www.tno.nl/en/digital/data-sharing/privacy-enhancing-technologies/

Version 24 April 2026

PET Decision Tree Guide

Contents

1 Introduction

According to the European Union Agency for Cyber Security, *Privacy-Enhancing Technologies (PETs)* is a term that covers the broader range of technologies that are designed for supporting privacy and data protection. These technologies incorporate the data protection principles by minimizing personal data use, maximizing data security and/or giving control to data subjects over their data.

Examples of PETs include pseudonymization, multi-party-computation, differential privacy and homomorphic encryption. The development and emerging popularity of using PETs in data processing operations aligns with current discussions around the idea of shaping technology according to privacy principles, as new technologies may bring about unforeseen risks. At the same time, legislation is updated to catch up with these developments, such as the General Data Protection Regulation (GDPR) being a main regulator of data privacy within the European Union. The GDPR obliges organizations i.a. to take (technical) measures to ensure privacy by design and default as data protection principles. The use of PETs, depending on the context and case, can help organizations to comply with these principles.

The benefit of using PETs becomes evident when an organization wishes to tackle challenges in relation to data sharing with another party. Legal regulations like GDPR may render this exchange impossible with traditional data-exchange based approaches. Next to privacy, also (other) data confidentiality reasons can prevent data sharing, even when legislation does not explicitly prohibit this. For instance, due to other regulations or organizational interests, e.g. due to commercial interests or agreements with customers. We therefore do not limit the scope of PETs to only personal data. PETs enable a paradigm where organizations can *leverage the information* that is stored in sensitive data[1] *without revealing the sensitive data* itself.

So how does one transit from theory to practice? Acknowledging that PETs might help you to solve a business challenge is only the first step to applying a specified PET to that challenge. Among the many steps that need to be taken, certainly a crucial one is to understand the business challenge at hand and investigate which PET can facilitate a solution. Which data is processed? What is the intended result? Which regulations apply? What are the technical constraints? Which PETs could be applicable? How do we balance technical guarantees (PET characteristics) and legal guarantees (formal agreements)? Although PETs are technologies, this is not a discussion between engineers and developers alone – it is a conversation between various stakeholders with diverse expertises and perspectives. Just like PETs enable privacy-enhanced solutions for single- and multi-organizational challenges, our work enables a multi-disciplinary discussion about PETs in the context of such business challenge.

The tool that we present is a Decision Tree that is designed to support the choice of a *PET in the context of inter-organizational data analysis*. We expect the tool to be useful when performing a Data Protection Impact Assessment (DPIA), or when considering which PET to use.

This document serves as a guide for an organization to use the Decision Tree efficiently and successfully. Its purpose is to facilitate a discussion that involves technical and legal aspects;

[1] Note: multiple times we will refer in this guide as well as in the decision tree to 'sensitive data'. This is meant in a broad setting: data which may be sensitive due to many reasons (privacy, commercial confidentiality, etc.).

however, note that it is not a legal document and you should always conduct your own legal assessment before using PETs. The tool itself is available on https://decisiontree.mpc.tno.nl/.

The Decision Tree and this document were created by the CBS, KNB, Rabobank, TNO and the University of Maastricht (who was involved in an earlier phase) in a use case of Brightlands Techruption[2]. Both technical as well as legal and compliance experts from the different organizations actively contributed. Brightlands Techruption helps corporate companies, governmental organizations and knowledge institutes partner up, so they can develop innovative solutions through the application of disruptive technologies like AI, MPC, SSI and blockchain.

1.1 Reading guide

The Decision Tree and this document are written for innovation departments and stakeholders with an interest in PETs to solve a given business challenge. All readers were assumed to have a somewhat technical background in the initial stage, but since then we tried to broaden the scope and terminology to also include non-technical readers – particularly with background in management or law and policy.

The intended use of the document, however, relates to complex, multidisciplinary challenges and likewise it is challenging to make the entire document easily readable to everyone. Instead, we think that the best results are obtained if all stakeholders scan this document, read those parts that relate to their expertise, and vocalize those parts in the joint discussion.

The document is structured as follows. Section 2 contains guidance on how the Decision Tree should be used. Section 3 gives a summary of the different PETs that are considered here. Section 4 contains several relevant Ethical, legal, social, and business considerations regarding PETs. Section 5 describes levels of data-sharing risk mitigations when using PETs. Section 6 contains disclaimers and explanations about the decision tree nodes when needed.

Appendix A describes different security scenarios when considering PETs. Appendix B describes a non-exhaustive list of potential attack vectors when considering PETs.

[2] https://www.brightlands.com/en/brightlands-smart-services-campus/brightlands-techruption

2 Decision tree usage

We encourage the reader to visit our webpage, look around, get insights and start a discussion. Before doing so, however, it is important to understand a bit more about the usage of the tree; both conceptually and technically.

Objective of the tree

It is crucial to understand that the main objective of the tree is to facilitate the process of exploring the potential of several PETs as solution to your business problem. It achieves this objective by presenting the user with multidisciplinary questions and topics that need to be discussed already in an early stage of the process. At the same time, the user can explore the impact of the answers to these questions; for example, how the answers lead to a suggested PET solution. The discussions and explorations will often benefit by involving all stakeholders (business, risk officers, privacy officers, compliancy, data scientists, IT architects, end users). Although the decision tree always concludes by suggesting a PET suggestion, the real value is in the questions that lead you there.

Multiple problem-solutions

Related to the final remark above, there may be multiple solutions to any problem. Rather than running through the tree once and expecting a definite outcome, it is more likely that multiple paths are feasible. For example, there may be a problem where one possible solution is to use a PET e.g. MPC, whereas another path might lead you to a different solution e.g. to outsource processing operations to a trusted third party. Both solutions have their advantages in terms of data protection compliance, data processing efforts, governance, flexibility and so on. From that perspective, the questions in the decision tree highlight some of the possibilities and potential requirements of a solution. These questions can thus facilitate the user in formulating, exploring, and discussing the requirements of the use case. Such discussions may in turn lead to adjusted use case requirements, which then lead to alternative paths in the decision tree.

It is quite likely that several of these iterations are made in the development of a use case and exploring the possible solutions; in particular, you might walk through the tree many times while gradually improving your understanding of the problem and the applicable PET technologies. The tool checklist also facilitates this process.[3]

Multidisciplinary at its core

It is important that all organizations that participate in a collaboration should traverse the tree from their own perspective (again, involving all stakeholders). Different organizations may have different roles in the solution, e.g. they may share different types of data (if any) and their role from a legal or compliancy perspective may yield other restrictions or responsibilities. Every organization should think about their role in various potential solutions and discuss the implications and preferences collaboratively to identify the solution that works best for the collaborators as a group. Additionally, this applies to single aspects in the tree as well. Every challenge can be viewed from a different

[3] https://decisiontree.mpc.tno.nl/documentation/Checklist.pdf

perspective and which of these is most important differs per organization. The main perspectives for addressing a challenge in the tree are depicted in the following figure.

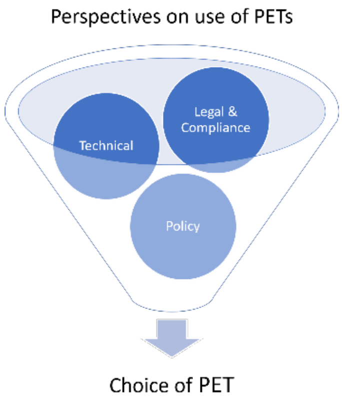

From a technical perspective, it is important to know both the fit and limitations in terms of end goal, IT setup or data science. Legal and compliance should be consulted to prevent potential legislative and privacy restrictions in a later stage. Lastly, there might be specific policy or supervisory guidelines within an organization that may limit the possibilities to use a specific implementation, for instance a restriction on setting up any third parties for the solution.

