January 2024

# America's DataHub Consortium: Privacy Preserving Technology Phase 1 – Environmental Scan

## Final Report

Prepared for

**National Center for Science and Engineering Statistics**
2415 Eisenhower Avenue, Suite W14200
Alexandria, VA 22314

Prepared by

**Zixin Nie, Rashonda Lewis, Allie Gartland-Grey, Amanda Flynn Riley**
RTI International
3040 E. Cornwallis Road
Research Triangle Park, NC 27709

Project Code ADC-PPT-PHASEI-23-N007

Project Code
ADC-PPT-PHASEI-23-N007

# America's DataHub Consortium: Privacy Preserving Technology Phase 1 – Environmental Scan

## Final Report

**January 2024**

Prepared for

**National Center for Science and Engineering Statistics**
2415 Eisenhower Avenue, Suite W14200
Alexandria, VA 22314

Prepared by

**Zixin Nie, Rashonda Lewis, Allie Gartland-Grey, Amanda Flynn Riley**
RTI International
3040 E. Cornwallis Road
Research Triangle Park, NC 27709

# Contents

The America's DataHub Consortium (ADC), a public-private partnership is being utilized to implement research opportunities that support the strategic objectives of the National Center for Science and Engineering Statistics (NCSES) within the U.S. National Science Foundation (NSF). This report documents research funded through the ADC and is being shared to inform interested parties of ongoing activities and to encourage further discussion. Any opinions, findings, conclusions, or recommendations expressed in this report do not necessarily reflect the views of NCSES or NSF. Please send questions to ncsesweb@nsf.gov. The OMB control number for this collection is 3145-0215. This product has been reviewed for unauthorized disclosure of confidential information under NCSES-DRN24-004.

# Figures

**Number** **Page**

# Tables

# Executive Summary

The CHIPS and Science Act of 2022 authorized the National Science Foundation (NSF) to establish a National Secure Data Service (NSDS) Demonstration (NSDS-D) project. To enable a wide range of perspectives and ensure that a variety of stakeholder needs are accounted for, the National Center for Science and Engineering Statistics (NCSES) partnered with federal and state government agencies as well as other organizations to help guide its implementation of the NSDS-D project. Many NSDS-D activities are being implemented and tested through the America's DataHub Consortium (ADC). These include, but are not limited to, data linkage and privacy preserving technologies (PPTs).

PPTs are defined as technologies and methodologies that mitigate privacy risks when using data that may contain identifying, sensitive, or confidential information. They attempt to strike the balance between privacy protection and data utility through controls on data systems and architecture and transformation of data. To improve understanding of the current PPT landscape, RTI International (RTI) (partnered with Flood Mason Holdings (FMH)) obtained a project agreement with ADC funded through NCSES to conduct an environmental scan of PPTs. The goal was to assess what PPTs are currently being developed, tested, and utilized across government, academia, and the private sector. Through the performance of a literature review, interviews with PPT practitioners and facilitated group discussions with practitioners and NCSES representatives, we gathered information and perspectives directly from those working in the field to answer the questions posed by NCSES regarding the current PPT landscape.

To help us navigate the broad range of technologies that are called PPTs, we developed a taxonomy to aid us in finding the ones that would best fit the intended use within a shared services environment. The technologies we identified are data privacy technologies (those that deal with relationships between data controllers and other controllers and processors) that can be classified as "hard" (technologies that do not necessarily rely on personal identifiable information (PII) to be shared) and protect both input and output privacy. These technologies include de-identification, differential privacy, synthetic data, secure multiparty computation, homomorphic encryption, trusted execution environments, privacy preserving record linkage, and federated learning. We provide details about each of these PPTs, including their definition, use cases found within literature, and their limitations within the report. Further details, including specific answers to the questions posed by NCSES, are provided in Appendix B.

We also provide information that was gathered from the interviews and facilitated group discussions that we conducted with PPT practitioners, detailing the common themes that arose from those conversations, as well as differences in perspective between federal PPT practitioners and nonfederal PPT practitioners. The focus of the interviews we conducted was not to delve deeply into the technical aspects of PPTs but rather to try and obtain information about their experiences working with PPTs, such as what worked well within their implementation process, barriers and challenges they faced, and any lessons they learned that could help inform future development.

The information we gathered as part of the environmental scan enabled us to create a framework for assessing the technological maturity of the PPTs that we detail within this report. The framework measures maturity according to three dimensions: (1) level of standards setting, (2) ease of use, and (3) public trust. This framework provides a snapshot of the current state of these technologies, providing decision-makers with a guide as to what technologies could be sufficiently mature to test for implementation within a shared services environment.

PPT practitioners with whom we engaged expressed a few next steps that they would like to see to help further the implementation of PPTs. Taking these recommendations as inspiration, we suggest the following next steps for the development of the NSDS: (1) setting up a sandbox environment for testing PPTs, (2) creating a community of practice to foster PPT expertise, (3) exploring the impacts of PPT usage toward data governance, and (4) creating communications materials to help inform those without technical background in PPTs. Each of these steps helps to address some of the challenges that PPT practitioners expressed that they have faced and could provide benefits for the establishment of a future NSDS.

# Introduction

Access to sensitive data is governed by laws, policies, and standards designed to protect the privacy of data subjects. Privacy preserving technologies (PPTs) attempt to strike the balance between privacy protection and data access and utility by transforming data or protecting data via systems and architecture. These technologies include, but are not limited to, de-identification, differential privacy, synthetic data, secure multiparty computation (sMPC), homomorphic encryption, trusted execution environments (TEEs), privacy preserving record linkage (PPRL), and federated learning. Common PPT use cases include application testing and data analysis, financial transactions, use and exchange of electronic health record data, and multiparty data transfer. While PPTs hold great promise for moving the needle on data accessibility, organizations face challenges as they adopt or consider adoption of these technologies, including inadequate understanding of privacy technology risks and benefits, lack of consensus standards, inconsistent definitions and taxonomies, and lack of clarity around regulatory compliance.

