

OPEN
ARTICLE

An Open Source Python Library for Anonymizing Sensitive Data

Judith Sáinz-Pardo Díaz & Álvaro López García

Open science is a fundamental pillar to promote scientific progress and collaboration, based on the principles of open data, open source and open access. However, the requirements for publishing and sharing open data are in many cases difficult to meet in compliance with strict data protection regulations. Consequently, researchers need to rely on proven methods that allow them to anonymize their data without sharing it with third parties. To this end, this paper presents the implementation of a Python library for the anonymization of sensitive tabular data. This framework provides users with a wide range of anonymization methods that can be applied on the given dataset, including the set of identifiers, quasi-identifiers, generalization hierarchies and allowed level of suppression, along with the sensitive attribute and the level of anonymity required. The library has been implemented following best practices for integration and continuous development, as well as the use of workflows to test code coverage based on unit and functional tests.

Introduction

The development of data driven applications is a growing field that requires large volumes of data for its successful evolution and performance. In some cases, these models are powered by data that may contain information about individuals, and therefore it is critical to focus on the privacy of such data. Numerous studies have collected information on the possible bias that algorithms and in particular Artificial Intelligence (AI) models can potentially exhibit based on the data used during training^{1,2}. When it comes to data containing personal information, performing proper pre-processing and curation steps are therefore key for two reasons: (1) to protect the privacy of the individuals associated with the data (2) to mitigate the bias that data may contain.

Substantial legislative work has been done to advance these technologies with attention to the security of personal data of individuals. Specifically, in Europe the General Data Protection Regulation (GDPR) was put into effect on May 25, 2018, and it aims to protect natural people regarding the processing of their personal data and the free movement of such information within the European Union and the European Economic Area (EEA). The GDPR establishes within its article 5 principles regarding data protection, including that data must be processed in such a way as to ensure adequate security of personal data. Likewise, consideration number 26 of the same states that this regulation does not affect the processing of anonymous information, including for statistical or research purposes, considering as anonymous information that which does not relate to an identified or identifiable natural person³. This motivates the need to promote software applications for use by the research community in order to carry out an adequate process of anonymization of research data for enhancing reproducibility and open science.

Recently, with the rise of technologies based on artificial intelligence and the risks they may entail (among others in relation to security and privacy of the individual) the European Commission proposed on April 21, 2021 to launch the EU AI Act⁴, which was approved on March 13, 2024, composing the first regulation on artificial intelligence. Specifically, this regulation defines different risks related to AI and different levels for them: Unacceptable, high, limited or minimal. As an example, AI systems identified as high-risk include AI technology employed in administration of justice and democratic processes. These data-driven models require that the data have been properly curated and subjected to the necessary anonymization processes to reduce the risk of extracting sensitive information from them.

Going back to the legislation on privacy and data security, in the USA the Privacy Act⁵ (1974) establishes a code of fair information practices that governs the collection, maintenance, use, and dissemination of information about individuals which is maintained by federal agencies. More specifically in the field of health, HIPAA⁶

Instituto de Física de Cantabria (IFCA), CSIC-UC Avda. los Castros s/n, 39005, Santander, Spain. e-mail: sainzpardo@ifca.unican.es

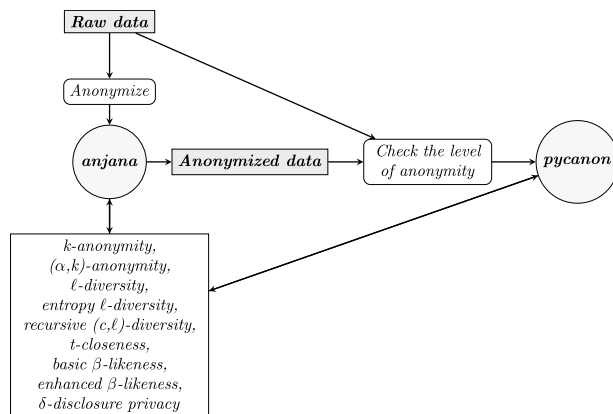


Fig. 1 Data anonymization and checking the level of anonymity: connection between *pycanon* and *anjana*.

is a suite of federal regulatory standards that outline the use and disclosure of protected health information in the United States.

In view of the above, it is essential to have tools to ensure privacy through data anonymization, specifically in case of sensitive data or associated with individuals. This is even critical if there is a potential use of such data through AI techniques as the basis of data-driven applications, like machine or deep learning (ML/DL) models. To this end, in this article we present *anjana*⁷, an open source Python library that implements different anonymization techniques that can be applied on tabular data without the need to externalize, upload to different service providers or share them with third parties.

This paper is structured as follows: in the Methods Section we explore the state of the art and motivate the development of this open source tool. We also present the anonymity techniques implemented, their connection and some important definitions. In the Results Section we explain the functionality of the library and some examples, together with important remarks on software development. Finally, the Discussion Section provides insights on the management of multiple sensitive attributes, how to create the hierarchies and get the anonymity transformation applied, a comparative analysis on the efficiency of the techniques implemented and the incorporation of these anonymity methods in data science workflows. Finally, the study concludes with some directions for future work.

Methods

Regarding tabular data anonymization, numerous theoretical papers review the implementation and definition of the techniques that will be implemented in this library and that are reported in the following subsection. In particular, each of these methods each of them focusing on different issues or prevents different attacks, as presented in the diagram of Fig. 2.

Specifically, in 2002 a definition of *k-anonymity* was given by L. Sweeney⁸ and in 2006 (α, k) -anonymity was introduced by R. C. Wong *et al.* in order to “protect the association of individuals to sensitive information”⁹. In 2007 A. Machanavajhala *et al.* presented *l-diversity* as a novel approach to protect against some of the attacks to which *k-anonymity* proved to be sensitive¹⁰. In that paper the concepts of *entropy l-diversity* and *recursive (c, l)-diversity* are also introduced. In the same year (2007) N. Li *et al.* presented *t-closeness* as a privacy technique beyond *l-diversity* and *k-anonymity*¹¹. In 2008 J. Brickell and V. Shmatikov¹² already define the concept of δ -disclosure privacy and they measure the trade-off between privacy and utility using the adult dataset (which is also used for testing purposes in the current work). Finally, in 2012 J. Cao and P. Karras proposed the use β -likeness as a “robust privacy model for microdata anonymization”, introducing the definitions of both *basic β -likeness* and *enhanced β -likeness*¹³.

One of the main concerns arising nowadays in relation to data privacy stems from the important role played by data as the basis for artificial intelligence (AI), machine and deep learning (ML/DL) models¹⁴. These kind of models can perform classification, clustering, inference, pattern recognition or anomaly detection tasks based on data that in many cases are subject to privacy restrictions. In the same direction, the development of data-driven AI models is extensively supported in Python by numerous frameworks for both ML and DL, as is the case for *scikit-learn*¹⁵, *keras*¹⁶, *TensorFlow*¹⁷ and *PyTorch*¹⁸ among others. However, in relation to the prior anonymization of the data from which the models available in these libraries are based, the computing facilities to apply the above techniques using open source Python frameworks are limited. We can find some projects on GitHub that implement some of the techniques mentioned above, especially in the case of *k-anonymity*, for which we can find different algorithms to implement it, such as *mondrian*, *datafly* or *incognito*^{19,20}.