Fourth, it should be recognized that many challenges are composed of multiple smaller challenges. One should traverse the tree as many times as applications you wish to develop. For example, in the context of Machine Learning, you may run through once for model training and once for model evaluation.

Disclaimer

We repeat and emphasize that the tool is no (legal) advise. It does not provide or replace professional legal counsel, nor should it be relied upon for compliance decisions. It is extremely challenging to capture the complexity and context-dependency of generic real-world challenges and the subtleties in the variety of PET solutions in a tool. The effectiveness and appropriateness of any given PET will always depend on the specific organizational, technical, and regulatory context in which it is applied.

Instead, aiming at a first step in applying PET solutions, the purpose of the tool is to facilitate both the internal and external interdisciplinary discussion for organizations that are interested in using PETs for the challenges they face. In this sense, the tool should be viewed as a conversation starter and orientation aid rather than a prescriptive solution. Its main value lies in helping organizations identify relevant PETs, clarify their needs, and foster collaboration between technical experts, legal advisors, and decision-makers. By creating a shared foundation for discussion, it supports more informed, transparent, and responsible consideration of PETs as part of broader data protection and innovation strategies.

2.1 Scope

In this guide, as well as in the accompanying decision tree, the focus is on application of PETs in the context of inter-organizational data analysis – i.e., multiple organizations (which could also be part of one larger organization) aim to perform joint analyses on their data sets. For this purpose, we focus on a relevant subset of PETs in our decision tree. These PETs are:

- Federated Learning (FL)
- Secure Multi-Party Computation (MPC), divided into
  - homomorphic encryption
  - secret sharing
- Differential Privacy
- Trusted Secure Environments (TSE, also known as Trusted Execution Environments, TEE)
- Synthetic Data Generation (SDG)

This does not mean that other PETs are not relevant, but we feel these are several of the more important categories when considering the purpose of inter-organizational data analysis. Furthermore, next to privacy of individuals there can be other reasons to apply these PETs, e.g., when dealing with commercially confidential data – these other possible applications are also in scope.

When considering protection of information in data analysis, there are two views of privacy, namely:

1. Input Privacy
2. Output Privacy

Input Privacy refers to the process of keeping the input data to a computation private. The necessity of input privacy is self-evident in the context of sharing sensitive data, as lack thereof is equivalent to sharing raw data among parties, which defeats the purpose. On the other hand, *Output Privacy* refers to techniques used to ensure that the output of a computation does not reveal information about the input data. Preserving output privacy is not always trivial. In a recent case, the US Census Bureau showcased that they were able to use their own publicly released statistics to reconstruct the raw data of the individuals described by these statistics[4]. This result motivated the Bureau to incorporate Differential Privacy in their computations to achieve output security.

In this document, we focus primarily on input privacy. We do include one state-of-the-art output security technique in our list of PETs, namely Differential Privacy. However, we should stress that there are other disclosure avoidance techniques that detect whether the output of an analysis leaks information about the input data. These other output privacy techniques are out of scope.

Several times we refer to 'sensitive data' in both this document as well as the decision tree. This is meant in a broad setting: data which may be sensitive due to many reasons (privacy, commercial confidentiality, etc.). This is a broader definition than the term 'sensitive personal data' used in the GDPR, where it refers to specific categories of personal data (medical, ethnic background, etc.).

2.2 Choosing a view

When accessing the tool, you immediately see that the tree is available in two different formats: an interactive tree graph and a questionnaire. Both views contain the same information/decision paths.

[4] https://www2.census.gov/about/policies/2020-03-05-differential-privacy.pdf

Which view is more useful depends on the user's preference. Our rough suggestion on how to profit from both is:

1. First, use the interactive tree. Explore different paths to get a feeling of the tool's utility and the reasoning of the line of questions and answers.
2. Then, use the questionnaire view to get a suggestion from the tool and export said suggestion along with the choices that led you there (by clicking the Save as PDF button).

We now elaborate further on both views.

2.2.1 Interactive tree

The interactive tree is a graph where the user is asked to answer questions regarding the challenge they seek to tackle and its characteristics by clicking on the equivalent visual elements. The interactive tree view allows the user to explore different paths rapidly and visualizes the many possibilities. Exploring the tree in this view can assist the user with understanding how certain (early) choices give direction to the proposed solution for various scenarios.

More specifically, the elements in the tree are:

- **Rectangular:** A question and nodes that link it to the possible answers. The text contains a question mark element that when hovered over pops up a window where the question in hand is explained.
- **Rhombus:** A possible answer. This element is clickable and the user can hence select an answer. Some of the answers can also be hovered over to produce pop-up explanations. When an answer is selected, the next question in the path is revealed.
- **Rectangular with rounded edges:** A decision on the PET technology based on the traversed path. The text contains a question mark element that when hovered over pops up a window where the decision is explained.

The other buttons seen on screen are:

- *Toggle style:* Choice to change colors.
- *Toggle contrast:* Choice to shift to black and white view.
- *Expand tree:* Choice to see the entire tree with all possible paths.
- *Collapse tree:* Choice to collapse the tree, i.e. reset the path.
- *Auto collapse alternative routes:* When selected, alternative routes automatically collapse when a different route is selected.

2.2.2 Questionnaire

The questionnaire is a linear list that contains the same questions and answers as the interactive tree. It consists of a rectangular element that contains the question and the possible answers. When an answer is chosen, the next question appears and the previous non-chosen answers disappear. Like the interactive tool, the text of the questions (and sometimes the answers) also contains a question mark that pops up an explanation when hovered over. The final element that appears will be the PET technology suggested based on the previous answers to the questions. The utility of the buttons on screen is:

- **Back**: The last answer is unchosen and hence the previous question appears again.

- **Reset**: The path resets, hence all the previous answers are unchosen and the initial question appears.
- **Save as PDF**: A document is saved in the user's device that includes all the choices made and the suggested PET.

The questionnaire view is simpler than the interactive tree.

2.3 Type of problem

The first question being reached in both the decision tree as well as the questionnaire is: what is the type of problem being considered? In this section we will give a brief explanation on the types of problems that are included:

Machine learning

For the Machine Learning type of problem, we consider problems where an actual machine learning or AI model is involved. This can be via different ways: either the parties aim to *train* a new machine learning model on data, or the parties aim to *evaluate* or apply an existing (already trained) machine learning model on data. Machine learning is a broad concept, containing a broad range of classification or regression models, intended to make a prediction based on a number of features. We do not consider other forms of statistical analysis as they are separately examined in the Statistical Analysis type of problem. It may be that the specific problem you consider fits both machine learning and statistical analysis – in that case, it may be good to traverse both paths in the tree.

Set intersection

In this type, we consider problems where two (or more) parties have a list of items (e.g. persons like patients or customers), and wish to determine the overlap between these two lists (the intersection). Set Intersection can be either a subproblem of any of the other problems, or a problem by itself. An example of the second scenario is when organizations wish to match their datasets without planning to perform a specific analysis per se. In this case, only the Set Intersection route shall be traversed. On the contrary, when additional analysis is intended, the tree shall be traversed twice: once for Set Intersection and once for said analysis (Machine Learning, Statistical Analysis or Synthetic Data Generation).

Statistical analysis

By statistical analysis, we refer to cases where one or more parties wish to compute a set of statistical metrics (e.g. counts, averages, standard deviations, quantiles, histograms, frequency plots, etc.) on their data and receive the results. Also other simple computations on the data, even if not strictly statistical in nature, may be considered when traversing this path.

Synthetic data generation

Synthetic data generation refers to cases where one wishes to generate new (fake) data based on existing data's distribution and characteristics. For instance, to validate and test models, or to train machine learning models. It can also be required for data retention or sharing. Here we assume that the original data used to generate the synthetic data is sensitive. If the original data is not, it is very straightforward to synthesize data without having to employ some privacy preserving technology to protect the original data from potential reconstruction by using the synthetic data.