PPTs have been identified as a strategic priority by the U.S. federal government, with the White House Office of Science and Technology Policy putting out a call for input on advancing adoption of PPTs within the U.S.;[1] publishing a national strategy to advance PPT usage;[2] and, more recently, within the Executive Order issued by the Biden Administration prioritizing development and usage of PPTs, protecting the privacy of Americans for managing the risks presented by artificial intelligence (AI).[3] A primary focus of both the legislation and the recommendations of the Advisory Committee on Data for Evidence Building (ACDEB) is data security, ensuring that data confidentiality is maintained and that PPTs are used whenever possible to achieve this. Within the ACDEB *Advisory Committee on Data for Evidence Building: Year 2 Report*, PPTs are identified as a key component of a future National Secure Data Service (NSDS), helping to enable tiered access to federal data, working with data in its original place, and conducting PPRLs with data from state, local, territorial, and tribal jurisdictions.[4] PPTs have also been identified as a key priority by national statistical organizations worldwide, who have been working together at forums such as the United Nations (UN) to determine how they can be used within official statistics.[5]

The purpose of this study is to gather information to inform future exploration and testing of PPTs as potential shared services in support of the NSDS Demonstration (NSDS-D) project as authorized under Section 10375 of the CHIPS and Science Act of 2022. To improve understanding of the current PPT landscape, RTI International (RTI) (the project agreement holder) partnered with Flood Mason Holdings (FMH) to conduct an environmental scan of PPTs currently being developed, tested, and utilized across government, academia, and the private sector, through the performance of a literature review, engaging practitioners of PPTs through interviews and hosting facilitated group discussions inviting PPT practitioners to speak with representatives from the National Center for Science and Engineering

Statistics (NCSES). We gathered information and perspectives directly from those working in the field, fostering a sense of community between the practitioners and NCSES to inform the development of an NSDS, and answering the five questions posed by NCSES within the original solicitation.

We present the results of this environmental scan within this report, which contains information and perspectives directly from key experts and experienced practitioners that answer the five questions that NCSES has posed, helping inform the usage and testing of PPTs within NSDS. Details about the approach taken for conducting this environmental scan are provided in Appendix A.

## Background: National Center for Science and Engineering Statistics (NCSES) within the National Science Foundation (NSF) and America's DataHub Consortium (ADC)

The CHIPS and Science Act of 2022 authorized NSF to establish the NSDS-D project. Responsibility for operation of this project was given to NCSES. To enable a wide range of perspectives and to ensure that a variety of stakeholder needs are accounted for, NCSES partnered with federal and state government agencies as well as other organizations to help guide its implementation of the NSDS-D project.

**Five questions**

1. What projects and pilots are currently testing or implementing (or have previously been done) privacy preserving technologies, including (but not limited to) sMPC, synthetic data, differential privacy methodologies, homomorphic encryption, and validation servers?
2. What lessons learned are available from using PPT, including what has worked or is working and under what contexts, purposes, or various types of data users; what challenges or barriers have been discovered with PPT in using these technologies from the data provider and data user perspectives; and what potential next steps are there in implementing these technologies?
3. What do we know about how to evaluate whether a certain PPT is a good fit for a specific use case? What features of the data or users might indicate one PPT approach over another?
4. What are best practices for effective communication strategies or user training on how to conduct research or program evaluation projects that leverage these new PPT?
5. What use cases exist for each of these technologies when applied to evidence-building research, policymaking, and program evaluation?

Many NSDS-D activities are being implemented and tested through the ADC. The ADC was established in 2021 by NCSES as a public-private partnership to address key research questions through innovative solutions in a rapid and streamlined acquisitions process. Topics include, but are not limited to, data linkage and PPTs. Given this overlap with the goals of the NSDS-D project, the ADC is being leveraged to inform the requirements under the demonstration project.

## Background: Research Triangle Institute (RTI) and Flood Mason Holdings (FMH) Privacy Preserving Technologies (PPT) Project

In July 2023, RTI received ADC's award for the Privacy Preserving Technologies Phase 1: Environmental Scan Project. RTI is an independent, nonprofit research institute dedicated to improving the human condition. RTI's vision is to address the world's most critical problems with science-based solutions in pursuit of a better future. With work in more than 75 countries—tackling hundreds of projects each year to address complex social and scientific challenges on behalf of governments, businesses, foundations, universities, and other client partners . The Environmental Scan Project aligns with RTI's work to improve the movement and use of data through technology and to support decision making throughout the data lifecycle. RTI conducted the environmental scan by (1) performing a literature review to identify PPTs and expert practitioners; (2) performing outreach to practitioners; (3) conducting interviews and group discussions with expert practitioners; (4) analyzing the data obtained from the literature review, interviews, and group discussions; and (5) organizing findings into this report.  To accomplish these steps, the Project Team leveraged team member privacy, outreach, and project management expertise, RTI's organizational resources and industry outreach support from Flood Mason Holdings, LLC ("FMH"). FMH served as RTI's project partner, providing outreach support for industry PPT experts. FMH is an advisory and consulting firm targeting innovative solutions to address some of the world's most challenging problems FMH's team brings over 200 years of experience spanning over 2,000 investments in life science and technology companies.  FMH has access to leading experts across a multitude of disciplines, including healthcare, technology, and finance.

# Environmental Scan Findings

This section of the report details the information gathered from the environmental scan of PPTs. We first provide details about how we developed a taxonomy for PPTs to help us find the PPTs that would be most beneficial toward usage within shared services environments. We then provide details about each of the PPTs we identify, including their definition, use cases found within literature, and their limitations. Further details, including specific answers to the questions posed by NCSES, are provided in Appendix B. We then map the technological maturity of each PPT based on an assessment we developed from the information we had gathered. We conclude this section by providing information gathered from the interviews and facilitated group discussions that we conducted with PPT practitioners, detailing the common themes that arose from those conversations, as well as the differences in perspective between federal PPT practitioners and nonfederal PPT practitioners.