As for Python libraries concerning data anonymization, we can highlight *ANONYPY* (<https://github.com/glassionion1/anonymypy>) which implements the *mondrian* algorithm supporting *k-anonymity*, *l-diversity* and *t-closeness*. On the other hand, the *anonym* (<https://gitlab.com/datainnovatielab/public/anonym>) library is designed to anonymize dataframes and it operates by replacing real data with fake ones, while maintaining the structure and format of the original data. *dicognito* (<https://github.com/blairconrad/dicognito>) is a Python library and command line interface (CLI) that anonymizes medical records given in DICOM format by

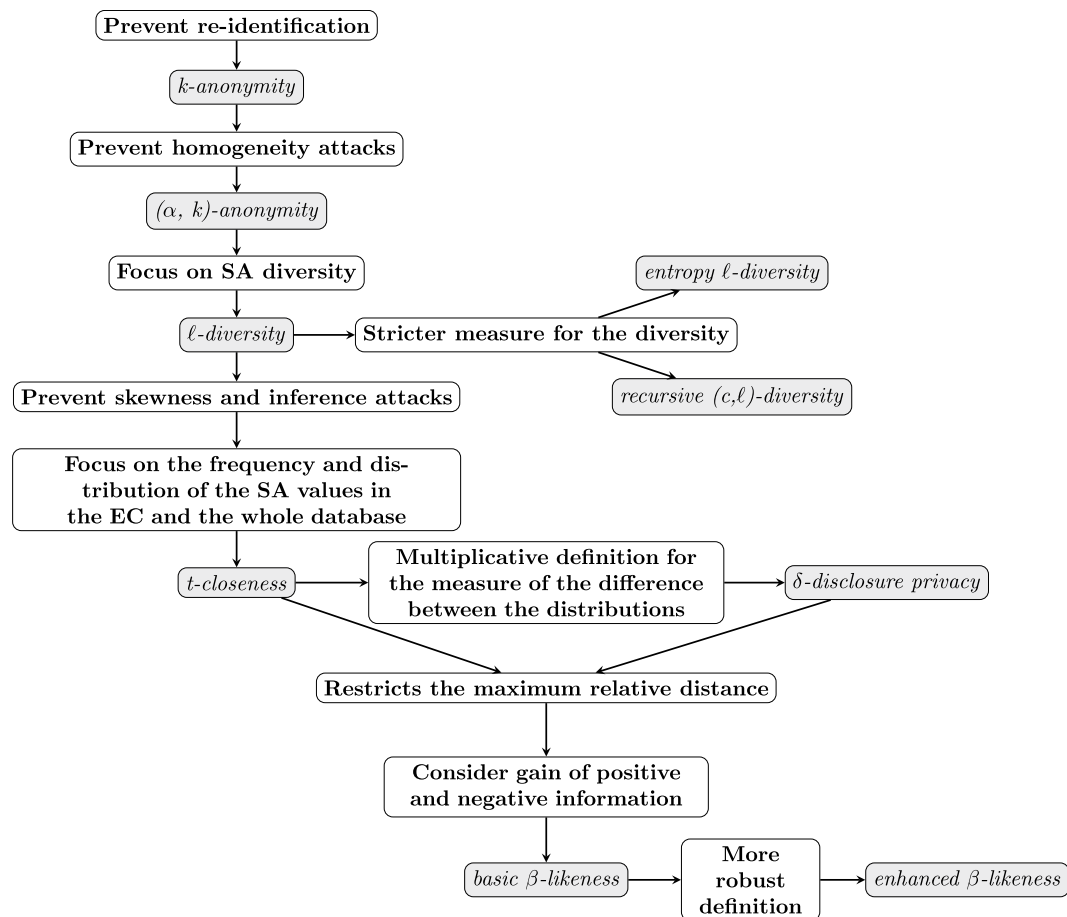


Fig. 2 Workflow: select the anonymity technique to be applied depending on the privacy objective.

removing or substituting different fields. The `privapy` (<https://github.com/vincentmin/privapy>) library includes methods for text anonymization by removing identifying or quasi-identifying words, and also provides methods for image cleanup, such as blackout or pixelate. Besides, `python-anonymity` is a master thesis project in which k -anonymity was implemented by means of both `datafly` and `incognito`, and some preliminary versions for ℓ -diversity and t -closeness were provided (this project is read only and archived, since it was created for academic purposes²¹).

In terms of open source tools for anonymizing sensitive data, ARX²², written in Java, stands out. It is a comprehensive open source software for sensitive data anonymization which supports several privacy models (such as the ones exposed in this work), together with risk and quality models. In addition, ARXaaS (<https://github.com/navikt/arxaas>) claims to be an “anonymization as a service” project built on top of the ARX library, which uses HTTP by default. Additionally, Amnesia (<https://amnesia.openaire.eu/>) is a framework deployed in Java for data anonymization that, in addition to k -anonymity implements km -anonymity, which can provide better assessments in case of high dimensional data. It can be installed locally but also includes an online version that allows to test its functionality. Other software solutions are available, as the ones reviewed in²³, such as μ -ARGUS, τ -ARGUS, PRIVAaaS or `g9 Anonymize`. In view of the current review of free software for anonymizing sensitive data carried out in the context of this work, ARX excels in terms of the multiple functionalities available, which can be used through an interactive interface.

Finally, it is important to mention `pycanon`²³, which will be used as the basis for the library proposed in this paper, and which is an open source Python library that allows to check the anonymity level of a tabular dataset according to the techniques mentioned at the beginning of this section. `Pycanon` and `anjana` are implemented as separate tools in order to keep the different functionalities independently developed. The main difference between `pycanon` and `anjana` is that the former can be used on a known dataset to find out its anonymity level and thus be aware of the possible risks derived from its sharing or publication, and can be employed on data that has already been anonymized, without the need to have knowledge of the raw data. However, `anjana` can be used by users who want to anonymize their data from scratch with different techniques, evaluating which are the most appropriate ones in each case, allowing to carry out the whole anonymization process as soon as the raw data is available. By keeping these two tools separated we can provide the dependencies for each tool to the minimum that is needed, keeping their installation and maintenance simpler. The connection between these two libraries is shown in Fig. 1.

Technique	Function and input
k -anonymity	$k_anonymity(data, id, qi, k, supp, hier)$
(α, k) -anonymity	$alpha_k_anonymity(data, id, qi, sa, \alpha, k, t, supp, hier)$
ℓ -diversity	$\ell_diversity(data, id, qi, sa, k, \ell, supp, hier)$
entropy ℓ -diversity	$entropy_l_diversity(data, id, qi, sa, k, \ell, supp, hier)$
recursive (c, ℓ) -diversity	$recursive_c_l_diversity(data, id, qi, sa, k, c, \ell, supp, hier)$
t -closeness	$t_closeness(data, id, qi, sa, k, t, supp, hier)$
δ -disclosure privacy	$delta_disclosure(data, id, qi, sa, k, \delta, supp, hier)$
basic β -likeness	$basic_beta_likeness(data, id, qi, sa, k, \beta, supp, hier)$
enhanced β -likeness	$enhanced_beta_likeness(data, id, qi, sa, k, \beta, supp, hier)$

Table 1. *Anjana* main functions to apply the different anonymity techniques. Note that id is the list of identifiers, qi that of the quasi-identifiers, sa is the sensitive attribute, $supp$ the suppression level, and $hier$ the dictionary with the hierarchies. In all the cases k is the desired value for k -anonymity.

age	education	marital-status	occupation	sex	native-country	race	salary-class
23	Some-college	Never-married	Protective-serv	Male	United-States	White	<=50 K
60	HS-grad	Widowed	Other-service	Female	United-States	White	<=50 K
60	Assoc-voc	Widowed	Other-service	Female	United-States	Black	<=50 K
48	Assoc-acdm	Married-civ-spouse	Prof-specialty	Female	United-States	White	>50 K
35	Bachelors	Separated	Craft-repair	Male	United-States	White	<=50 K
53	Some-college	Widowed	Adm-clerical	Female	United-States	Black	<=50 K
28	Some-college	Married-civ-spouse	Protective-serv	Male	United-States	White	<=50 K
19	HS-grad	Never-married	Handlers-cleaners	Male	United-States	White	<=50 K
23	10th	Never-married	Other-service	Female	United-States	White	<=50 K

Table 2. Sample extracted from the adult dataset with a selected subset of columns.