3 Introduction per PET

For traditional analysis, we assume that all data is locally available (centralized data), thus allowing us to focus on the system efficiency, model performance, and applicability of some analysis of interest. Additionally, we assume that the end user of the output has no evil intentions of any kind; for example, we like to assume that the end user has no incentive to reverse-engineer the training data in the context of machine learning. Although this setting suffices in many cases, it also poses severe limitations to collaboration with sensitive data. PETs enable such collaborations.

First consider the assumption of centralized data and instead assume that data is distributed over different organizations or even different devices (smartphones). In this case, it may be time-consuming or even infeasible to send all the distributed data to a central location due to constrained resources, such as high latency or small bandwidth. These constraints motivate the development of techniques that allow for analysis on distributed data. We will shortly describe federated learning, which is a type of learning from distributed data with lower latency, less power consumption, and enhanced end users' privacy.

Another reason why not all data might be available in a central place is a reserve of sharing data, presumably motivated by the private or confidential nature of the data. This is easily perceived as a reason to not collaborate; however, in many machine learning scenarios only the raw data is sensitive and not the trained model. Several techniques therefore aim to compute (an approximation of) the trained model that respects the privacy of the underlying data. Differential privacy and secure multi-party computation are designed for this purpose and provide mathematical guarantees of the type of privacy that they provide. Federated learning was originally introduced to address issues that arise in the context of distributed analysis involving many devices. One of the improvements over centralized analyses is improved privacy since the data of any device is no longer shared in raw form. However, unlike the mathematically guaranteed privacy in secure multi-party computation protocols, the privacy benefits of federated learning are often hard to quantify. Federated learning is therefore regularly combined with secure multi-party combination or differential privacy or both to boost privacy.

We now describe the techniques in greater detail. Note that all techniques are relatively novel and are rapidly improved upon. The attacks that we refer to are described in more detail in Appendix B. Organizations and individuals that act independently of each other are all referred to as *parties*.

3.1 Federated Learning (FL)

Federated learning specifically targets the issue of learning on distributed data. The core concept is, for every party, to obtain a partial model by training on the data that is locally available. Then, the partial models are aggregated (by a dedicated aggregator party or in a peer-to-peer architecture) into a global model that captures the information of all data. In doing so, instead of sharing all raw data, only the local models are shared (Konečný et al. 2016; McMahan et al. 2017).

Communication of these local models is not quite as demanding as communication of the raw data, which was one of the main objectives. Complementary to (1) limited communicational resources, federated learning problems are characterized by (2) systems heterogeneity, (3) statistical heterogeneity, and (4) privacy concerns (Li et al. 2019). Note that these characteristics also set Federated Learning apart from distributed learning on multiple servers in a data farm.

For the privacy concerns, we note that federated learning achieves some level of privacy because the local models only represent aggregated information of the raw data. In practise, however, Federated Learning is known to be susceptible to several type of attacks including backdooring (Bagdasaryan et al. 2019) and reconstruction attacks (Li et al. 2020; Zhu et al. 2019) and it is generally complex to quantify the privacy that is obtained.

While the field of federated learning is evolving, the scope itself is becoming broader. In the past years, the term federated learning has often been used to describe any type of learning where the data is partitioned among parties (Fernandez et al. 2024; Lee et al. 2023; Fernández et al. 2022). In this document and designed Decision Tree, we mainly consider the original federated architecture proposed by Google, and hence our claims and paths chosen in the tree are based on that. A next step of our work should be to broaden the scope and account for different federated solutions that have been proposed and may be of interest when considering the choice of a PET.

3.2 Secure Multi-Party Computation (MPC)

MPC also considers multiple parties that collaboratively wish to evaluate a function, exemplarily to train a machine learning model or the intersection of records in two databases. The promise of MPC is that, from a functional perspective, it is indistinguishable from an ideal trusted third party who receives the data from all parties, performs the computation and returns the result (Lindell 2020; Yao 1982). A direct consequence of this functionality is that a party's data is in no form revealed to any other party apart from what can be deduced from deliberately shared information (e.g. result of the computation).

MPC has been investigated in the academic world for several decades. In contrast to federated learning, MPC is a set of cryptographic techniques that focus on mathematically verifiable security guarantees and usually achieves it at the cost of (considerably) higher system requirements. The first MPC solutions were too involved to be applied to real-world challenges, but advancements in cryptography and computing and networking capabilities have reached a point where the several MPC solutions have been applied successfully. Currently, MPC solutions are tailored to a specific problem. This indicates that it is relatively time consuming to implement MPC solutions, but also implies that MPC solutions can be used only for the purpose that they were designed for. From a legal or compliancy perspective, this can be quite advantageous.

There are MPC protocols for various security models. Some of them only prevent honest parties to deduce information that they should not obtain; others additionally provide security against colluding (or hacked) parties and even parties that maliciously deviate from the protocol. The different security models are briefly described in Appendix A. To be noted is that the guarantees given by an MPC protocol, e.g. no data is revealed to any other party even if he acts malicious, always relate to the computation phase and not the result. So any information that can be gained from the output result is no longer secure. For example, if the intersection of two databases is securely computed and the result is revealed, then surely the result reveals membership of all records in the intersection. That said, in an MPC protocol different choices can be made regarding what happens with the output result – for instance, it may be that only one party is allowed to see the output, or that the output is only revealed if it satisfies certain criteria.

There exist different types of MPC protocols. Two important categories are Homomorphic Encryption and Secret Sharing. Within these protocols there are different roles which each party will play. These can be best described as:

- *Input party*
  A party that inputs data in the computation (encrypted, or not)
- *Compute party*
  A party that does the computation (In an encrypted domain, or not)
- *Output party*
  A party that receives the results from the joint computation

Each role (or combination of roles) needs to be separately considered when researching the best solution and its complications.

3.2.1 Homomorphic Encryption

A homomorphic encryption protocol uses a public and a private key. The public key is known to everyone and can be used by the parties to encrypt data. The encryption protects the underlying data and can only be lifted using the private key. All parties can therefore encrypt their own data and share the encrypted data with one another. The homomorphic property ensures that analyses can also be performed on the encrypted data. Only when the analyses have been performed is encryption lifted using the private key. During all of the intermediate steps, the data remains encrypted and no secrets are revealed. Please note that the party which holds the private key can decrypt all of the encrypted data. This key is therefore very powerful and a potential privacy risk. It is crucial that this key be handled correctly. A common approach is to divide the private key into pieces so that no single party has access to the entire key.

3.2.2 Secret Sharing

Secret-sharing involves dividing secret data into pieces (shares) in such a way that a single share does not contain any information on the secret data. The shares can therefore be spread among the participating parties without revealing the secret data. Ironically, secret-sharing does not mean that a secret is shared with other parties. All parties distribute the shares of their own input data in this way. The second step is to carry out the analysis. Instead of one party performing the analysis of all data, all parties perform the same analysis of the shares they received from the other parties. All parties receive a different outcome from which nothing meaningful can be derived. Only when the parties combine these local, intermediate results can the analysis result be revealed. This is the third and final step of the MPC approach. This three-step approach is also called the share-compute-reveal approach.

3.3 Differential Privacy (DP)

Differential privacy is a mathematical framework that limits the amount of information about the input data that can be deduced from the result of a computation. The protection that it provides thus focusses on the output of a computation rather than the computation itself, which sets it apart from federated learning or MPC.