## Development of a PPT Taxonomy

One of the key results of our environmental scan is the development of a PPT taxonomy to help find the PPTs that could inform the establishment of an NSDS and be useful within a shared services environment (see Figure 1). Most existing taxonomies split PPTs into groups such as "soft" versus "hard";[6],[7] "cryptographic" versus "statistical";[8] "input privacy" versus "output privacy";[9]–[11] "altering data" versus "shielding data" versus "systems and architecture";[12] and "data obfuscation tools" versus "encrypted data processing tools" versus "federated and distributed analytics" versus "data accountability tools."[13] A paper published by Garrido et al. established a comprehensive multilayered taxonomic system that broke down PPTs by layers of the Internet of Things technology stack.[14] From the taxonomies that we examined, we developed the taxonomic system shown in the following figure to break down and classify the field of PPTs in a way that best aided us in finding the PPTs that would be useful for this project. Using this taxonomic system, we were able to find the types of PPTs that best satisfied the inclusion or exclusion criteria we established for PPTs that could best support the development of the NSDS.

**Figure 1. PPT Taxonomy**

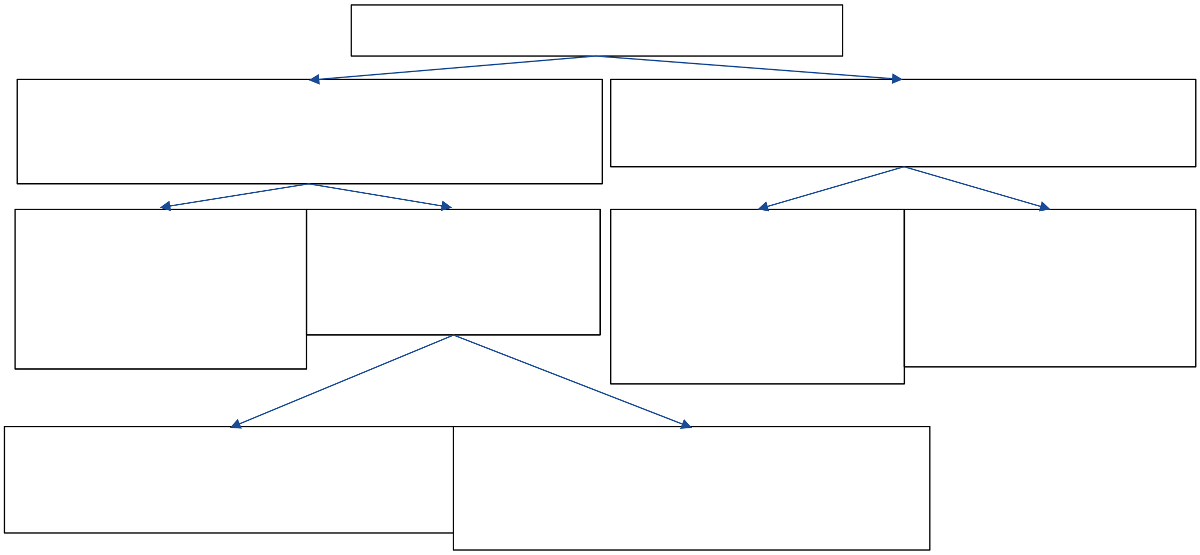

*Source: RTI based on research*

At the first layer, all PPTs can be broken down into data privacy technologies and personal privacy technologies. Data privacy technologies are concerned with protecting the privacy of data subjects within data that have already been collected. Personal privacy technologies are concerned with individuals protecting their own privacy. Data privacy technologies manage relationships between data controllers and other controllers or data processors, whereas personal privacy technologies manage relationships between data subjects and data controllers.

Personal privacy technologies can be broken down into technologies that enable transparency or obfuscation. Transparency-enabling technologies are for situations where individuals trust data controllers to manage their personal private data, but they also want restrictions based on their privacy rights. These technologies include consent management tools, privacy dashboards, and data subject rights managers. Obfuscation technologies for situations where individuals do not trust the data controllers and processors and wish to protect their personal data. The classic example of such a technology is the anonymous ballot. Digital examples of these kinds of technologies include The Onion Router (Tor) browser and public-private key encryption when used in blockchain ledgers.

Data privacy technologies can be broken down into “soft” privacy technologies and “hard” privacy technologies. Soft privacy technologies are technologies used to protect data subject privacy when the data controller and data processors can be trusted to handle personal data, either because they have the consent of the data subject or because the processing falls under a legitimate use. Most of these technologies use encryption in transit to ensure that the data are protected while they are being transferred between data controllers and processors, such as transit layer security (TLS), secure file transfer protocol (SFTP), and hypertext transfer protocol secure (HTTPS).

Hard privacy technologies are for situations where the data controllers and processors are not trusted with private data, either due to a lack of consent from data subjects, or when data controllers do not wish to reveal the private data they hold to other parties. These types of technologies can further be broken down into technologies that protect input privacy and output privacy. Technologies that protect input privacy protect the privacy of individuals when multiple parties are collaborating in computation and analyses, enabling them to share their data and perform analyses without seeing private information that they do not already control. Most of these technologies involve encryption and cryptography to obfuscate data, and many of them enable computations on encrypted data. Technologies that protect output privacy protect the privacy of individuals when releasing data to other parties or to the public. Most techniques that protect output privacy utilize statistical methods and transformations to the data to manage the risk of identifying data subjects.

Utilizing this taxonomy, we identified that the PPTs that could be useful for a future NSDS are data privacy technologies that can be classified as hard that can be used to protect input and output privacy. Table 1 gives the types of PPTs that we have found that fall into these groups.

**Table 1. Types of PPTs, by Input and Output Privacy**

| Input Privacy | Output Privacy |
|---|---|
| ▪ Secure multiparty computation (sMPC)<br>▪ Homomorphic encryption<br>▪ Trusted execution environment (TEE)<br>▪ Federated learning<br>▪ Privacy preserving record linkage (PPRL) | ▪ De-identification<br>▪ Noise addition (differential privacy)<br>▪ Synthetic data |

## PPT Technological Maturity Assessment

Based on the information we gathered as part of this environmental scan, we created a framework for assessing the technological maturity of the PPTs that we detail within this report. This framework was inspired by the framework developed by the UN Big Data Working Group within their *Handbook on Privacy Preserving Computation Techniques*,[15] although the UN framework has a different purpose. This framework tries to provide a snapshot of the current state of these technologies, providing decision-makers with a guide as to what technologies are sufficiently mature to test for implementation within a shared services environment.