Extensive development has taken place in recent years in terms of frameworks for AI/ML/DL data-driven model development built on top of Python. In addition, this review makes it evident that it is necessary to provide users, and in particular researchers working with sensitive tabular data, with a tool that can be easily installed locally and used in an intuitive way to have at their disposal a wide range of anonymization techniques, covering different types of attacks, and that this is done in the form of an open source Python library. This motivates the implementation of the *anjana* open source Python library, which aims to provide a solution to this need.

Anonymity techniques. The objective of the open source framework presented in this work is to provide researchers and general users which work with tabular data containing sensitive information, with an easy-to-use Python based tool that allows them to anonymize them according to different techniques. Then, in this section we will introduce the different anonymity techniques available in *anjana*.

First of all, let's introduce some basic definitions in relation to the attributes or columns of tabular databases that must also be entered as input to the anonymization techniques implemented.

Definition 1. Given a tabular database, each column composes an attribute in relation to the data, and they can be classified as follows:

Identifiers (ID), which contain information that allow to unequivocally identify a user (e.g. name, ID card number, national insurance number, etc.). This information must be removed in the anonymized version of the database.

Quasi-identifiers (QI), these are the information that although by themselves do not allow to uniquely identify a user, they do allow to identify him/her through a combination of a series of these (examples: age, city of birth, level of studies, zip code, etc.).

Sensitive attributes (SA), which are the information about the individuals in the database that we want to protect and that should not be inferred by an attacker (example: diseases, salary class, medical records, police records, etc.).

Insensitive attributes (IN), which do not require special processing and can be left in the same form (example: randomly generated IDs).

Definition 2. Given a database \mathcal{D} and a set of quasi-identifiers Q , we define an **equivalence class (EC)** as a subset of rows that are identical regarding the given set of quasi-identifiers Q ²³.

Formally, let $\mathcal{S} \subseteq \mathcal{D}$ be a subset of rows of \mathcal{D} , $|Q|$ the number of quasi-identifiers in \mathcal{D} , S form an equivalence class regarding Q (EC_Q) if for every pair of entries $x, y \in \mathcal{S}$ and for every quasi-identifier $q_i, i \in \{1, \dots, |Q|\}$, $x[q_i] = y[q_i] \forall i \in \{1, \dots, |Q|\}$ and in addition $\nexists z \in \mathcal{D} \setminus \mathcal{S}$ verifying $z[q_i] = x[q_i] \forall i \in \{1, \dots, |Q|\}$ and $x \in \mathcal{S}$.

Regarding the anonymity techniques implemented, in the version 1.0.0 of the library, the definition of one sensitive attribute is supported. In this line, *anjana* provides users with the following nine anonymity tools defined as stated in²³ based on the assumption of only one SA: *k-anonymity*, (α, k) -*anonymity*, *l-diversity*, *entropy l-diversity*, *recursive (c, l)-diversity*, *t-closeness*, *delta-disclosure privacy*, *basic beta-likeness* and *enhanced beta-likeness*. As stated in²³, these techniques are complementary as they prevent from different types of attacks (see also²²). We can see the summary of the main attacks prevented in Table 2 from²³. However, each technique can prevent more attacks than the ones marked in that table, although some are more suitable than others for certain attacks. The workflow that can be followed to choose the appropriate anonymization technique is shown in the diagram presented in Fig. 2. It is important to note that although these functions have been defined for a single sensitive attribute, it is possible to apply them in case of multiple sensitive attributes as will be explained in the Methods section, following an iterative process depending on the desired paradigm: *harmonization of the quasi-identifiers* or *quasi-identifiers update*²³.

The idea of these techniques is to remove the identifiers and apply recursive transformations on the quasi-identifiers that allow them to be generalized, making complicated to extract information about a particular individual. To apply these transformations, a set of hierarchies must be created for each QI. Usually, if we allow the suppression of a quasi-identifier, this will be the last level of hierarchy.

Definition 3. Let Q be the set of quasi-identifiers of a database \mathcal{D} , we define a **hierarchy** h_i over a certain quasi-identifier q_i ($i \in \{1, \dots, |Q|\}$) as a mapping from $D(q_i)$ to $D(q_i^{(1)})$, $h_i: D(q_i) \rightarrow D(q_i^{(1)})$, with $D(q_i)$ the initial set of values of q_i in D , and $D(q_i^{(1)})$ being the new set of new values of q_i once generalized (first transformation). Thus, let $\{x_1, \dots, x_m\} \in D(q_i)$ be the set of unique values taken by the attribute q_i in \mathcal{D} , the function h_i assigns $h_i(x_i) = y_j \forall i \in \{1, \dots, m\}$. Note that different values of x_i can take the same value y_j $j \in \{1, \dots, p\}$, $p \leq m$. Also, note that h_i is supposed to operate upon all the values of $D(q_i)$.

Thus, we can define n mapping $h_i \forall i \in \{1, \dots, n\}$ to generalize the attribute q_i as much as necessary by applying a new hierarchy over each state of the database: $h_1: D(q_i) \rightarrow D(q_i^{(1)})$, $h_2: D(q_i^{(1)}) \rightarrow D(q_i^{(2)})$, \dots , $h_n: D(q_i^{(n-1)}) \rightarrow D(q_i^{(n)})$. Then, for the quasi-identifier q_i we define a set of hierarchies $H_{q_i} = \{h_1, \dots, h_n\}$, with n the number of possible transformations for q_i .

If we perform only the first transformation, the quasi-identifier q_i will have been anonymized with hierarchy level 1, while if we apply n transformations (until h_n), we will have a transformation level of n for q_i .

Results

Functionality and use examples. As already stated, *anjana* provides user with the most common anonymization techniques for their application on tabular databases. These methods have been performed according to the definition given in²³, and the *pycanon* Python library has been used as an auxiliary tool to check that the level of anonymization required by the user is verified.

Except for *k-anonymity*, for which the sensitive attribute is not introduced as it focuses only on the quasi-identifiers, the input of the functions that apply the available anonymity techniques is as follows:

- **Data:** a pandas dataframe with the data to be anonymized. Each column can contain: identifiers, quasi-identifiers or sensitive attributes.
- **Identifiers:** a list of strings containing the names of the identifiers in the dataframe, in order to suppress them (substituting by '*').
- **Quasi-identifiers:** a list of strings, containing the names of the quasi-identifiers in the dataframe.
- **Sensitive-attributes:** a string with the name of the sensitive attribute in the dataset (only one). This parameter must be introduced in case of applying other techniques than *k-anonymity*.
- **Privacy level:** an int or float (depending the function applied) with level of anonymity to be applied, e.g. k (int) for *k-anonymity*, α (float) and k (int) for (α, k) -*anonymity*, ℓ (int) for *l-diversity* or *entropy l-diversity*, c (int) and ℓ (int) for *recursive (c, l)-diversity*, t (float) for *t-closeness*, β (float) for *basic or enhanced beta-likeness* or δ (float) for *delta-disclosure privacy*. Note that in all the cases a value of k for *k-anonymity* must be introduced, following the same structure as ARX.
- **Suppression level:** a float between 0 and 100 with the maximum level of record suppression allowed.
- **Hierarchies:** dictionary containing the hierarchies and the generalization level for each quasi-identifier to be generalized. The dictionary contains one dictionary for each quasi-identifier with the hierarchies and the levels.

In Table 1, the different available techniques are shown, together with their name in the library and the input expected for applying them.