As discussed, neither federated learning nor MPC prevents parties from learning something that can be deduced from the result of the computation. For example, knowing the average annual salary of your department in 2018 and in 2019, combined with the fact that there was only one change in the department members, allows you to deduce the salary of the newest colleague. Differential Privacy strategically introduces specific mathematical uncertainty (noise) somewhere in a computation such that, given the (perturbed) result, it is impossible to make high-confidence deductions about the data of an individual (Dwork 2008). That is, if the average annual salary is computed with a differentially private mechanism, we will only be able to deduce a range of salaries that *probably* includes the actual salary of our newest colleague. However, the addition of noise to the computation (to improve privacy) often reduces the usefulness of the outcome (e.g., reduced accuracy). A tuning parameter allows the user to trade-off privacy and quality of the model.

The above example only exhibits part of the power of differential privacy. When a differentially private mechanism is used to publish several statistics about the department then it is impossible for someone outside the department to deduce with certainty whether some individual works in the department. That is, differential privacy protects against membership inference. What makes differential privacy stand out is that this protection is independent of any background knowledge available to the attacker. Many other anonymizations and pseudonymization techniques fail if the attacker has access to other (public) databases whose information can be used to infer details from the published data that were supposed to be protected.

To summarize, differential privacy limits the information about input data that can be extracted form a computation output. An attractive property of differential privacy is that no amount of post-processing or background knowledge can reduce the privacy that is achieved by a differential privacy mechanism. Although it may be hard to show that a protocol is differentially private, the reward of doing so is that the protocol provides a strong protection against reconstruction, membership inference and background knowledge attacks - ideally without jeopardizing the model accuracy.

3.4 Trusted Secure Environment (TSE)

A trusted secure environment (TSE) is also known as a trusted execution environment (TEE) or a secure enclave. A TSE is a set of software and hardware features that provide an isolated execution environment to enable strong security guarantees for applications running in the TSE (Rashid 2020). Specifically, TSEs can provide confidentiality, integrity, and attestation. They enable a program to run secure computations over confidential data while providing strong isolation from other applications, the operating system, and the host (Brandão et al. 2021). TSEs establish an isolated execution environment that runs in parallel with a standard operating system, its aim is to defend sensitive code and data against privileged software attacks from a potentially compromised operation system. Data is stored in the TSE, where it is impossible to view the data or operations performed on it from outside, even with a debugger. The TSE ensures only authorized code can access and compute on the data. The TSE can for example be used to protect the data once it is on the device: while the data is protected during transmission by using encryption, the TSE protects the data once it has been decrypted on the device.

3.5 Trusted Third Party (TTP)

The traditional means of performing analyses on sensitive data from multiple stakeholders concerns the involvement of a trusted third party (TTP). Such a TTP collects all input data necessary for the analyses, and provides output back to the involved parties. In this case, the party is aware of all of the assets and is responsible for privacy and confidentiality of the data – data processing agreements should be made between the data input parties and the TTP. Note: TTP is not a PET, but does appear in the decision tree as one of the options. In certain situations, a TTP may still be the most feasible solution. Also hybrid situations are possible, in which a TTP no longer has access to the data itself, but still can play a role in performing the computation done via PETs.

3.6 Synthetic Data Generation (SGD)

Synthetic data generation as a Privacy-Enhancing Technology (PET) refers to a family of techniques designed to create artificial datasets that preserve the statistical properties and utility of the original sensitive data, while minimizing or eliminating the risk of exposing individual records. These methods aim to replicate patterns, correlations, and distributions found in real data as loyally as possible without retaining any direct identifiers or exact copies of original entries. By doing so, synthetic data enables organizations to share, analyse, and develop models in a privacy-compliant manner, supporting tasks such as machine learning, testing, and research without compromising confidentiality or regulatory requirements. It is possible to generate Synthetic Data for a wide range of modality, for example tabular data (arguably the most explored) as well as images and text. SDG is an active research field and modern approaches leverage a range of techniques, including **deep learning models** (i.e. Generative Adversarial Networks (GANs) and Variational Autoencoders (VAEs)), as well as **statistical and Bayesian approaches** (like copulas, probabilistic graphical models, and differential privacy-based methods).

4 Ethical, legal, social, business implications and considerations

The use of PETs goes well beyond technical choices. Their implementation brings along ethical and societal implications, as well as legal and business considerations, all of which influence how organizations handle data responsibly and build trust. Understanding these dimensions helps ensure that PETs are not only compliant with laws and regulations but also aligned with broader values such as fairness, transparency, and accountability.

Here, we build on the ELSA concept, which stands for Ethical, Legal, and Societal Aspects [5]. This framework provides a structured way to reflect on the broader implications of technologies beyond their technical function. To capture the organizational realities of PET adoption, we complement ELSA with business considerations, recognizing that sustainable and responsible implementation also depends on strategic and economic factors. In sum, this section aims to highlight these interrelated aspects and to encourage a balanced, interdisciplinary perspective on how PETs can be effectively and responsibly applied in practice.

Note that the tree itself only considers some legal considerations. We discuss these first. In the future, the tree may account for other aspects.

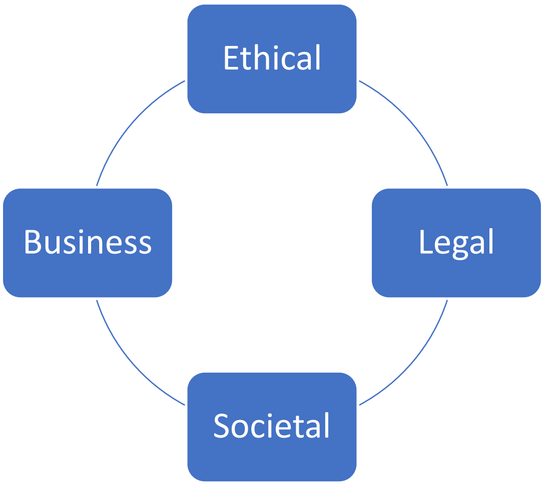

4.1 Legal considerations

PETs can be used for collaborative analysis of data while guaranteeing data confidentiality. PETs can be applied to both enhance the protection of data privacy and increase confidentiality of trade secrets or any other type of confidentiality that needs to be upheld. The identification of the type of data you want to apply the PET technologies to is crucial for the determination of the legal framework the PET technology has to operate in. When using personal data this legal framework would be the applicable data protection law, such as the GDPR and respective national privacy laws. When using business confidential information this may be a different framework, such as competition law or (internal) rules and regulations with regard to trade secrets or intellectual

[5] (Veenstra et al. 2021)

property. Also, there may be other legislation that may apply, e.g., with regards to Machine Learning the Artificial Intelligence Resolution from the European Commission. In the following paragraphs we dive into the legal framework of the GDPR.

Disclaimer

The Decision Tree only includes legal considerations at certain specific decision markers. Neither the tree itself nor the current document should be taken as legal advice. Therefore, it is strongly advised to consult your legal department at an early stage when considering the use of PETs. Even more, it is highly recommended to discuss the different outcomes of the Decision Tree with your legal department.

4.1.1 Personal data

The processing[6] of personal data within the EU and EEA requires compliance with the General Data Protection Regulation (GDPR) and national legislation of EU member states. GDPR defines personal data as any information relating to an identified or identifiable natural person. Information about an identified natural person means information about someone who's identity is already known. Information about an identifiable natural person means information about someone who's identity can be derived directly or indirectly by that information in particular by reference to an identifier such as a name, an identification number, location data, an online identifier or to one or more factors specific to the physical, physiological, genetic, mental, economic, cultural or social identity of that natural person.

Within personal data there is the distinction between regular personal data and special categories of personal data revealing racial or ethnic origin, political opinions, religious or philosophical beliefs, or trade union membership, and the processing of genetic data, biometric data for the purpose of uniquely identifying a natural person, data concerning health or data concerning a natural person's sex life or sexual orientation and personal data relating to criminal convictions and offences. GDPR restricts the processing of special categories of personal data and requires additional safeguards when processing such data or even prohibits such processing. The use of PETs in the processing operations of such data can be extra helpful when trying to tackle compliance issues while processing this kind of data.