The framework is built on the following three dimensions:

1. **Level of Standards Setting:** Standards set by major standards-setting bodies such as ISO/International Electrotechnical Commission (IEC), Institute of Electrical and Electronics Engineers (IEEE), the National Institute of Standards and Technology (NIST), and American National Standards Institute (ANSI) are a sign of technological maturity because these standards arise from common practices and processes developed through experience and consensus from expert practitioners. The standards-setting process promotes maturation of technologies through bringing together experts within the field to find and discuss aspects of a technology that require improvement, creating strategic direction for the further development of a technological field. Where a technology sits within the standards-setting process is thus indicative of the maturity level of the technology.

   For this dimension, we look at the standards that have been set for each PPT, where the PPT sits within the standards-setting process, and what aspects of the PPT still require standardization as the criteria for measuring technological maturity.

2. **Ease of Use:** Technologies that are more mature tend to be easier to use because efforts to make the technology user friendly tend to come toward the end of technological development lifecycles. We therefore measure PPT ease of use on the aspects of availability of commercial tools, level of expertise needed to implement and use the technology, and the amount of customization and optimization required by the end user to determine the level of maturity on this dimension.

3. **Public Trust:** Public trust can be an indicator of technological maturity because technologies that are more mature tend to be better known and more trusted by the public. Trust is an especially important factor for PPTs because the public needs to trust that a technology can protect their privacy for it to be broadly adoptable.

   We measure public trust with three criteria, using a scale of high, medium, and low. The first criterion is the level of understanding the public has about how the PPT operates (i.e., if you ask a layperson, what they tell you about the PPT). The second criterion is how much knowledge and scrutiny the public has on the PPT. The third is the level of difficulty in informing the public about how the PPT works, which includes evaluating the availability of materials that help explain the PPT.

Note that these dimensions are intertwined (i.e., standardization promotes ease of use and increases public trust, public scrutiny is a required step in setting standards, and tools that are easier to use and understand are better trusted by the public).

The maturity scale is broken down into three levels: emerging, maturing, and mature. Emerging technologies are those PPTs that generally are in the early phases of standards setting, require significant technical expertise to implement and use, and are relatively unknown to the public. Maturing technologies have some standards set, though there remain aspects of the technology that still require standards, require less technical ability for use due to the development of some user-friendly tools, and have some exposure to public scrutiny. Mature technologies have relatively complete sets of standards, are user friendly due to availability of tested commercial tools or open-source solutions, and have a high degree of public knowledge, understanding, and trust.

Using the three dimensions defined in this section and the information we gathered within the environmental scan, we performed a qualitative assessment of where each PPT we identified sits upon the maturity scale we developed. Figure 2 shows the maturity scale and the location of each PPT upon the scale. Details as to why each PPT was given the location that it was are provided in Appendix B.

**Figure 2. PPT Technological Maturity Spectrum**

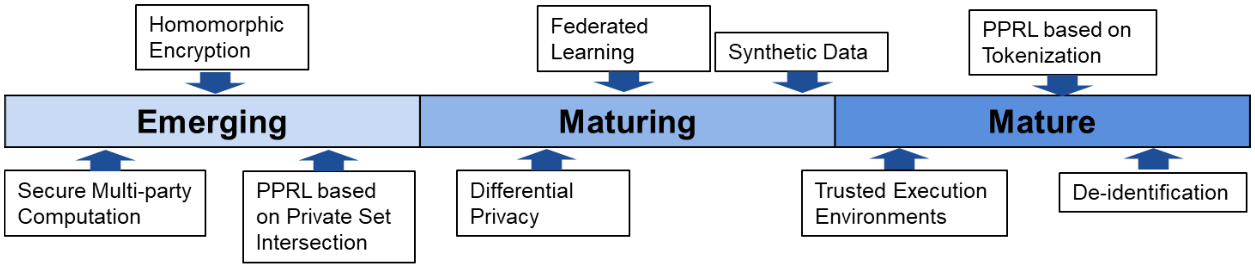
Homomorphic Encryption
Federated Learning
Synthetic Data
PPRL based on Tokenization
Emerging
Maturing
Mature
Secure Multi-party Computation
PPRL based on Private Set Intersection
Differential Privacy
Trusted Execution Environments
De-identification

*Source: RTI based on research*

## PPTs That Protect Input Privacy

We provide details about the PPTs that protect input privacy that we have found within this section.

### *Secure Multiparty Computation (sMPC)*

sMPC covers a broad family of computational techniques involving data from multiple parties that prevent any party from learning about data that are not their own beyond the results of the computation. These techniques use cryptography to enable privacy protected data sharing between the parties, though depending on the architecture of the system and the privacy mechanism, the computations may or may not be performed on encrypted data.

Different sMPC architectures reflect the different needs of the systems in which they are deployed. Use cases where no party can be trusted to hold all the data and run the computation system, even if said data are all encrypted, may require a distributed system. In contrast, certain use cases require that all the data be centralized because one of the parties performing the computation has the authority or the requirement to do so. Distributed sMPC systems rely on computations performed on encrypted data with a shared compute protocol agreed on by all parties, possessing high levels of privacy and security guarantees; however, they have limited flexibility and large computational overhead (see

Figure 3). Centralized sMPC systems perform computations on decrypted data within trusted environments, providing flexibility and reducing overhead, but they may have a lower level of privacy and security (see Figure 4).