In order to test the functionality of the library, two well known datasets have been used: the hospital dataset (defined as given in Table 4) and the adult dataset²⁴. The simplest one was used for the first proofs of concept due to its simplicity (hospital dataset), while the second one was used to perform multiple tests in a real scenario with a dataset with more than 30000 rows (adult dataset). Specifically, with respect to the adult dataset, in Table 2 we show an extraction of ten rows and eight columns that have been used as ID, QI and SA during the testing phase:

Given the dataset presented in Table 2, we can consider *race* as identifier in order to suppress it and the *salary-class* as sensitive attribute (since it is the attribute that we want to prevent an attacker from extracting). In

addition, the other columns (*age*, *education*, *marital-status*, *occupation*, *sex* and *native-country*) will be the quasi-identifiers. Note that this data extraction does not contain identifying information. Therefore, in order to apply the anonymity techniques previously exposed, we will need to introduce the list of QI, ID and the SA, together with the hierarchies to be applied to the quasi-identifiers, the maximum level of records suppression allowed and the anonymity level desired. As an illustrative example, let's start by applying *k-anonymity* for $k=10$, *l-diversity* for $\ell=2$ and *t-closeness* for $t=0.5$ using *anjana*. The steps to follow are detailed in the code given in Code 3, once the hierarchies presented in²⁵ have been imported. Regarding creating and defining the hierarchies more details will be given in the [Method](#) Section.

```
import pandas as pd
import anjana
from anjana.anonymity import k_anonymity, l_diversity, t_closeness

# Read and process the data
data = pd.read_csv("adult.csv")
data.columns = data.columns.str.strip()
cols = [
    "workclass",
    "education",
    "marital-status",
    "occupation",
    "sex",
    "native-country",
]
for col in cols:
    data[col] = data[col].str.strip()

# Define the identifiers, quasi-identifiers and the sensitive attribute
quasi_ident = [
    "age",
    "education",
    "marital-status",
    "occupation",
    "sex",
    "native-country",
]
ident = ["race"]
sens_att = "salary-class"

# Select the desired level of k, l and t
k = 10
l = 2
t = 0.5

# Select the suppression limit allowed
supp_level = 50

# Import the hierarchies for each quasi-identifier. Define a dictionary containing them
hierarchies = {
    "age": dict(pd.read_csv("hierarchies/age.csv", header=None)),
    "education": dict(pd.read_csv("hierarchies/education.csv", header=None)),
    "marital-status": dict(pd.read_csv("hierarchies/marital.csv", header=None)),
    "occupation": dict(pd.read_csv("hierarchies/occupation.csv", header=None)),
    "sex": dict(pd.read_csv("hierarchies/sex.csv", header=None)),
    "native-country": dict(pd.read_csv("hierarchies/country.csv", header=None)),
}

# Apply k-anonymity:
data_kanon = k_anonymity(
    data, ident, quasi_ident, k, supp_level, hierarchies
)

# Apply l-diversity once k-anonymity is implemented:
data_kanon_ldiv = l_diversity(
    data_kanon, ident, quasi_ident, sens_att, k, l, supp_level, hierarchies
)

# Apply t-closeness once k-anonymity and l-diversity are implemented:
data_kanon_ldiv_tclos = t_closeness(
    data_kanon_ldiv, ident, quasi_ident, sens_att, k, t, supp_level, hierarchies
)
```

Example Code 1. Example: applying *k-anonymity*, *l-diversity* and *t-closeness* for the adult dataset ($k=10$, $\ell=2$ and $t=0.5$).

Technique	Anonymity values applied	Anonymity values calculated	Suppressed records (%)
<i>k</i> -anonymity	$k = 10, \text{supp_level} = 50\%$	$k = 10$	43.71%
(α, k) -anonymity	$k = 10, \alpha = 0.8, \text{supp_level} = 100\%$	$k = 10, \alpha = 0.8$	87.99%
ℓ -diversity	$k = 10, \ell = 2, \text{supp_level} = 50\%$	$k = 72, \ell = 2$	43.71%
<i>t</i> -closeness	$k = 10, t = 0.5, \text{supp_level} = 50\%$	$k = 72, t = 0.4737$	43.71%
δ -disclosure privacy	$k = 10, \delta = 3, \text{supp_level} = 50\%$	$k = 392, \delta = 2.159$	43.71%
Basic β -likeness	$k = 10, \beta = 0.5, \text{supp_level} = 100\%$	$k = 2098, \beta = 0.4178$	72.74%
Enhanced β -likeness	$k = 10, \beta = 0.5, \text{supp_level} = 100\%$	$k = 2098, \beta = 0.4178$	72.74%

Table 3. Anonymity techniques applied to the adult dataset: values for the privacy suppression levels applied using *anjana*, real anonymity values calculated using *pyCANON* and percentage of suppressed records.

name	age	gender	city	religion	disease
Ramsha	29	Female	Tamil Nadu	Hindu	Cancer
Gabu	24	Male	Tamil Nadu	Hindu	Cancer
Sabu	23	Male	Tamil Nadu	Hindu	Cancer
Jonas	22	Male	Tamil Nadu	Hindu	Cancer
Yadu	24	Female	Kerala	Hindu	Viral infection
Salima	28	Female	Tamil Nadu	Muslim	TB
Sunny	27	Male	Karnataka	Parsi	No illness
Joan	24	Female	Kerala	Christian	Heart-related
Bahuksana	23	Male	Karnataka	Buddhist	TB
Rambha	19	Male	Kerala	Hindu	Cancer
Kishor	29	Male	Karnataka	Hindu	Heart-related
Johnson	17	Male	Kerala	Christian	Heart-related
John	19	Male	Kerala	Christian	Viral infection

Table 4. Example: data from the *hospital_extended.csv* file.

The result in this case will be an anonymized dataset verifying *k*-anonymity for $k = 10$, $\ell = 2$ for ℓ -diversity and $t = 0.5$ for *t*-closeness, which can be stored as a *csv* file or other format for later use. Note that this dataset contains 32561 rows and the application of this technique using *anjana* provides the anonymized dataset in less than one second (carrying it out in a machine with 12th Gen Intel(R) Core(TM) i5-1, 16 GB RAM and running under Ubuntu 22.04.4 LTS). Specifically, applying *pycanon*²³ we get the following values once anonymized with the three techniques: $k = 72$, $\ell = 2$ and $t = 0.47370$.

During the functionality test phase and examples, all the techniques of the library have been applied to this dataset with the values of Table 3. The actual value calculated using the *pycanon* library is shown in that table. It should be noted that for *entropy* ℓ -diversity and *recursive* (c, ℓ) -diversity it has not been possible to anonymize it given the hierarchies introduced, so they are not shown. In addition, the percentage of records eliminated in each case is shown.

Software development. *Dependencies management and installation.* To manage the dependencies of the library and its installation, *poetry* (<https://github.com/python-poetry/poetry>) has been used. Specifically, a *pyproject.toml* file has been created which, in addition to these dependencies, contains the metadata associated with the library (version, authors, license, classification, etc). The dependencies are distinguished between those that must be installed together with the library for its correct operation (described in `[tool.poetry.dependencies]`), and the dependencies necessary to run the tests, create the documentation, check the code style, etc.

CI/CD and publish pipelines. In the current version, the development of the library is done on GitHub and GitHub actions have been used for creating a CI/CD pipeline and a workflow for publishing the current release in PyPI. The CI/CD pipeline (named *cicd.yml*) is executed with the following stages: lint, test and build. The lint stage is responsible for checking the code style and syntax errors. In the test stage we check the correct execution in Python 3.9 3.10 3.11 and 3.12. Finally, the build stage checks that the package is built correctly and the dependencies are properly installed.