4.1.2 Principles and legal grounds

The application of PETs are meant to enhance the privacy of lawfully processed data. Hence, every processing must have a lawful basis regardless the use of a PET. The GDPR states that there are six lawful bases that can be used to justify the processing (Table).

1) Consent	The data subject has given consent to the processing of his or her personal data for one or more specific purposes;

[6] The term 'processing' is an umbrella term used in the GDPR and means any operation or set of operations which is performed on personal data or on sets of personal data, whether or not by automated means, such as collection, recording, organisation, structuring, storage, adaptation or alteration, retrieval, consultation, use, disclosure by transmission, dissemination or otherwise making available, alignment or combination, restriction, erasure or destruction

2) Contractual obligation	Processing is necessary for the performance of a contract to which the data subject is party or in order to take steps at the request of the data subject prior to entering into a contract;
3) Legal obligation	Processing is necessary for compliance with a legal obligation to which the controller is subject;
4) Vital interests	Processing is necessary in order to protect the vital interests of the data subject or of another natural person;
5) Public interest	Processing is necessary for the performance of a task carried out in the public interest or in the exercise of official authority vested in the controller;
6) Legitimate interests	Processing is necessary for the purposes of the legitimate interests pursued by the controller of by a third party, except where such interests are overridden by the interests or fundamental rights and freedoms of the data subject which require protection of personal data, in particular where the data subject is a child. [7]

Table 1 Six lawful grounds following from Article 6 (1) GDPR

Once the question regarding legal basis is answered, be sure to also answer the question whether the processing relates to specific national data protection legislation that goes beyond the GDPR. For example, is it legally allowed for the data to be shared or even leave the premises when it concerns medical data (medical confidentiality)?

Next, GDPR requires that the processing of personal data must be in accordance with the basic data protection principles. These principles include: transparency, lawfulness, fairness, purpose limitation, data minimisation, accuracy, storage limitation, integrity and confidentiality, accountability and the principle of Data Protection by Design and Default (DPbDD). These principles are governed by the overarching notion of proportionality. This means that the processing is proportionate to the purposes of the intended and that these purposes are obtained in the least intrusive way or most privacy-friendly manner. The use of PETs can help make the processing more proportionate. Example: the processing of personal data by a municipality is needed to make a policy to help multi-problematic households (criminality, poverty, debt, low income, etc.). Data from different databases is needed to know how big this problem actually is. Sharing of the data is difficult, risky and would entail a disproportionate infringement of fundamental rights and freedoms of all citizens of a particular municipality. The use of PETs can help reduce the infringement of individual privacy rights

DPbDD is the obligation to organisations to implement appropriate technical and organisational measures which are designed to implement the before mentioned data protection principles and to

[7] Regulation (EU) 2016/679 General Data Protection Regulation (GDPR), European Parliament and the Council (27 April 2016) art. 6

integrate the necessary safeguards into the processing in order to meet the requirements of the GDPR and protect the rights and freedoms of people. A technical or organisational measure and safeguard can be anything from the use of advanced technical solutions, such as PETs to the basic training of personnel. Making use of PETs is therefore a way to implement such technical measures in your processing operations and can thus be helpful to realise compliance with these principles, specifically the DPbDD principle. However, once you have decided to make use of a PET this does not automatically guarantee complete compliance with the DPbDD principle, let alone GDPR compliance as a whole. Be aware that use of PETs may entail risks of using too much data to maximise the (learning) process which can impede the abovementioned principles such as purpose limitation. These and other risks need to be categorised and addressed from a technical and a legal perspective in a data governance document (for example a joint controller agreement) along with a Data Protection Impact Assessment (DPIA), see Sections 4.1.3 and 4.1.4. In the table in Section 4.2 an overview is given of the data protection principles in relation to the possibility of PETs as useful measures to meet compliance with abovementioned GDPR principles.

Only in case no personal data is processed – for example if it should concern data of legal persons, or data can be anonymized before processing, the GDPR does not apply.

4.1.3 Controller and processor

Processing personal data can entail risks for people. These risks may result in infringements with their fundamental rights and freedoms, such as discrimination, exclusion, stigmatization and loss of control over owns data. Compliance with the before mentioned data protection principles are meant to mitigate these risks. Besides, organizations which process personal data not in accordance with applicable legislation may get hefty fines and experience reputation loss. GDPR identifies multiple actors that have compliance responsibilities and obligations. Two of the most important actors are the ‘controller’ and the ‘processor’. If you determine the purposes (which data is processed to what end and on what grounds) and the means (how the processing will take place) of the processing operation in a decisive way, you are a controller. In a situation where multiple parties determine together purposes and means, these parties are joint controller under the GPDR. If your involvement in the processing operations is limited to technical service and technical support on behalf of the controller and you don't determine yourself the purposes and the means of the processing, you are a processor under the GDPR.

4.1.4 Data Protection Impact Assessment

One of the obligations of the controllers is to assess the risks that may rise when personal data is processed. Commonly this is done by performance of data inventory. Sometimes a Data Protection Impact Assessment (DPIA) has to be performed (article 35 GDPR). Part of the DPIA is to conclude the technical and organisational measures that need to be taken to mitigate the identified risks in the assessment. In this part of the DPIA process PETs may be the solution to some of the identified risks. The controller or joint controllers, rather than a processor, are the organisation(s) responsible for choosing the PET that is going to be used. PETs can therefore be a valuable asset when mitigating risks and fulfil its obligations. The decision tree tool can therefore be useful when performing a DPIA.

4.1.5 GDPR principles in relation to PETs

In the schedule below we provide insight how the PETs that are suggested in the decision tree can help contribute to the fulfilling a data protection principle set out in the GPDR. We also give insights on how this relates to the use of the Decision Tree and what expertise you might need to consult in your organisation if you wish to proceed in applying a PET.

Principle/ obligation	What is it?	Can a PET[8] help? Yes/No/Partially	Expertise
Proportionality	The processing is proportionate in relation to the intended goals and is done in the least privacy intrusive way.	Yes, classic processing operations that are needed to achieve legitimate goals may seem disproportionate. PETs in general can help solve these problems.	Technical + legal
Lawfulness	One of the main principles of the GDPR is that the processing is lawful (article 5 GDPR). In order to assess if the processing is lawful, one of the legal ground for processing personal data (article 6 GDPR) should apply	No, neither suggested PET will create a legal ground such as consent, legal obligation or the performance of a task of public interest.	Legal
Fairness	One of the main principles of the GDPR is that the processing is fair (article 5(1)(a) GDPR)	Partially, MPC and Federated Learning can be used to restrict use and technically prevent data is used unintentionally for purposes that negatively impact an individual, which is an important part of fair processing. But especially SDG is applicable; it is even mentioned explicitly in the Data Strategy.	Technical + legal + security
Transparency	Inform people how personal data is collected, used, consulted or otherwise processed and to what extent the personal data are or will be processed. (article 5(1)(a) GDPR, but also 12-14 GDPR).	No, the suggested PETs do not in itself enhance compliance with GDPR or give people control over the data.	Technical + legal + security
Purpose limitation	Personal data collected for specified, explicit and legitimate purposes and	Partially/Yes, MPC and Federated Learning can be adopted to restrict access and	Technical + legal + security

[8] PET has the meaning of the PETs used in the decision tree: MPC (secret sharing and homomorphic encryption), federated learning, trusted secure environment, differential privacy.