**Figure 3. Distributed sMPC Architecture Diagram**

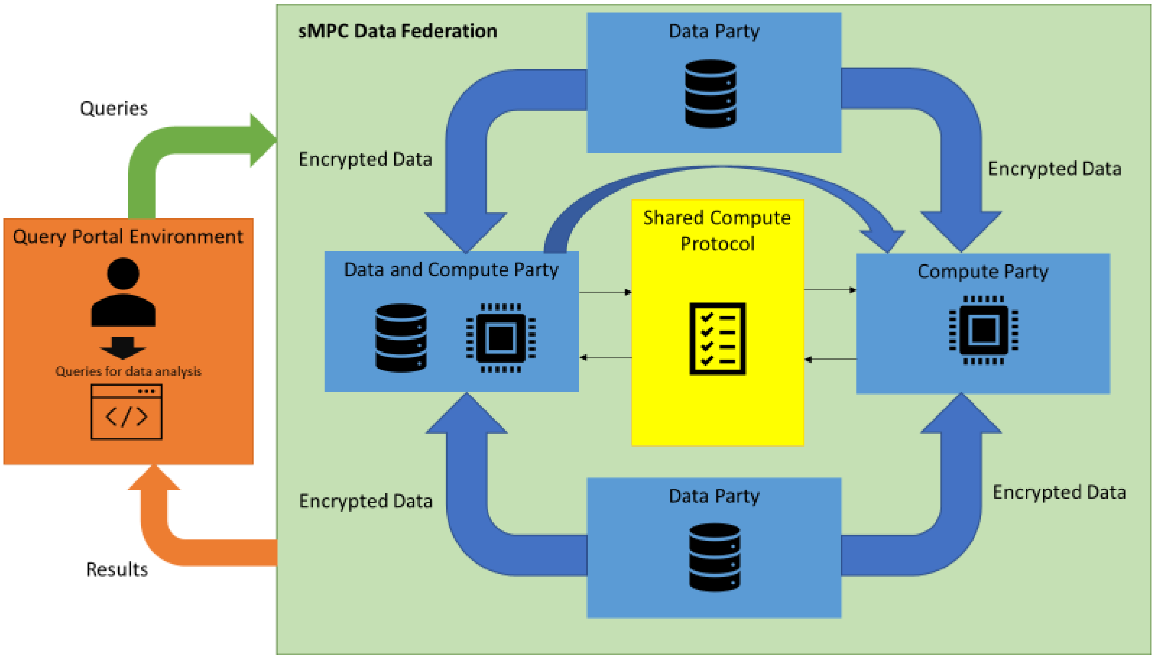

*Source: RTI based on research*

**Figure 4. Centralized sMPC Architecture**

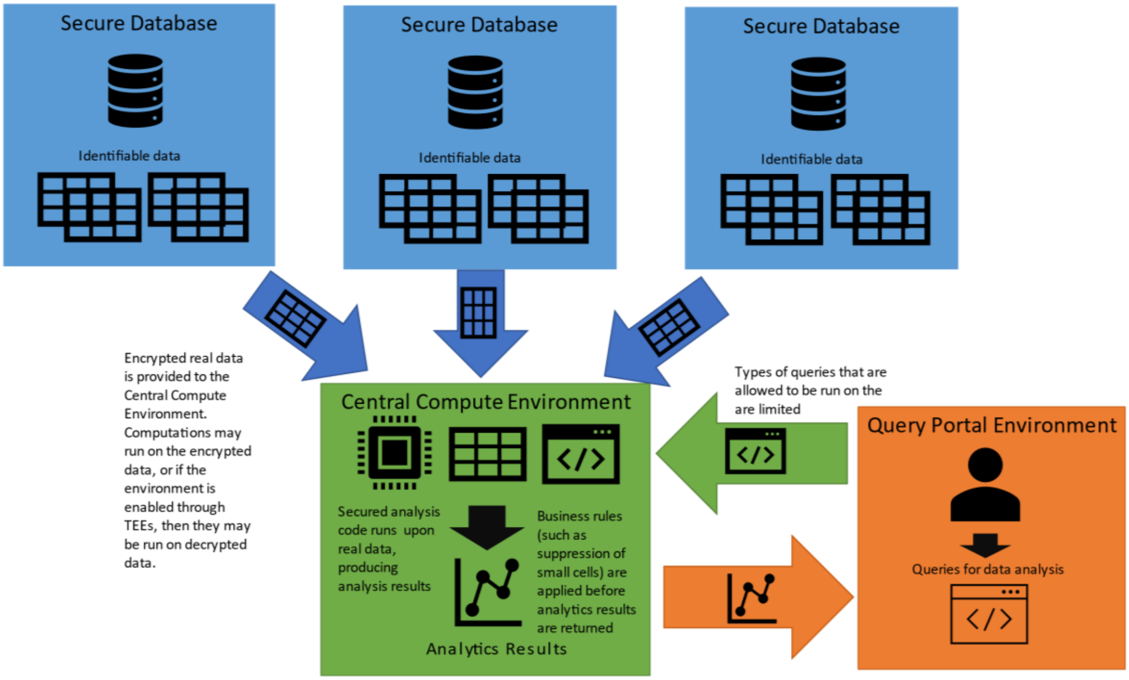
Secure Database
Identifiable data
Secure Database
Identifiable data
Secure Database
Identifiable data
Encrypted real data is provided to the Central Compute Environment. Computations may run on the encrypted data, or if the environment is enabled through TEEs, then they may be run on decrypted data.
Central Compute Environment
Secured analysis code runs upon real data, producing analysis results
Business rules (such as suppression of small cells) are applied before analytics results are returned
Analytics Results
Types of queries that are allowed to be run on the there are limited
Query Portal Environment
Queries for data analysis

*Source: RTI based on research*

A few use cases of sMPC within evidence-building research, policymaking, and program evaluation can be found within the literature. The most notable case is with the Boston Women's Workforce Council's Gender and Racial Wage Gap Studies, which have successfully used sMPC to perform computations over data sourced from multiple organizations across the Boston area for multiple years.[16] There are also a few demonstration projects conducted using human services administrative records,[17] Department of Education data,[18] and health data.[19] Internationally, sMPC protocols have also been explored by the UN through their project sharing trade data across different countries;[5] they have also been explored by the Italian National Institute of Statistics and by Statistics Canada.[20],[21]

As sMPC has a broad range of different implementations, there are also a broad range of different outcomes and limitations possible with different sMPC techniques and architectures. The sMPC implementations that have a distributed architecture and are reliant on cryptography to protect input privacy have the best privacy protections because there is no single point of failure within the system, and data can be encrypted to a high standard of protection. However, performing computations on encrypted data is resource intensive in terms of both computation power and time (see [17], [18] for detailed comparisons in real use cases) and requires expertise from experienced cryptographers to design a secure protocol. The sMPC implementations that do not require computations on encrypted data generally use TEEs to perform computations. These implementations have less security guarantees than encryption methods and may require centralized processing, creating a single point of failure.

