, where k is the size of the dataset and g is the calculated number of groups.

These two requirements ensure that the resulting generalization maintains the desired level of abstraction and prevents the formation of groups that consist only of outliers, which would otherwise be easily identifiable.

Randomization. In the randomization process, noise is added to attribute values based on a normal distribution. The intensity of this noise depends on the distribution of the dataset: data points in sparse regions receive increased noise to enhance their anonymity.

The parameter i represents the number of instances that would be assigned to a bucket if generalization were applied. It is calculated by dividing the dataset size k by the number of buckets g . The salt applied to a value x is then computed by multiplying a standard normal random variable with zero mean and unit variance $\mathcal{N}(0, 1)$, by the distance between x and its i -th closest value x_i . The formula is defined as follows:

$$salt_x = \mathcal{N}(0, 1) * \text{distance}(x, x_i) \quad (2)$$

To generate an anonymised value, the salt is added to the original value. If the resulting value falls outside the range of values of the original dataset, the salt is instead subtracted from the original value.

Running Example. Figure 2 illustrates an anonymisation example with the data of an REC. The input dataset contains the attributes entryDate, longitude, latitude, and energyProduction. anonymisation is applied to the first three attributes, while energy production remains unchanged. As a result, the energy data is no longer linkable to a specific individual.

The anonymisation service is invoked with two primary components: the input data to be anonymised and a reference URL pointing to the anonymisation configuration. The configuration defines the anonymisation methods applied to each attribute. Despite the anonymisation, data mining remains feasible; for example, the dataset can still be used to analyze energy production patterns across different geographic locations. A sample dataset and detailed instructions for reproducing the example are available in the public GitHub repository.

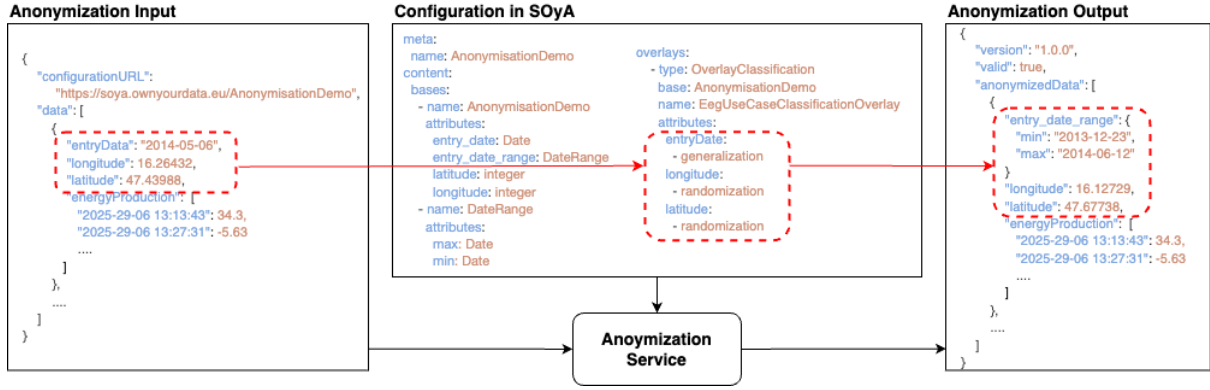


Figure 2: An example anonymisation on REC data

4. Feasibility Evaluation

To assess the data security of the service, a series of tests was conducted using synthetic test datasets. These datasets were generated to replicate the personal data of members of a Renewable Energy Communities (RECs) in the Austrian state Burgenland. Both generalization and randomization techniques were applied to these datasets, with varying sample sizes ranging from 100 to 10,000. This section explains the synthetic test data generation process and outlines the method used to calculate a benchmark value. Additionally, we present the test results and evaluate the compliance with existing regulatory standards.

Test Data Generation. Due to the limited availability of real personal data, synthetic data was generated. The created data sets include three attributes: longitude, latitude, and the date of an individual’s registration in the energy community. The geographic coordinates were selected to approximate locations within the Austrian state of Burgenland. Registration dates were randomly generated within the range from 2005 to 2025, with a higher probability assigned to more recent years, reflecting the assumption that membership registrations have increased in recent years.

Benchmark Calculation. To evaluate the anonymisation service with respect to data confidentiality, k-anonymity is employed as the metric for generalization-based techniques. For randomization-based anonymisation, a modified approach is required, that assesses the similarity between original and anonymised instances to establish an analogous notion of k-anonymity.

Similarity in this context is defined as follows: an original instance $A_{original}$ is considered similar to an anonymised instance $B_{anonymised}$ if $A \neq B$ and all randomized attribute values in $B_{anonymised}$ lie within an acceptable range of deviation from the corresponding values in $A_{original}$. To determine this acceptable range, the distribution of distances between original values and their anonymised counterparts is analyzed. Since the randomization process adds normally distributed noise (salt), a threshold of 2σ (twice the standard deviation) is used as the similarity criterion. If the distance between an anonymised value and the corresponding original value is below this threshold, the two are considered similar for that attribute.