	not further processed in a manner that is incompatible with those purposes	prevent data is used unintentionally for purposes that are out of scope i.e. business rules of the MPC. Trusted secure environment can help to limit the amount of data to be processed specifically for the intended purposes. Note that results from a PET analysis are not controlled by PETs (anymore) and can in principle be used for other (unintended) purposes. For each newly considered purpose, a new legal assessment needs to be done.	
Data minimisation	Processing personal data is adequate, relevant and limited to what is necessary in relation to the purposes for which they are processed.	Yes, MPC, Federated Learning and trusted secure environment can allow you to precisely identify the data to be used for the intended outcome.	Technical + legal + security
Accuracy	Personal data is processed accurately and, where necessary, kept up to date; every reasonable step must be taken to ensure that personal data that are inaccurate, having regard to the purposes for which they are processed, are erased or rectified without delay.	Yes, MPC allows you to develop specific business rules to double check the algorithmic calculations. No, DP and SGD in general will (intentionally) result in less accurate outcomes when processing personal data in order to safeguard output privacy.	Technical + legal + security
Storage limitation	Personal data is kept in a form which permits identification of data subjects for no longer than is necessary for the purposes for which the personal data are processed; personal data may be stored for longer periods insofar as the personal data will be	Partially. PETs do not directly contribute to storage limitation, however, often PETs offer the opportunity to not copy personal data but to process the data in memory. This results in no or very short storage periods, other than archiving purposes. Additionally, Synthetic Data Generation (SDG) can be	Technical + legal + security

	processed solely for archiving purposes in the public interest, scientific or historical research purposes or statistical purposes.	deployed for data retention if the original has to be removed. Features such as automated deletion are not result of using PETs and can also be implemented with non-PET AI or other algorithms.	
Integrity and Confidentiality	Personal data processed in a manner that ensures appropriate security of the personal data, including protection against unauthorised or unlawful processing and against accidental loss, destruction or damage, using appropriate technical or organisational measures.	Partially, all PETs contribute very strongly to confidentiality. FL and MPC keeps data confidential without leaving the premises. Trusted secure environment keeps data confidential and integer while leaving the premises. SGD with no link to the original dataset can guarantee that no individual persona data is further processed in the tasks, thus lowering confidentiality breaches. Differential privacy makes it possible to make analysis on data sets while withholding information about the individuals in the dataset. Integrity is only slightly improved with MPC and FL, since no data needs to be copied to another trusted third party, which results in new integrity risks.	Technical + legal + security
Accountability	The controller shall be responsible for, and be able to demonstrate compliance with, the data protection principles.	Yes, using PETs will make it possible to proof compliance with the GDPR data protection principles, most importantly confidentiality. Often when using PETs, there are joint controllers with different responsibilities, instead of one controller with all responsibilities. This is especially true for techniques like MPC, and may be less	Technical + legal + security

		straightforward with SDG and FL.	
Data Protection by Design and Default (DPbDD)	Taking into account the state of the art, the cost of implementation and the nature, scope, context and purposes of processing as well as the risks of varying likelihood and severity for rights and freedoms of natural persons posed by the processing, the controller shall, both at the time of the determination of the means for processing and at the time of the processing itself, implement appropriate technical and organisational measures, such as pseudonymisation, which are designed to implement data-protection principles, such as data minimisation, in an effective manner and to integrate the necessary safeguards into the processing in order to meet the requirements of this Regulation and protect the rights of data subjects.	Partially, PETs can be employed as a measure in accordance with the DPbDD requirements if appropriate in a risk based approach. PETs in themselves do not necessarily cover the GDPR compliance as a whole or DPbDD entirely.	Technical + legal + security
International transfer (outside EEA)	Transfer of (meta) data from one country to another for research purposes.	Partially. MPC, HE and FL keep data confidential, but the encrypted or aggregated data being shared in these PETs can still be considered personal data in some cases. In case of an international transfer to a country outside the EEA faces challenges regarding the continuity of the same level of data protection, the PETS might help as supplementary	Legal + Security

		measures to overcome these challenges.	

4.2 Ethical implications

PETs, like other emerging technologies, raises important ethical considerations beyond legal compliance. One key concern is the intended use of the outcomes generated by PET-enabled computations. Even though PETs can limit the exposure of sensitive data, they cannot prevent the misuse of analytical results. Improper application could include using insights for questionable product or service development, targeted marketing that unfairly profiles clients or data subjects, personal gain, or the creation of ethically questionable systems. Organizations must therefore ensure that clear guidelines and ethical principles govern the use of outcomes, recognizing that the technology itself cannot fully mitigate fundamental ethical objections.

Ethical considerations also extend to the PETs themselves and the security of their implementation. Insufficient attention to security risks can result in unintended exposure of personal data, undermining the ethical premise of the technology. To address this, organizations should formalize data governance through agreements that define the integrity and use of input data, and establish shared responsibilities among all participating parties. For high-risk data processing activities, measures such as a data privacy impact assessment should be considered the minimum requirement. By embedding governance, oversight, and security measures alongside PET deployment, organizations can better align their collaborative analytics initiatives with ethical standards while responsibly managing potential risks.

4.3 Societal implications

PETs carry societal implications that are broadly similar to those observed with other emerging technologies, such as artificial intelligence, cloud computing, or big data analytics. Like these technologies, PETs can influence economic activity, innovation patterns, and competitive dynamics, as organizations gain the ability to extract insights from data in ways that were previously difficult or impossible. They can also impact the workforce, creating new roles in data management, analysis, and collaborative governance, while requiring broader digital literacy across participating organizations. In general, PETs share the common societal considerations of technology adoption, including questions around access, fairness, equity, and potential unintended consequences of widespread deployment.

What sets PETs apart, however, is their built-in focus on data protection and privacy, which directly addresses a key societal concern in the modern digital environment. By enabling joint analysis without exposing individuals' sensitive information, PETs reduce risks associated with data misuse, breaches, or surveillance, supporting public trust in organizations and collaborative ecosystems. This protective layer can foster greater willingness to share data for research, public health, or industry-wide insights, thereby unlocking societal benefits that might otherwise be constrained by privacy concerns. In this sense, PETs not only provide the usual societal advantages and challenges of technological innovation but also offer a built-in mechanism to safeguard privacy, making them a unique tool for responsible data-driven collaboration.

Practical examples highlight how PETs can deliver societal benefits while mitigating privacy risks. For instance, in healthcare, multiple hospitals can use federated learning to develop predictive models for cancer prediction without sharing patient data (Pati et al. 2022), enabling better public health responses while protecting individual privacy. In the financial sector, banks can collaborate via secure multi-party computation to detect fraud patterns without exposing sensitive customer data (Bogdanov et al. 2015), enhancing security across the system. Similarly, city governments can leverage differential privacy when releasing mobility or census data for research (Abowd 2018), ensuring that useful insights reach urban planners and researchers without compromising citizens' identities. These examples demonstrate that PETs not only enable valuable cross-organizational analytics but do so in a way that preserves public trust and minimizes societal risks associated with data exposure.

4.4 Business considerations

As denoted, the focus of this work is on application of PETs in the context of inter-organizational data analysis – i.e., multiple organizations (which could also be part of one larger organization) aim to perform joint analyses on their data sets. In this context, we highlight considerstaions that are relevant from a managerial, business perspective.

4.4.1 Strategic alignment

Before entering into a PET-enabled collaboration, organizations must establish a clear strategic alignment with both the initiative and their prospective partners. This begins with defining the overarching business objectives that the collaboration is intended to achieve, such as improving predictive models, performing benchmarking across organizations, or generating actionable insights that would be impossible to obtain independently. Equally important is assessing how each organization's data assets complement one another and whether the combined insights can create tangible business value. PETs enable this alignment by allowing partners to leverage collective data for analysis without sharing sensitive information, making it possible to pursue joint goals while protecting each organization's proprietary assets. Organizations must also ensure that the collaboration fits into their broader strategic roadmap and contributes to long-term growth (via a sustainable collaboration), competitive positioning, or innovation initiatives.