### *Privacy Preserving Record Linkage (PPRL)*

PPRL enables multiple parties to compare their data sets without giving up individual privacy, computing the intersection of their data using encrypted identifiers for record linkage. PPRL can be enabled through a variety of methods; some rely on sMPC techniques, linking and performing computations on encrypted data (also known as private set intersection).[34] Other PPRL techniques rely on tokenization, a process by which the PII used in traditional cleartext linkages is transformed into a token that is assigned across multiple sets of data, enabling them to be linked.[35] PPRL often relies on having an honest broker, who holds keys or lookup tables, to facilitate the data linkage (see Figure 5).

**Figure 5. PPRL Process Diagram**

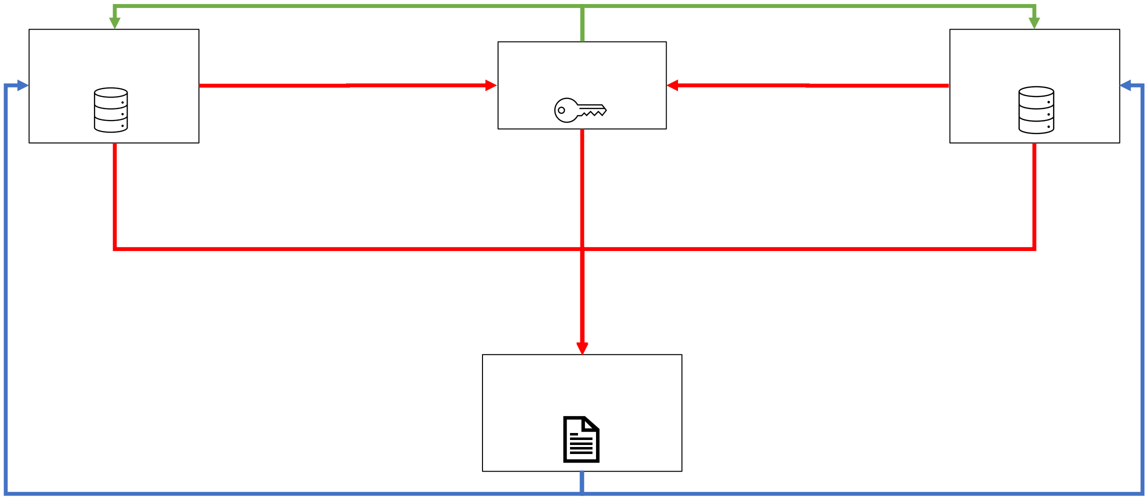

*Source: RTI based on research*

Typical usages of PPRL include de-duplication of records, creating data sets that contain information gathered from multiple sources for research and analyses, and generation of a common individual token within data systems with multiple data sources. There have been many PPRL projects within the federal government from the National Institutes of Health, General Services Administration, and Centers for Disease Control. A selection of these use cases is detailed in Appendix B.

PPRL may enable organizations to share data across programs and jurisdictions where sharing identifiable data is not allowed, such as when sharing sensitive data about vulnerable populations. The ability to share record level data without exposing data subject identities can encourage organizations to re-evaluate the scope and application of policies that prohibit disclosures of identifiable data and open the door for more affirmative data sharing decisions. Administrative and policy needs, such as establishment of data use and data sharing agreements and ethics reviews, should be addressed early in the design process of PPRL systems. Clear explanations of the PPRL security models should be provided to administrators to obtain their understanding and support. Considerations for long-term system sustainability should be taken into account as well, as systems may rely on vendors to provide the tokenization software that enables PPRL, or third parties to function as honest brokers to facilitate linkage.

It is important to note that even though linkage through PPRL should not enable sharing of any information that would directly identify a data subject (such as a name), the resulting linked data set could have a higher risk of re-identification arising from having a richer set of indirectly identifying information. An example of how this can occur can arise from linkage of electronic health record data with employment information, which can add information about a person's occupation and income to information present in health data such as diagnoses, procedures, age, and race. While each individual data set could be considered de-identified, the combined file may raise the risk of re-identification beyond the acceptable risk threshold, requiring the addition of further controls.[35],[36] These controls can include storage of the dataset in environments with more stringent privacy and security controls, or application of output privacy techniques. An evaluation by a qualified expert should be performed to evaluate privacy risks that can arise from linkages and recommend risk mitigation methods.

### *Homomorphic Encryption*

Homomorphic encryption refers to cryptographic techniques that allow for computation over encrypted data. Using such techniques, no party other than the party providing the data can learn anything about the data. Outputs from computations are encrypted as well so that only the party providing the data can decrypt and view them. Homomorphic encryption protocols are often the encryption techniques used to enable sMPC.

Homomorphic encryption techniques come in three different types. Partial homomorphic encryption allows for a single operation to be performed on encrypted data an indefinite number of times. Somewhat homomorphic encryption allows for multiple operations to be performed a limited number of times. Fully homomorphic encryption allows for multiple operations to be performed an unlimited number of times.[22] Fully homomorphic encryption allows for the greatest flexibility in terms of computation; however, it is also the hardest to implement, and effective and usable protocols are still an area of active research.

While there are several software libraries that can enable fully homomorphic encryption,[23]–[26] the technique has numerous limitations. Most notably, the issues of message expansion (where the encrypted ciphertext could be several times larger than the inputted plaintext) and computational overhead (where computations could be slowed tens to thousands of times) present barriers to utility. Implementation of homomorphic encryption also requires cryptographic expertise to ensure security.

### *Trusted Execution Environments (TEEs)*

TEEs originally referred to a feature of modern central processing unit (CPU) hardware that allows for execution of code in a way that mitigates input privacy, code privacy, and code assurance by creating an execution environment that is separate from the rest of the computer system. The definition has been expanded in more recent times to include software enabled TEEs, such as Amazon Web Services (AWS) Nitro Enclaves,[27] which are available from many cloud providers.