An anonymised instance is considered similar to an original instance if it satisfies the similarity condition for all attributes. The k-anonymity of the anonymised dataset is then defined as the minimum number of original instances that are similar to any given anonymised instance.

Evaluation Result. The test results are presented in Table 1. Both randomization and generalization were evaluated using sample sizes of 100, 1,000, and 10,000. Each configuration was executed ten times with different synthetic datasets. For each setup, the median and minimum k-anonymity values across the ten runs are reported. In addition, the table includes the number of buckets (groups) used for each sample size, calculated according to the method described in Section 3.2.

	Size: 100 Groups: 2		Size: 1.000 Groups: 4		Size: 10.000 Groups: 8	
	Median	Minimum	Median	Minimum	Median	Minimum
Generalization	9.0	6	7.0	5	7.0	6
Randomization	24.5	18	17.5	8	16.5	12

Table 1

Median and minimum k-anonymity for Generalization and Randomization across different dataset sizes

Our experiment confirms the anonymised data complies with GDPR requirements, ensuring that Data Intermediaries can share information in a legally compliant manner while maintaining the trust of data subjects and ecosystem stakeholders. Across all tested sample sizes and techniques, no instance was uniquely identifiable; each anonymised instance satisfied a minimum k-anonymity of 5. However, the analytical utility varies with dataset size and the number of groups used. Figure 3 illustrates spatial groupings by dataset size for generalization, e.g., a sample size of 100 results in only four spatial groups, limiting analytical potential.

Overall, randomization achieved greater anonymisation than generalization. The selected group count represents a balanced trade-off between data confidentiality and analytical utility. For stronger privacy guarantees, fewer groups can be used, albeit reducing analytical detail.

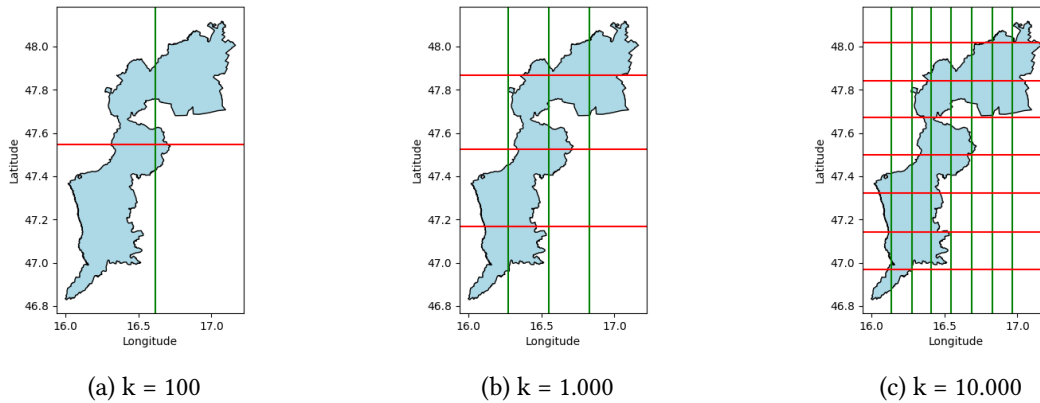


Figure 3: Spatial Groups according to Data size

5. Summary and Future Work

This paper presents a configurable and extensible anonymisation service for semantically annotated datasets, addressing key compliance aspects of GDPR and the EU Data Governance Act (DGA). Through integration into Renewable Energy Communities (RECs), we demonstrated automated, rule-based anonymisation pipelines that effectively mitigate re-identification risks while preserving data utility. The novelty lies in combining semantic annotation via SOyA with flexible anonymisation techniques, notably generalization and randomization, to ensure broad regulatory alignment and ease of adoption.

Our service supports diverse deployment scenarios, from public SaaS platforms to self-hosted environments, catering to various data sovereignty requirements. Future developments will introduce additional anonymisation techniques, enhance anonymisation outputs with Key Performance Indicators (KPIs) and risk assessments, and improve the user interface based on systematic feedback to enhance accessibility and usability. Additionally, assessing operational impacts on data intermediaries, including scalability and performance, will guide practical implementations within real-world data governance infrastructures.

Acknowledgments

This work was conducted as part of the USEFLEDs project. This project has received funding in the program “Datenökosysteme für die Energiewende” by the Federal Ministry for Climate Action, Environment, Energy, Mobility, Innovation and Technology (BMK) under grant number 905128.

Declaration on Generative AI

During the preparation of this work, the author(s) used X-GPT-4 in order to: Grammar and spelling check. After using these tool(s)/service(s), the author(s) reviewed and edited the content as needed and take(s) full responsibility for the publication’s content.