Competition versus collaboration, coopetition?

Engaging in PET-enabled collaborations often requires organizations to embrace a form of "coopetition," where they cooperate with other organizations — which could include competitors — to achieve shared analytical benefits while maintaining their own competitive advantage (Woolley 2023). This approach allows participants to jointly develop insights or predictive models that none could create alone, but it requires a careful balance between collaboration and competition. Organizations must evaluate the potential risks and rewards of working alongside partners who may have overlapping interests, ensuring that shared insights enhance the collective outcome without undermining proprietary advantages. Specific PETs – such as federated learning – serve as a critical enabler in this context, as they allow partners to participate in shared analysis while preserving control over their sensitive data, thus creating a framework where coopetition is both feasible and strategically valuable (Barbereau et al. 2025). Underlying this framework indeed needs to lie some kind of collaborative business model; one that allows for sustained profitability of a technological cooperation.

4.4.2 Governance and trust

Before collaboration begins, it is essential to establish governance structures and trust mechanisms that will guide the partnership. Effective governance includes defining decision-making processes, roles, responsibilities, and clear rules (Weber et al. 2009). Trust between participants is reinforced not only through PETs — which technically protect data and computations — but also through agreements on operational practices, transparency of processes, and fair allocation of value.

Organizations should determine how contributions will be recognized and how benefits will be distributed to ensure that all parties are motivated to participate fully. By addressing governance and trust upfront, organizations can reduce uncertainty, prevent conflicts, and set the stage for a smooth, productive collaboration (for FL see Barbereau et al. 2025).

Liability and responsibility

Establishing clear links of liability, responsibility, and accountability is a critical step before initiating any collaborative setting. Even though PETs protect sensitive data during joint analysis, organizations must explicitly define who is responsible for each aspect of the collaborative process, including data preparation, model training, results validation, and decision-making based on shared insights. Clear accountability ensures that every participant understands their obligations and the consequences of errors, omissions, or misuse of the shared analytical outputs. This includes defining liability in cases where outcomes are inaccurate, incomplete, or misinterpreted, as well as responsibilities for maintaining the integrity and security of the collaborative environment. By codifying these roles and responsibilities upfront, organizations reduce ambiguity, mitigate potential disputes, and build confidence among all partners, creating a transparent and trustworthy foundation for effective and sustainable collaboration.

Risk management

Effective risk management is essential throughout the lifecycle of PET-enabled inter-organizational collaboration, beginning before engagement and continuing during the partnership. Prior to entering collaboration, organizations must identify potential strategic, operational, and reputational risks associated with joint data analysis, including uncertainties about partner reliability, data quality, and the alignment of business objectives. This phase involves establishing mitigation strategies, such as selecting trusted partners, defining governance structures, and setting clear expectations for contribution and outcome sharing. During collaboration, risk management shifts to monitoring and responding to emerging challenges, including inconsistencies in data inputs, deviations from agreed processes, or unexpected analytical outcomes. Maintaining active oversight, transparent communication, and mechanisms to address errors or disagreements ensures that risks are managed in real time and that the collaboration remains productive and trustworthy. By integrating risk management into both planning and execution phases, organizations can safeguard value, maintain confidence among partners, and enhance the resilience and sustainability of the collaborative effort.

5 Levels of data-sharing risk mitigation

If personal data is involved in a collaborative solution, then usually (processed) personal data will also need to leave the premises in some form. PETs generally handle this data in some encrypted or aggregated form that greatly reduces the risk of data subject identification and data breaches. They are therefore also a measure adding to the obligation for taking sufficient technical measures preventing personal data security breaches.

PETs ensure varying levels of risk mitigation. Some technologies provide mathematical guarantees of the security that they provide and the types of attacks that they protect against. Some PETs provide protection in many scenarios, some against few. Often, however, it is hard to provide formal, generic guarantees and an assessment should be conducted for the specific challenge at hand. *It is important to be aware of these differences both from a technical and from a legal perspective; if one PET does not adhere to your constraints, there might be another one that does. Also, in the legal assessment, it is important to be aware of the security scenarios and the fact that PETs provide different levels of protection.*

To get a better feel of security scenarios, please refer to Appendix B and Appendix C.

Unfortunately, as the security requirements grow the pool of feasible technologies shrinks. If your solution requires data sharing in e.g., an aggressive, competitive environment, it might be that only few PETs satisfy your criteria. In that case it may help to explore various routes: stronger PETs, in combination with organisational measures or legal agreements that mitigate part of the risks. Perhaps a slight alteration to your proposed solution results in the exchange of less sensitive information (e.g., computations performed by another party). In the end, risks can be mitigated in various ways and keeping a wider view of the possibilities helps finding the best solution to your challenge.

6 Decision tree guide

In this section we will include disclaimers and explanations about the decision tree nodes when needed.

6.1 Data sources independence

In a node we ask whether the data sources are independent of each other. Independent in this context means that the analysis can technically be performed on the data set of any of the data sources without including samples from another.

6.2 Federated Analytics/Learning

In the tree, we use the term Federated Analytics to describe statistical analysis (e.g., aggregation) that is performed in some federated manner. We use Federated Learning to signify specifically that the process of *model training* occurs in a federated manner.

6.3 Data vs Model and Output sensitivity

A dataset can be sensitive, but also a trained model can be sensitive. When it comes to a dataset, it corresponds to the data itself being under protection. This can be due to different reasons, including privacy, commercial reasons or simply due to other policies of the organization owning the data. See also Section 4.

By model sensitivity, we mean that the trained model itself is to be protected. This can occur either if there are serious concerns that the model can leak information about the data on which it was trained or if the model owner wishes to keep the model private for organizational reasons, such as commercial confidentiality.

Model sensitivity is a specific but important example of the broader concept of output sensitivity or output privacy. Often the focus of PETs is on keeping the input data, as well as the computation, hidden. But also, the output of a computation may be sensitive, specifically if it can lead back to the sensitive input data. For example, knowing the average annual salary of your department in 2018 and in 2019, combined with the fact that there was only one change in the department members, allows you to deduce the salary of the newest colleague. This closely links to the concept of output privacy, see Section 3. *Output Privacy* refers to techniques used to ensure that the output of a computation does not reveal information about the input data.

6.4 Set Intersection

The type of problem Set Intersection can be either a subproblem of any of the other problems, or a problem by itself. An example of the second scenario is when organizations wish to match their datasets without planning to perform a specific analysis per se. In this case, only the Set Intersection route shall be traversed. On the contrary, when additional analysis is intended, the tree shall be traversed twice: once for Set Intersection and once for said analysis (Machine Learning, Statistical Analysis or Synthetic Data Generation).

6.5 Sensitivity of locally computed values

If you traverse the tree for the case of Machine Learning or Statistical Analysis, it might be that you come across the question: “Are the locally computed values to be exchanged sensitive?”. This question might seem vague and/or unclear. Let us consider an extreme example of local sensitivity to outline the reasoning of this question.

Let us assume that two parties wish to calculate collaborative statistics on their data. Let us also assume that each party is aware of the number of samples the other party owns. Let us finally assume that party A only owns 2 samples. Then, we can immediately realize that if party A shares with party B the local mean and variance of a feature they own, then party B will be able to fully reconstruct that feature for both samples owned by party A.

As mentioned, this is an extreme example and in practice the risk of freely exchanging locally computed values is nuanced and often hard to quantify. That being said, various recent research papers have focused on reconstructing data from local model updates, even in far less extreme examples of local sensitivity. As a result, any party participating in data sharing need to consider such dangers and carefully plan the protocol of sharing intermediate values that may leak much more information than immediately obvious.

7 Bibliography

Abowd, John M. 2018. "The U.S. Census Bureau Adopts Differential Privacy." *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining* (New York, NY, USA), KDD '18, July 19, 2867. https://doi.org/10.1145/3219819.3226070.