The TEE privacy model differs from the model used by cryptographic systems. Privacy guarantees from cryptographic systems arise from computations being performed on encrypted data, which should be meaningless to any party that does not have the decryption key. TEEs perform computations on unencrypted data in an environment that is not visible to any party. This comes with the advantages of better performance and scalability than cryptographic methods and of greater ease of implementation. However, the privacy and security guarantees are weaker as users of TEEs are reliant on trusting the providers of the environment. Usage of TEEs could also result in vendor lock-in because code and processes customized to run on one TEE may not be able to run on a different TEE. Hardware-based TEEs have the added overhead of physical delivery, installation, and storage. TEEs may also have security exploits that compromise the guarantees that they are supposed to provide.

TEEs are being used in many real-world use cases, such as secure key storage in computers, phones, and smart devices; secure enclaves for storage and computations on sensitive data; and sMPC. A demonstration project using human services administrative data used TEEs for sMPC that showed no significant increase in computational overhead while maintaining input privacy for all parties that supplied data, demonstrating that TEE-based sMPC could be viable within shared services environments.[17]

### Federated Learning

Federated learning is a method of training machine learning models by sending copies of the model to each place data reside and performing training on-site, eliminating the necessity of moving large amounts of data to a central location. The central server only receives updates to the model from each location. These updates are then aggregated to make the global model. Split learning is a subset of federated learning where instead of each location sending the full versions of their locally trained models, only a part of the model is sent to update the central model. Split learning provides a stronger privacy guarantee because it can prevent someone from trying reverse local models to obtain input data.[28],[29]

Federated learning is commonly used within smart devices; for example, both Apple and Google have implemented federated learning in prediction models used for keyboard software on smartphones.[30],[31] It is also employed when using sensitive smaller individual data sets that are stored locally on a network of entities or organizations with limited resources to collaboratively train a machine learning solution in a variety of domains, such as health care, finance, and logistics. One notable use case in the latter category arose during the COVID-19 pandemic, where a group of 20 health care institutes used federated learning to collaboratively train a machine learning model to predict COVID-19 outcomes.[32]

Federated learning enables organizations to pool their data for training machine learning models without having to share their data. However, it does not necessarily protect privacy, as private or sensitive information may be leaked through reversal attacks on local models or through the output of the trained model. As such, a common approach to strengthening the privacy and security guarantees is to combine federated learning with other PPTs.[33] Local models may undergo encryption before being sent to the central server, with aggregation being performed on the encrypted models to provide a guarantee against a malicious central coordinator. Outputs from the model may undergo output privacy protection using techniques like differential privacy.

## PPTs That Protect Output Privacy

We provide details about PPTs that protect output privacy that we have found within this section.

### *De-Identification*

De-identification (or anonymization) refers to a suite of techniques for transforming data sets to remove identifying information. These include statistical methods such as K-anonymity (transforms a given set of k records in such a way that in the published version, individuals are indistinguishable from others), and rules-based methods such as the Health Insurance Portability and Accountability Act (HIPAA) Safe Harbor (removal of 18 types of identifying information). De-identification is officially defined under U.S. regulations under section 164.514(a) of the HIPAA Privacy Rule, which lays out the two acceptable methods of Expert Determination and Safe Harbor (see Figure 6).[37]

**Figure 6. De-Identification Methods under HIPAA [37]**

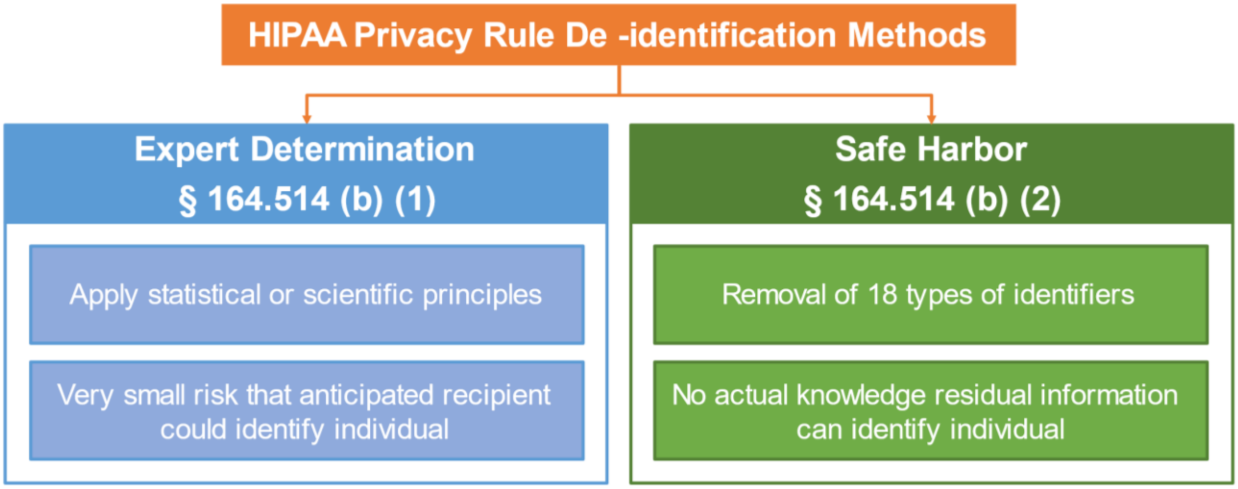
HIPAA Privacy Rule De -identification Methods
Expert Determination
§ 164.514 (b) (1)
Apply statistical or scientific principles
Very small risk that anticipated recipient could identify individual
Safe Harbor
§ 164.514 (b) (2)
Removal of 18 types of identifiers
No actual knowledge residual information can identify individual

*Source: Guidance Regarding Methods for De-identification of Protected Health Information in Accordance with the Health Insurance Portability and Accountability Act (HIPAA) Privacy Rule (https://www.hhs.gov/hipaa/for-professionals/privacy/special-topics/de-identification/index.html)*

De-identification is the most widely used method for protecting output privacy because it is understood, well studied, and relatively simple to implement. The transformations that are applied to de-identify data, such as suppression, generalization, noise addition, and record swapping, have been used to success in various applications such as data releases by the U.S. Census prior to 2020, health data sharing, and data releases from research publications. The techniques used for de-identification can be considered mature, with many vendors selling software that automate the process,[38] and open-source software packages that are well-supported and tested in real situations (such as the sdcmicro package in R, which was used to de-identify and release COVID-19 data by the CDC).[39]