Unit and functional tests. Regarding the validation of the correct functionality of the library, two types of tests have been carried out: unit and functional tests. In the first ones, we have tested the boundary cases, the correct management of exceptions, the correct typing of the function inputs, and the extreme cases. In the functionality tests, we have verified that the implemented anonymity functions work correctly, returning the expected results based on the desired level of anonymity applied. For the latter, the *pycanon* library has been used as auxiliary library, since it already has its own tests that ensure the proper performance of the same. To implement these

name	age	gender	city	religion	disease
*	[20, 30)	Female	Tamil Nadu	Hindu	Cancer
*	[20, 30)	Male	Tamil Nadu	Hindu	Cancer
*	[20, 30)	Male	Tamil Nadu	Hindu	Cancer
*	[20, 30)	Male	Tamil Nadu	Hindu	Cancer
*	[20, 30)	Female	Kerala	Hindu	Viral infection
*	[20, 30)	Female	Tamil Nadu	Muslim	TB
*	[20, 30)	Male	Karnataka	Parsi	No illness
*	[20, 30)	Female	Kerala	Christian	Heart-related
*	[20, 30)	Male	Karnataka	Buddhist	TB
*	[10, 20)	Male	Kerala	Hindu	Cancer
*	[20, 30)	Male	Karnataka	Hindu	Heart-related
*	[10, 20)	Male	Kerala	Christian	Heart-related
*	[10, 20)	Male	Kerala	Christian	Viral infection

Table 5. Example: *hospital_extended.csv* anonymizing with *k*-anonymity with *k* = 2.

tests `pytest` (<https://github.com/pytest-dev/pytest>) (version $> = 7.1.2$ and $< 9.0.0$) has been used. These tests are located in the test folder of the library's repository.

Additionally, we are interested in determining the level of code coverage obtained with these tests, for which `coverage` (<https://about.codecov.io/>) has been used. In this sense, a workflow has been created in GitHub using a YML pipeline that creates a Python 3.10 environment, installs `pytest`, `poetry`, the dependencies and uses the `coverage run` command to run the tests and obtain the code coverage. Subsequently, using a `coverage` token, the code coverage level is published and can be analyzed within the <https://app.codecov.io/gh/IFCA-Advanced-Computing/anjanainteractive> interface, along with the scan that shows the lines that are fully covered, partially covered and uncovered. For version 1.0.0, the code coverage is greater than 92%.

Release automation. The `release-please` (<https://github.com/googleapis/release-please>) library has been used to create automated releases. A workflow has been defined using GitHub actions to create the releases when there is a push on the main branch. Finally, in case a new tag has been included, it is automatically published in PyPI, so that the installation can simply be carried out running the `pip install anjana` command.

Documentation. To create the library documentation, `sphinx` (<https://www.sphinx-doc.org/>) has been used, in order to generate the documentation of the implemented functions from the docstrings. For this purpose, the docstrings have been written using reStructuredText. Moreover, additional pages have been included to the documentation such as an installation and getting started guide, information about hierarchy management and how to manage multiple sensitive attributes, among others. The documentation has been published in readthedocs in an automated way and it is available in: <https://anjana.readthedocs.io/en/latest/>.

In addition to the availability of a comprehensive documentation (which will be enhanced in future updates adding further examples), in order to collect the user feedback and promote collaboration, the repository contains a file `CONTRIBUTING.md`, with instructions on how to open a pull request to ask for a bug fix, request a new feature, etc.

Requirements. In addition to the requirements for the tests, lint and build phases, the requirements for the proper operation of the library are the following: `Python`²⁶ (version $> = 3.9$), `numpy`²⁷ (version 2.0.1), `pan-das`²⁸ (version 2.2.2), `pycanon`²³ (version 1.0.1.post2), `typing_extensions` (https://github.com/python/typing_extensions) (version 4.12.2), `beartype` (<https://github.com/beartype/beartype>) (version 0.18.5) and `docutils` (<https://docutils.sourceforge.io/>) (version 0.21.2). Note that `dependabot` has been used for keeping the dependencies up to date, so the versions of the dependencies may change (see the workflow implemented in <https://github.com/IFCA-Advanced-Computing/anjana/blob/main/.github/dependabot.yml>).

Discussion

Multiple sensitive attributes. In version 1.0.0, *anjana* allows the incorporation of a single sensitive attribute (SA) for the anonymization process with the different techniques implemented. If there are more than one SA, two approaches can be followed as stated in²³: *harmonization of the quasi-identifiers* or *quasi-identifiers update*. In the first one we should apply each method to each sensitive attribute recursively, i.e., we perform the anonymization for the first SA and then anonymize the resulting dataset with respect to the next one, and so on until all of them are covered. In the second case (*quasi-identifiers update*) it is considered that some of the sensitive attributes can act as a quasi-identifier for the rest of the sensitive attributes. In this case, we shall anonymize the data for the first SA adding to the list of quasi-identifiers the remaining sensitive attributes that may act as QIs in the case of the current SA. Once anonymized for that SA, the process is repeated for the remaining ones updating in each case the list of QIs according to the data and attributes requirements.

Creating the hierarchies. Regarding the hierarchies used for generalizing the quasi-identifiers, all the anonymity functions available in *anjana* receive a dictionary with the hierarchies to be applied to the QIs. In

name	age	gender	city	religion	disease
*	[20, 30)	Female	*	Hindu	Cancer
*	[20, 30)	Male	*	Hindu	Cancer
*	[20, 30)	Male	*	Hindu	Cancer
*	[20, 30)	Male	*	Hindu	Cancer
*	[20, 30)	Female	*	Hindu	Viral infection
*	[20, 30)	Female	*	Muslim	TB
*	[20, 30)	Male	*	Parsi	No illness
*	[20, 30)	Female	*	Christian	Heart-related
*	[20, 30)	Male	*	Buddhist	TB
*	[10, 20)	Male	*	Hindu	Cancer
*	[20, 30)	Male	*	Hindu	Heart-related
*	[10, 20)	Male	*	Christian	Heart-related
*	[10, 20)	Male	*	Christian	Viral infection

Table 6. Example: *hospital_extended.csv* anonymizing with k -anonymity with $k = 2$ and ℓ -diversity with $\ell = 2$. Note that this table verifies k -anonymity for $k = 3$.

particular, this dictionary has as key the names of the columns that are QIs to which a hierarchy is to be applied (it may happen that a user does not want to generalize some QIs and therefore no hierarchy is to be applied to them, just do not include them in this dictionary). The value for each key (QI) is formed by a dictionary in such a way that the value 0 has as value the raw column (as it is in the original dataset), the value 1 corresponds to the first level of transformation to be applied, in relation to the values of the original column, and so on with as many keys as levels of hierarchies have been established.

For a better understanding, let's look at the following example. Suppose that we have the following simulated dataset given in Table 4 (corresponding with the *hospital_extended.csv* dataset used for testing purposes) with *age*, *gender* and *city* as quasi-identifiers, *name* as identifier, *disease* as SA and *religion* as insensitive attribute. Regarding the QIs, we want to apply the following hierarchies: interval of 5 years (first level) and 10 years (second level) for the *age*. Suppression as first (and only) level for both *gender* and *city*.

Then, in order to create the hierarchies we can define the following dictionary for *age*, *gender* and *city*, using the auxiliary function *generate_intervals()* available in the *utils* subpackage from the library to generate interval-based hierarchies. This function receives as input the original data to be generalized, the lower and upper bounds for the interval set and the spacing between the values of the intervals (step). Note that following the requirement previously mentioned, for the quasi-identifiers *gender* and *city* only suppression is applied, while for *age* intervals of five and ten years are used as generalization levels 1 and 2 respectively, as stated in Code 2.

```
import numpy as np
from anajana.anonymity import utils

hierarchies = {
    "age": {
        0: data["age"].values,
        1: utils.generate_intervals(data["age"].values, 0, 100, 5),
        2: utils.generate_intervals(data["age"].values, 0, 100, 10),
    },
    "gender": {
        0: data["gender"].values,
        1: np.array(["*"] * len(data["gender"].values)) # Suppression
    },
    "city": {
        0: data["city"].values,
        1: np.array(["*"] * len(data["city"].values)) # Suppression
    }
}
```

Example Code 2. Example: creating hierarchies for the quasi-identifiers of a database.