Bagdasaryan, Eugene, Andreas Veit, Yiqing Hua, Deborah Estrin, and Vitaly Shmatikov. 2019. "How To Backdoor Federated Learning." *arXiv:1807.00459 [Cs]*, August 6. http://arxiv.org/abs/1807.00459.

Barbereau, Tom, Joaquin Fernandez Delgado, and Sergio Potenciano Menci. 2025. "The Governance of Federated Learning: A Decision Framework for Organisational Archetypes." *Data & Policy* 7 (53). https://doi.org/10.1017/dap.2025.10020.

Bogdanov, Dan, Marko Jõemets, Sander Siim, and Meril Vaht. 2015. "How the Estonian Tax and Customs Board Evaluated a Tax Fraud Detection System Based on Secure Multi-Party Computation." In *Financial Cryptography and Data Security*, edited by Rainer Böhme and Tatsuaki Okamoto. Springer. https://doi.org/10.1007/978-3-662-47854-7_14.

Brandão, André, João Resende, and Rolando Martins. 2021. "Hardening of Cryptographic Operations through the Use of Secure Enclaves." *Computers & Security* 108 (May): 102327. https://doi.org/10.1016/j.cose.2021.102327.

Dwork, Cynthia. 2008. "Differential Privacy: A Survey of Results." In *Theory and Applications of Models of Computation*, edited by Manindra Agrawal, Dingzhu Du, Zhenhua Duan, and Angsheng Li. Lecture Notes in Computer Science. Springer. https://doi.org/10.1007/978-3-540-79228-4_1.

Fernandez, Joaquin Delgado, Renan Lima Baima, Tom Barbereau, and Alexander Rieger. 2024. "Opportunities and Applications of Federated Learning in the Financial Services Industry." In *Decentralization Technologies: Financial Sector in Change*, edited by Gilbert Fridgen, Tobias Guggenberger, Johannes Sedlmeir, and Nils Urbach. Springer Nature Switzerland. https://doi.org/10.1007/978-3-031-66047-4_11.

Fernández, Joaquín Delgado, Sergio Potenciano Menci, Chul Min Lee, Alexander Rieger, and Gilbert Fridgen. 2022. "Privacy-Preserving Federated Learning for Residential Short-Term Load Forecasting." *Applied Energy* 326 (November): 119915. https://doi.org/10.1016/j.apenergy.2022.119915.

Konečný, Jakub, H. Brendan McMahan, Daniel Ramage, and Peter Richtárik. 2016. "Federated Optimization: Distributed Machine Learning for On-Device Intelligence." *arXiv:1610.02527 [Cs]*, October 8. http://arxiv.org/abs/1610.02527.

Lee, Chul Min, Joaquin Delgado Fernandez, Sergio Potenciano Menci, Alexander Rieger, and Gilbert Fridgen. 2023. "Federated Learning for Credit Risk Assessment." Paper presented at Hawaii International Conference on System Sciences. *Proceedings of the 55th Hawaii International Conference on System Sciences*, January 3. https://hdl.handle.net/10125/102676.

Li, Tian, Anit Kumar Sahu, Ameet Talwalkar, and Virginia Smith. 2019. "Federated Learning: Challenges, Methods, and Future Directions." *arXiv:1908.07873 [Cs, Stat]*, August 21. http://arxiv.org/abs/1908.07873.

Li, Zhaorui, Zhicong Huang, Chaochao Chen, and Cheng Hong. 2020. "Quantification of the Leakage in Federated Learning." *arXiv:1910.05467 [Cs]*, March 11. http://arxiv.org/abs/1910.05467.

Lindell, Yehuda. 2020. *Secure Multiparty Computation (MPC)*. No. 300. https://eprint.iacr.org/2020/300.

McMahan, H. Brendan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. 2017. "Communication-Efficient Learning of Deep Networks from Decentralized Data." *arXiv:1602.05629 [Cs]*, February 28. http://arxiv.org/abs/1602.05629.

Pati, Sarthak, Ujjwal Baid, Brandon Edwards, et al. 2022. "Federated Learning Enables Big Data for Rare Cancer Boundary Detection." *Nature Communications* 13 (1): 7346. https://doi.org/10.1038/s41467-022-33407-5.

Rashid, Fahmida Y. 2020. "The Rise of Confidential Computing: Big Tech Companies Are Adopting a New Security Model to Protect Data While It's in Use - [News]." *IEEE Spectrum* 57 (6): 8–9. https://doi.org/10.1109/MSPEC.2020.9099920.

Veenstra, Anne-Fleur, Liesbet van Zoonen, and Natali Helberger. 2021. *ELSA Labs for Human-Centric Innovation in Artifical Intelligence*. Dutch AI Coalition. https://nlaic.com/wp-content/uploads/2022/02/ELSA-Labs-for-Human-Centric-Innovation-in-AI.pdf.

Weber, Kristin, Boris Otto, and Hubert Österle. 2009. "One Size Does Not Fit All: A Contingency Approach to Data Governance." *J. Data and Information Quality* 1 (1): 4:1-4:27. https://doi.org/10.1145/1515693.1515696.

Woolley, Jennifer L. 2023. "Getting along with Frenemies: Enhancing Multi-Competitor Coopetition Governance through Artificial Intelligence and Blockchain." *Industry and Innovation* 30 (9): 1156–89. https://doi.org/10.1080/13662716.2023.2168519.

Yao, Andrew C. 1982. "Protocols for Secure Computations." *23rd Annual Symposium on Foundations of Computer Science (Sfcs 1982)*, November, 160–64. https://doi.org/10.1109/SFCS.1982.38.

Zhu, Ligeng, Zhijian Liu, and Song Han. 2019. "Deep Leakage from Gradients." *arXiv:1906.08935 [Cs, Stat]*, December 19. http://arxiv.org/abs/1906.08935.

Appendix A. Security scenarios

We will describe several characteristics for security models to illustrate the variety of scenarios that one can encounter.

An exhaustive list of security models would be quite long and technical, so we only list a couple of characteristics that are part of a security model. The more adversarial scenarios will often not be very likely, but they help in getting a feel of the possibilities. Be aware that organizations do not always behave malicious on their own accord – they also expose that behaviour when an attacker infiltrated (e.g. hacked) some of them and act on their behalf.

- *Security type.* The security of the PET can depend on the computing capabilities of an attacker. Some PETs are resistant against classical computers only, some against both classical and quantum computers, and some provide security against any (future) computer with unlimited computing power (information-theoretic security).
- *Conspiring organizations.* A group of organizations in a PET solution may together try to break security and infer information of another organization.
- *Untrustworthy organizations.* Often it is assumed that all organizations in a PET solution adhere to the agreements made and implement and execute the solution as agreed upon. Additional measures must be considered if this assumption cannot be made.

Appendix B. Attack vectors

PET solutions do not guarantee protection from all potential attacks at any point in the process. But what would such an attack entail? Specific types of attacks may be more serious than other types of attacks and not all of them may be equally likely to happen.

To facilitate this discussion, we include a non-exhaustive list of potential attack vectors:

- *Reconstruction attack.* Reconstruct input data from the output data.
- *Membership inference.* Find out whether a certain record (person) is present in the input or output data set.
- *Property inference*. Retrieve the value of a certain attribute of a record (person) in the input or output dataset.
- *Model poisoning or backdooring (in Machine Learning)*: temper with the training phase of the machine learning to poison the model. The poisoned model may infer detailed information of some training data or provide forced (malicious) outputs for certain inputs.
- *Infrastructure attack*: an attack that aims to weaken the infrastructure software to insert a malicious algorithm, or to weaken security checks for repetitions (e.g. to avoid reconstruction attacks) or authentication/authorization to the infrastructure.