References

- [1] Bundesministers für Wirtschaft, Familie und Jugend Österreich, Intelligente messgeräte-einführungsverordnung – ime-vo, 2012. <https://www.ris.bka.gv.at/eli/bgbl/II/2012/138>.
- [2] Parlament Österreich, Elektrizitätswirtschafts- und -organisationsgesetz (eiwog 2010), 2010. <https://www.parlament.gv.at/gegenstand/XXIV/I/2067>.
- [3] EU, General data protection regulation (GDPR) – legal text, 2016. URL: <https://gdpr-info.eu/>.
- [4] M. T. Dimakopoulos, Dimitris Tsitsigkos {and} Nikolaos, Amnesia anonymization tool - data anonymization made easy, 2025. URL: <https://amnesia.openaire.eu/index.html>.
- [5] ARX – data anonymization tool – a comprehensive software for privacy-preserving microdata publishing, 2025. URL: <https://arx.deidentifier.org/>.
- [6] P. Ponnaganti, R. Sinha, J. Pillai, B. Bak-Jensen, Flexibility provisions through local energy communities: A review, *Next Energy* 1 (2023) 100022. URL: <https://doi.org/10.1016/j.nxener.2023.100022>. doi:10.1016/j.nxener.2023.100022.
- [7] S. Zhang, S. Maharjan, L. A. Bygrave, S. Yu, Data sharing, privacy and security considerations in the energy sector: A review from technical landscape to regulatory specifications, *arXiv preprint arXiv:2503.03539*, 2025. URL: <https://arxiv.org/abs/2503.03539v1>.
- [8] European Parliament and Council, Regulation (eu) 2022/868 of the european parliament and of the council on european data governance (data governance act), *Official Journal of the European Union*, L 152, Article 12(e), 2022. URL: <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32022R0868>.
- [9] F. J. Ekaputra, C. Fabianek, G. Unterholzer, E. Gringinger, The semantic overlay architecture for data interoperability and exchange, in: 2023 IEEE International Conference on Data and Software Engineering (ICoDSE), IEEE, 2023, pp. 232–237.
- [10] S. Murthy, A. Abu Bakar, F. Abdul Rahim, R. Ramli, A comparative study of data anonymization techniques, in: 2019 IEEE 5th Intl Conference on Big Data Security on Cloud (BigDataSecurity), IEEE Intl Conference on High Performance and Smart Computing, (HPSC) and IEEE Intl Conference on Intelligent Data and Security (IDS), 2019, pp. 306–309. URL: <https://ieeexplore.ieee.org/abstract/document/8819477>. doi:10.1109/BigDataSecurity-HPSC-IDS.2019.00063.
- [11] A. Majeed, S. Lee, Anonymization techniques for privacy preserving data publishing: A comprehensive survey 9 (2021) 8512–8545. URL: <https://ieeexplore.ieee.org/abstract/document/9298747>. doi:10.1109/ACCESS.2020.3045700.
- [12] B. C. Fung, K. Wang, A. W.-C. Fu, P. S. Yu, *Introduction to privacy-preserving data publishing: Concepts and techniques*, Chapman and Hall/CRC, 2010.
- [13] I. E. Olatunji, J. Rauch, M. Katzensteiner, M. Khosla, A review of anonymization for healthcare data 12 (2024) 538–555. URL: <https://www.liebertpub.com/doi/abs/10.1089/big.2021.0169>. doi:10.1089/big.2021.0169, publisher: Mary Ann Liebert, Inc., publishers.
- [14] R. Bayardo, R. Agrawal, Data privacy through optimal k-anonymization, in: 21st International

Conference on Data Engineering (ICDE'05), 2005, pp. 217–228. URL: <https://ieeexplore.ieee.org/abstract/document/1410124>. doi:10.1109/ICDE.2005.42, ISSN: 2375-026X.

- [15] P. Samarati, L. Sweeney, Generalizing data to provide anonymity when disclosing information, in: PODS, volume 98, 1998, pp. 10–1145.
- [16] Z. Teng, W. Du, Comparisons of k-anonymization and randomization schemes under linking attacks, in: Sixth International Conference on Data Mining (ICDM'06), IEEE, 2006, pp. 1091–1096.
- [17] H. Li, L. Yu, W. He, The impact of gdpr on global technology development, 2019.
- [18] R. N. Zaeem, K. S. Barber, The effect of the gdpr on privacy policies: Recent progress and future promise, ACM Transactions on Management Information Systems (TMIS) 12 (2020) 1–20.
- [19] J. F. Marques, J. Bernardino, Analysis of data anonymization techniques., in: KEOD, 2020, pp. 235–241.
- [20] Y. Jiang, L. Mosquera, B. Jiang, L. Kong, K. El Emam, Measuring re-identification risk using a synthetic estimator to enable data sharing, PLoS One 17 (2022) e0269097.