In view of the hierarchies created in Code 2 for the three quasi-identifiers, if we anonymize Table 4 with *name* as identifier, *age*, *gender* and *city* as quasi-identifiers, and *disease* as SA, using a suppression level of 0% and k -anonymity with $k = 2$, the resulting dataset is shown in Table 5. In addition, Table 6 shows the resulting dataset when in addition to k -anonymity with $k = 2$, ℓ -diversity is also applied with $\ell = 2$. For the first one, only the second level of generalization has been applied to *age* (intervals of 10 years), and for the second one it has also been applied suppression for the QI *city* (first level of generalization created).

gender	age	hypertension	heart	ever	work	residence	average	bmi	smoking	
			disease	married	type	type	glucose level		status	stroke
Male	80.0	0	1	Yes	Private	Rural	105.92	32.5	never smoked	1
Female	49.0	0	0	Yes	Private	Urban	171.23	34.4	smokes	1
Female	79.0	1	0	Yes	Self-employed	Rural	174.12	24.0	never smoked	1
Male	81.0	0	0	Yes	Private	Urban	186.21	29.0	formerly smoked	1
Female	35.0	0	0	Yes	Self-employed	Rural	82.99	30.6	never smoked	0
Male	51.0	0	0	Yes	Private	Rural	166.29	25.6	formerly smoked	0

Table 7. Sample extracted from the stroke dataset (<https://www.kaggle.com/fedoriano/stroke-prediction-dataset>).

k	Suppressed records (%)	Suppressed QIs	Transformation applied	Elapsed time (s)
2	10.38%	2	[0, 3, 0, 0, 0, 1, 0, 3, 2, 1]	4.1364 s
3	17.27%	2	[0, 3, 0, 0, 0, 1, 0, 3, 2, 1]	4.4706 s
4	14.52%	3	[1, 3, 0, 0, 0, 1, 0, 3, 2, 1]	4.6565 s
5	18.03%	3	[1, 3, 0, 0, 0, 1, 0, 3, 2, 1]	4.8107 s
10	16.13%	6	[1, 3, 1, 1, 1, 1, 0, 3, 2, 1]	5.6093 s

Table 8. Anonymizing the stroke dataset using *k-anonymity*: efficiency analysis. Suppression limit: 20%.

k	Suppressed records (%)	Suppressed QIs	Transformation applied	Elapsed time (s)
2	76.31%	0	[0, 1, 0, 0, 0, 0, 0, 2, 1, 0]	1.2181 s
3	74.15%	0	[0, 2, 0, 0, 0, 0, 0, 3, 1, 0]	2.2031 s
4	71.81%	0	[0, 2, 0, 0, 0, 0, 0, 3, 2, 0]	2.6897 s
5	78.49%	0	[0, 2, 0, 0, 0, 0, 0, 3, 2, 0]	2.6529 s
10	69.89%	1	[0, 3, 0, 0, 0, 1, 0, 3, 2, 0]	3.7985 s

Table 9. Anonymizing the stroke dataset using *k-anonymity*: efficiency analysis. Suppression limit: 80%.

Anonymization	Suppressed records (%)	Suppressed QIs	Transformation applied	Elapsed time (s)
<i>t</i> -closeness, $t = 0.4$	84.80%	3	[1, 3, 0, 0, 0, 1, 0, 3, 2, 1]	2.5345 s
Enhanced β -likeness, $\beta = 5$	84.80%	0	[0, 1, 0, 0, 0, 0, 0, 1, 1, 0]	0.9677 s
δ -disclosure privacy, $\delta = 3.5$	84.80%	3	[1, 3, 0, 0, 0, 1, 0, 3, 2, 1]	2.6596 s

Table 10. Anonymizing the stroke dataset using *t*-closeness, enhanced β -likeness and δ -disclosure privacy: efficiency analysis.

Anonymity transformation applied. Regarding the transformation applied for anonymizing the data based on the hierarchy level applied to each quasi-identifier, in some cases it is important to extract it in order to transfer statistics on the processing performed on the data. Usually, this transformation will be denoted with a list of the same length as the number of quasi-identifiers. When performing the anonymization process, the quasi-identifiers are entered in a certain order, which will be the same as the order in which they are represented in the list with the transformation.

To obtain this transformation, the `get_transformation()` function from the `utils` sub-module can be used introducing the anonymizing dataset, the set of quasi-identifiers and the dictionary containing the hierarchies. This can be of particular interest in many cases, such as to transmit statistics about the anonymization process performed without sharing the data. Additionally, in an architecture where different data owners are anonymizing different datasets following the same hierarchies, if once the data is anonymized they will want to train a global model, either in a distributed way (for example under a federated learning architecture) or in a centralized way, it will be necessary to harmonize the transformations carried out. Thus, it will be of interest that all data owners have applied the same transformations to the quasi-identifiers in order to train the model (and, if necessary, aggregate it) in a consistent way.

Specifically, the function `get_transformation()` returns a list so that the value in the position i corresponds to the hierarchy level applied for the quasi-identifier entered in position i in the list of quasi-identifiers entered during the anonymization process, $\forall i \{1, \dots, n_{QI}\}$, with n_{QI} the number of quasi-identifiers introduced. For example, a transformation [1,0,2,3] represents that the first level of hierarchy has been applied to the first QI, that no transformation has been applied to the second, and that the second and third levels of hierarchies have been applied respectively to the third and fourth quasi-identifiers.

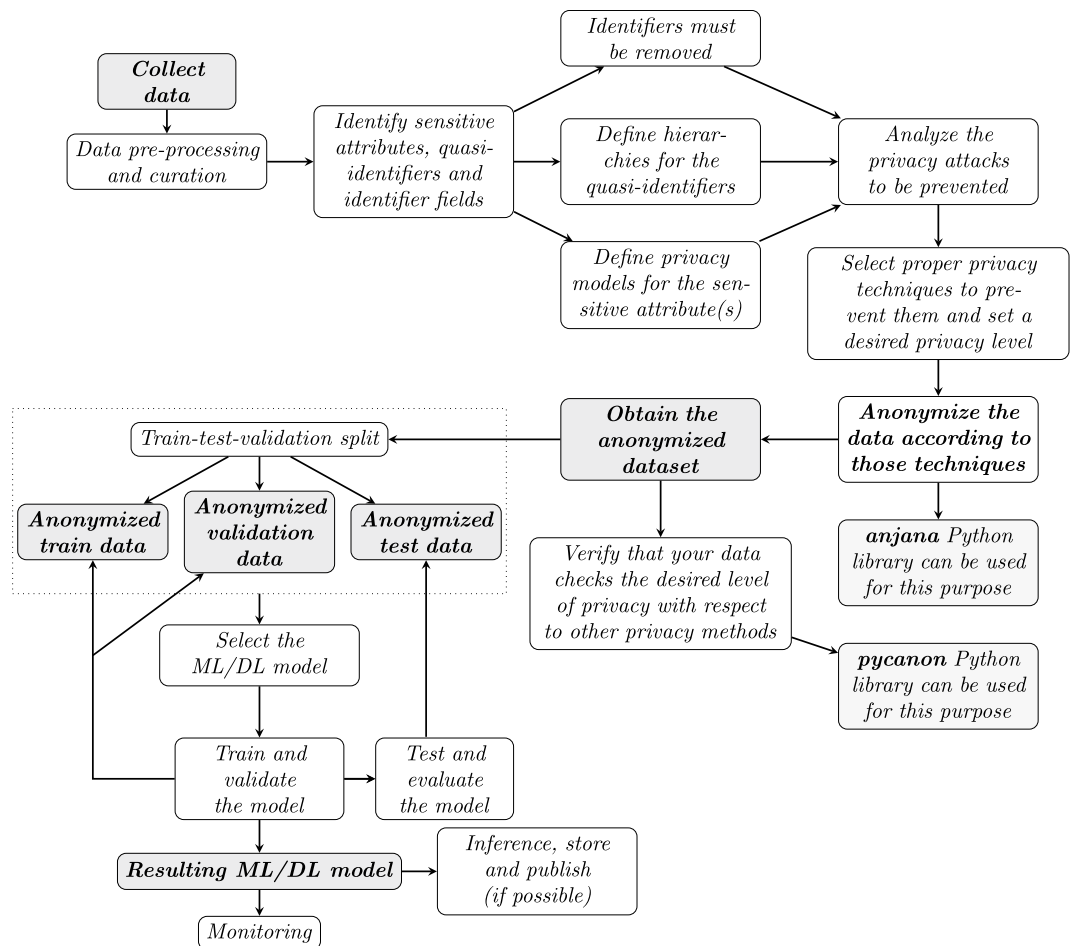


Fig. 3 Classic workflow of a data-based AI project including the training/validation and testing phase of ML/DL models and the data anonymization process.

In the example defined above, note that the transformation applied to Table 4 for obtaining Table 5 is [2,0,0], while for obtaining Table 6 the transformation is [2,0,1].

Efficiency comparative analysis. Once we know how to obtain the applied anonymity transformation, how to create the hierarchies and how to anonymize the data with the different techniques available in *anjana*, we are interested in comparing the performance of the different anonymization methods implemented and their impact on the usability of the data once anonymized.

To this end, we will use the widely known *stroke dataset*, which was also used in *pycanon* to test the functionality of the library (see²³). In particular, we are going to apply different anonymization techniques on this data, quantify the execution time required in *anjana* to achieve the desired anonymity level, and analyze the number of records eliminated in each case as well as the number of QIs suppressed. Table 7 shows an extraction of the data available in the original database.!

In this example we are dealing with a case with ten quasi-identifiers: *gender, age, hypertension, heart disease, ever married, work type, residence type, average glucose level, BMI* and *smoking status*, and one sensitive attribute: *stroke* (whether or not the patient has suffer a stroke). Two out of ten QIs take floating values (*average glucose level* and *BMI*) and one takes integer values (*age*). The hierarchies applied are suppression in the case of categorical and binary values, and generalization to intervals of 5, 10 and 20 for *age* and *average glucose level*, and intervals of 5 and 10 for *BMI*. Regarding the number of records, 4909 rows are available.

Anonymizing using *k*-anonymity. Let's start by analyzing the impact of the application of *k*-anonymity for different values of *k*. We are interested in evaluating the computation time required, taking into account the complexity of the quasi-identifiers that are of floating type. We are also interested in assessing the number of quasi-identifiers that have been completely suppressed to obtain the desired level of anonymity, as this may result in a high loss of information that makes the data unusable for analysis. We will evaluate two paradigms: in the first one, we want to keep a large volume of data, so we establish the suppression limit to 20%, while in the second one, we value the usefulness of the resulting anonymized data and therefore should make less generalization steps, so we set the suppression limit to 80%. The results obtained in each case are shown in Tables 8 and 9

respectively. Note that for those QIs with intervals as hierarchies, their maximum generalization level does not imply their suppression (they can not be suppressed with hierarchies established).

From Tables 8 and 9 we can draw several immediate conclusions: (1) it is trivial that allowing a higher suppression limit makes the anonymization process much faster, since it is not necessary to perform as many rounds of hierarchy application on the set of QIs; (2) a higher suppression limit enables less generalization of the QIs, which is appropriate if a large volume of data is available so that the data does not lose its usefulness when overgeneralized; (3) it is important to strike a balance between usability, generalization and the volume of data we can suppress, which goes hand in hand with the intended use and purpose of the data; (4) in all cases the anonymization is successfully completed within a few seconds.

Anonymizing focusing on the sensitive attribute. In the previous analysis we have only applied *k-anonymity*, which does not take into account the sensitive attribute. Then, we will carry out an analysis of the efficiency by applying techniques that focus on the sensitive attribute, which in this case is the *stroke* attribute.

It is important to note that with the established hierarchies we have observed that *ℓ-diversity* for $\ell = 2$ was not achieved even allowing a suppression limit of 95%. This makes sense from the data analytics and utility point of view: if in the same equivalence class the *stroke* attribute can take both values (0 and 1), the dataset would not be very suitable for use in training a ML or DL model that allows us to predict whether or not a patient will develop such a pathology based on the remaining features (the QIs), that is, for performing a binary classification task. Then, in order to test other anonymity functions with the purpose of assessing runtime, we have applied *t-closeness*, *enhanced β -likeness* and *δ -disclosure privacy* with very relaxed restrictions on the maximum number of records removed (an upper suppression limit of 90%), since a wide generalization of the QIs would make the data less useful for training data based models. The results obtained are shown in Table 10 below. Note that the value of *k* for *k-anonymity* is set to 2, and that up to 90% of the records can be removed.

Table 10 allows us to know the impact of the anonymization process carried out. If the objective is in fact to use these data as the basis of a ML/DL model for classification, in both cases we will have more than 700 records, which in many cases may be sufficient for an initial approximation with such models. However, the suppression of QIs is more critical, and it is likely that the anonymization obtained using *t-closeness* and *δ -disclosure privacy* will not be of interest in practice, since three quasi-identifiers have been suppressed and 3 have been completely generalized. However, in the case of $\beta = 5$ for *enhanced β -likeness*, any QI has been removed, being more useful for data analysis than even the case of $k = 2$ with 20% of suppression limit (it that case 2 QIs were removed). This is also reflected in the computation time, which is much lower.

Finally, in the absence of any additional technique applied to *k-anonymity* with $k = 2$ and 90% of suppression limit, the value of *t* for *t-closeness* was 0.4879, $\beta = 4.4175$ for *enhanced β -likeness* (note that $\beta = 5$ was already satisfied without applying further steps), and $\delta = 3.7244$ for *δ -disclosure privacy*, which could be easily checked using *pycanon* on the data anonymized with *anjana*.

For performing this efficiency and runtime analysis a machine provided with 12th Gen Intel(R) Core(TM) i5-1 and with 16 GB RAM running under Ubuntu 22.04.4 LTS was used.

Anonymization in data science workflows. Applying the previously exposed anonymity techniques appropriately with special attention to the privacy attacks that are to be prevented is a fundamental step in a data science process. In particular, the diagram in Fig. 3 details the flow of a machine or deep learning project, including the correct management of the privacy of the data involved.

From the diagram in Fig. 3, several factors must be taken into account: first, the fact that the initial dataset verifies a certain privacy constraint does not mean that once it has been split into train/val/test, it will continue to be verified. Therefore, it may be interesting to perform the train test split before anonymizing. In the same sense, if the train test split is performed before anonymizing, it should be considered that in order to train, validate and evaluate the ML/DL model, it is convenient that the same transformations have been applied at the time of anonymizing the data. In other words, if the model has been trained with data transformed by applying certain hierarchies, the validation and test data must have been anonymized with the same levels of hierarchies. It is important to note that these hierarchies may not allow to reach the desired level of anonymity, so it is necessary to be quite conservative in this part. Thus, depending on the objective and scope of the model to be developed, it will be interesting to follow the scheme shown in Fig. 3 or to invert it so that the anonymization is performed just before the train/validation/test split, harmonizing in each case the transformations applied in terms of the hierarchies on the quasi-identifiers (taking the strictest one in order to guarantee the required privacy level for the three data splits).

Future work. This paper presents the implementation of an open source Python library that allows to anonymize tabular data using the most commonly used techniques in an intuitive and user-friendly way. This tool aims to address the need to provide researchers and general users with a comprehensive tool that allows them to locally anonymize their sensitive data using Python, so that it can be easily integrated into the workflow of an ML/DL project. For this purpose, nine anonymization techniques can be applied on tabular data by introducing the data, identifiers, quasi-identifiers, sensitive attribute, maximum level of records' suppression allowed and privacy level required according with the method to be applied.

Regarding future lines of work, it is of particular interest to include additional anonymity tools together with metrics to quantify the loss of information during the anonymization process, as well as to analyze the risk of attribute disclosure. Also, in order to make the tool even easier and more attractive, the creation of an application that allows users to carry out the anonymization process through a local interface or dashboard (so that the data does not have to be uploaded to any web or external service) will also be explored. Finally, identifying ways

to quantify resistance to attacks based on auxiliary information is an interesting line of work, as well as incorporating additional techniques, including pseudonymization of identifiers.

In addition, within the pycanon library, we started to integrate the calculation of various metrics to measure the amount of information lost during anonymization and the usefulness of the data. As a future work, the idea is that this functionality becomes part of anjana but in an extended way, integrating different metrics for the usefulness of the data, being able to compare the raw data that was anonymized, and the resulting dataset after anonymization using anjana. It is thus more intuitive that this functionality belongs to anjana since pycanon can be used if only the anonymized data is available to know its privacy level, and to calculate these metrics it is also necessary to know the original data to quantify the loss. In addition to the average equivalence class size, classification metric and discernability metric, which are already integrated in pycanon, the incorporation of other techniques, such as the generalized information loss and the cell information gain among others will be explored and implemented.

Concerning users adoption, as future work we aim to prepare a set of tutorials and guides on the use of the library for generic data, but also in specialized sectors. It's important to note that the development of this software is framed in the context of the EOSC SIESTA project, in which five use cases belonging to quite different fields are involved: epidemiology, medical imaging, energy domain, text anonymization for sensitive data and demography. The objective is to work with these use cases as early adopters (those who are appropriate candidates to use this software), in order to collect their feedback and make further improvements based on it.

Data availability

No new data have been generated for this work. The data used to test the library and the corresponding anonymized data generated according to some techniques tested can be found in the [examples folder of the library's repository \(v1.0.0\)](#).

Code availability

Relevant links regarding the availability of the source code, documentation and installation of the software are listed below. Please note that the version referred to in this manuscript is v1.0.0.

- The code for the *anjana* library is openly available in GitHub: <https://github.com/IFCA-Advanced-Computing/anjana>.
- The documentation can be found in ReadTheDocs: <https://anjana.readthedocs.org>.
- A Zenodo DOI has been created to ensure code reproducibility and availability⁷. This DOI stands for all versions, and always resolves to the latest one: <https://doi.org/10.5281/zenodo.11184467>.

Received: 21 August 2024; Accepted: 18 October 2024;

Published: 26 November 2024

References

1. Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K. & Galstyan, A. A survey on bias and fairness in machine learning. *ACM Comput. Surv.* **54**, <https://doi.org/10.1145/3457607> (2021).
2. Hort, M., Chen, Z., Zhang, J. M., Harman, M. & Sarro, F. Bias mitigation for machine learning classifiers: A comprehensive survey. *ACM J. Responsib. Comput.* <https://doi.org/10.1145/3631326> (2023).
3. European Commission. Regulation 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. <https://data.europa.eu/eli/reg/2016/679/oj>. [Accessed 16-05-2024].
4. European Commission. Regulation of the european parliament and of the council laying down harmonised rules on artificial intelligence (artificial intelligence act) and amending certain union legislative acts. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52021PC0206>. [Accessed 16-05-2024].
5. Office of Privacy and Civil Liberties. U.S. Department of Justice. Privacy act of 1974. <https://www.justice.gov/opcl/privacy-act-1974>. [Accessed 16-05-2024].
6. Centers for Medicare & Medicaid Services. The Health Insurance Portability and Accountability Act of 1996 (HIPAA). Online at <http://www.cms.hhs.gov/hipaa/> (1996).
7. Sáinz-Pardo Daz, J. & López Garca, Á. Anjana, <https://doi.org/10.5281/zenodo.13320086> (2024).
8. Sweeney, L. k-anonymity: A model for protecting privacy. *International journal of uncertainty, fuzziness and knowledge-based systems* **10**, 557–570 (2002).
9. Wong, R. C.-W., Li, J., Fu, A. W.-C. & Wang, K. (α , k)-anonymity: an enhanced k-anonymity model for privacy preserving data publishing. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, 754–759 (2006).
10. Machanavajjhala, A., Kifer, D., Gehrke, J. & Venkatasubramanian, M. l-diversity: Privacy beyond k-anonymity. *Acm transactions on knowledge discovery from data (tkdd)* **1**, 3–es (2007).
11. Li, N., Li, T. & Venkatasubramanian, S. t-closeness: Privacy beyond k-anonymity and l-diversity. In *2007 IEEE 23rd international conference on data engineering*, 106–115 (IEEE, 2006).
12. Brickell, J. & Shmatikov, V. The cost of privacy: destruction of data-mining utility in anonymized data publishing. In *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '08, 70–78, <https://doi.org/10.1145/1401890.1401904> (Association for Computing Machinery, New York, NY, USA, 2008).
13. Cao, J. & Karras, P. Publishing microdata with a robust privacy guarantee. *arXiv preprint arXiv:1208.0220* (2012).
14. Domingo-Ferrer, J. & Soria-Comas, J. Anonymization in the time of big data. In Domingo-Ferrer, J. & Pejić-Bach, M. (eds.) *Privacy in Statistical Databases*, 57–68 (Springer International Publishing, Cham, 2016).
15. Buitinck, L. et al. API design for machine learning software: experiences from the scikit-learn project. In *ECML PKDD Workshop: Languages for Data Mining and Machine Learning*, 108–122 (2013).
16. Chollet, F. et al. Keras. <https://keras.io> (2015).
17. Abadi, M. et al. TensorFlow: Large-scale machine learning on heterogeneous systems Software available from tensorflow.org. (2015).
18. Paszke, A. et al. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems* **32**, 8024–8035 (Curran Associates, Inc., 2019).

19. Ayala-Rivera, V. *et al.* A systematic comparison and evaluation of k-anonymization algorithms for practitioners. *Trans. Data Priv.* **7**, 337–370 (2014).
20. Yuvaraj, N., Praghash, K. & Karthikeyan, T. Privacy preservation of the user data and properly balancing between privacy and utility. *International Journal of Business Intelligence and Data Mining* **20**, 394–411 (2022).
21. Madrazo Quintana, E. *Building a Python library for anonymizing sensitive data*. Master's thesis, University of Cantabria (2023). <https://repositorio.unican.es/xmlui/handle/10902/30791>. Supervisors: López García, Álvaro and Sáinz-Pardo Díaz, Judith.
22. Prasser, F. & Kohlmayer, F. Putting statistical disclosure control into practice: The arx data anonymization tool. *Medical data privacy handbook* 111–148 (2015).
23. Sáinz-Pardo Daz, J. & López Garca, Á. A python library to check the level of anonymity of a dataset. *Scientific Data* **9**, 785 (2022).
24. Becker, B. & Kohavi, R. Adult. UCI Machine Learning Repository <https://doi.org/10.24432/C5XW20> (1996).
25. Sáinz-Pardo Daz, J. & López Garca, Á. Comparison of machine learning models applied on anonymized data with different techniques. In *2023 IEEE International Conference on Cyber Security and Resilience (CSR)*, 618–623 (IEEE, 2023).
26. Van Rossum, G. & Drake, F. L. *Python 3 Reference Manual* (CreateSpace, Scotts Valley, CA, 2009).
27. Harris, C. R. *et al.* Array programming with NumPy. *Nature* **585**, 357–362, <https://doi.org/10.1038/s41586-020-2649-2> (2020).
28. Wes McKinney. Data Structures for Statistical Computing in Python. In S. van der Walt & J. Millman (eds.) *Proceedings of the 9th Python in Science Conference*, 56–61, <https://doi.org/10.25080/Majora-92bf1922-00a> (2010).

Acknowledgements

The authors would like to thank the funding through the SIESTA project “Secure Interactive Environments for Sensitive daTa Analytics”, funded by the European Union (Horizon Europe) under grant agreement number 101131957.

Author contributions

Both authors defined and conceived this work. J.S.-P.D. performed the formal analysis, methodology, validation and software development. A.L.G. was responsible for supervision, project administration and funding acquisition. Both authors contributed to writing and reviewing the manuscript.

Competing interests

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Additional information

Correspondence and requests for materials should be addressed to J.S.-P.D.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2024, corrected publication 2024