As de-identification is in common use, studies about the vulnerabilities of de-identification have been conducted that show its flaws. De-identification cannot remove all the risk of re-identification from data subjects. It is about managing risk levels to make them below an acceptable threshold.[40] Assumptions must be made about what information a malicious party attempting re-identification could know when performing these evaluations. As more information becomes available through social media and data breaches, re-identification vulnerability increases, to the point where, in many instances, de-identification of data becomes not feasible due to the high impact on data utility. Additionally, de-identification is not compositional, so multiple releases of the same data set de-identified in different manners can result in the original data being reconstructed by colluding parties through the mosaic effect, undoing all privacy protections.[11]

To ensure that privacy protections remain in place with de-identified data, additional legal, regulatory, and contractual controls defining acceptable data use and preventing re-identification, technical controls that limit the types and number of queries that can be conducted, and data storage within secure and monitored environments are recommended. The necessity of having these controls, and the uncertain nature of the ability of de-identification to continue protecting data privacy in the future, has led to exploration of other techniques for protecting output privacy, which will be detailed in the following sections.

### *Noise Addition (Differential Privacy)*

Noise addition is a traditional de-identification technique that involves adding an amount of randomness to data values to render them less likely to re-identify a data subject and to protect sensitive data values such as incomes. The introduction of differential privacy added framework to guide the amount of noise to be added to protect individual privacy. Differential privacy is a formal mathematical definition of privacy based on the idea of the “differencing attack,” preventing the results of data queries from isolating the information about any single individual in the data through the addition of randomness to query results based on a privacy budget (defined by the parameter epsilon).[41] The formalism allows one to create a proof that a noise-addition algorithm will prevent a person from isolating a single person’s data values within a certain number of data queries.

Differentially private algorithms have been successfully used to share and release data by the 2020 U.S. Census, which used a custom differential privacy algorithm.[42] Both Apple and Google use differentially private algorithms to mask user data before they are sent for processing and analysis within their central servers.[30],[43] NIST is gathering a library of differentially private algorithms that they are testing and making open for use within their Privacy Engineering Program’s Prize Challenges and Collaborative Research Cycles.[44],[45]

Differential privacy assumes all information is identifying information, eliminating the challenging (and sometimes impossible) task of accounting for all identifying elements of the data. It is resistant to privacy attacks based on auxiliary information, effectively preventing linking attacks that are possible on de-identified data.[46] Differential privacy is also compositional in that it is possible to calculate the privacy loss from two separate computations performed on a data set by adding up the individual privacy loss from each computation.[41]

Despite its promising capabilities, differential privacy has certain limitations that have prevented it from being put into broader use. There is no consensus on how to choose the value of epsilon, which determines the privacy budget, nor is there agreement on how to approach this and other key implementation decisions.[47] Also, because differential privacy adds noise to data, it has the side effect of removing outliers and introducing inaccuracy within small, diverse populations, which makes it inappropriate for certain studies.[48] Differentially private algorithms require technical expertise to both design and implement, as they require customization to both the data set and the use case. Of note is that while certain algorithms may be differentially private, the resulting data after applying the algorithm may not be able to satisfy other privacy metrics, meaning that they may not have as strong privacy protections as supposed.

### *Synthetic Data*

Synthetic data involve creating a data set using statistical or machine learning techniques that contains brand new records with similar aggregate statistical properties as the original data set. As a privacy protection technique, the argument for synthetic data is that the records within the data set are not the records in the real data, and thus the synthetic data can be shared and released without compromising privacy. There are two families of approaches to generating synthetic data; the traditional method uses statistical methods such as generation of values from multidimensional models and imputation; more recently, methods using generative adversarial neural networks have become popular.[49]

Creating synthetic data for data sharing has been in use for a long time. The U.S. Census Bureau has actively worked on generating synthetic data products such as the Survey of

Income and Program Participation (SIPP) Synthetic Beta for decades.[50] The National Center for Health Statistics (NCHS) has created and released partially synthetic linked mortality files for public use.[112] Other national statistical organizations (e.g., Statistics Canada,[51] National Statistical Office of the United Kingdom[52]) have worked on creating privacy preserving synthetic data sets. Outside of national statistical offices, synthetic data have been used in place of sharing sensitive administrative data such as tax information.[53]

Synthetic data have several limitations that must be kept in mind. Synthetic data may not be suitable for general analysis because not all relationships that are present within the real data can be preserved in synthetic data. Synthetic data are not useful when one wants to ask questions in the future which are beyond the scope of the requirements when creating the data sets because the synthetic data algorithm cannot guarantee that the characteristics required to answer those questions will be preserved. Synthetic data also may not be able to capture outliers or accurately represent small populations with significant variations within the original data. Due to these limitations, synthetic data are better used for testing purposes, such as helping researchers develop queries that they can then compute on real data, saving time and resources and enhancing security within secure computing environments. Usage of synthetic data in this fashion could support the idea of a tiered access model as mentioned in the Evidence Act, where the synthetic data can reside in an access tier with lower controls to be used for testing, before providing users access to real data in a higher access tier.

There are also some privacy concerns with generation and usage of synthetic data. Artificial intelligence models used by synthetic data generators may remember some personal information, especially when the original data are sparse, which is likely in high-dimensional data sets such as images, text, or series of events and the model has a large learning capacity. Very flexible models can overfit the data, leading to generation of records that contain potentially sensitive information. Measurement of the levels of privacy protection offered by synthetic data generation is an active area of research, with organizations like NIST working on gathering ideas and coming up with standards for privacy metrics.[54] The Federal Chief Data Officers Council has also put out a request for information regarding synthetic data generation.[113]

## Information Gathered from Interviews and Discussion Panels

The focus of the interviews[1] we conducted with PPT practitioners was to obtain personal experiences working with PPTs, such as what worked well within their implementation process, barriers and challenges they faced, and any lessons they learned that could help inform future development. Practitioners from across the government, academic, and industry sectors expressed common themes regarding their experiences, with nuanced

[1] OMB control number 3145-0215.

differences between practitioners within and outside the federal government. We detail these common themes and differences in this section.

### Common Themes Expressed by PPT Practitioners

#### *Theme 1: PPTs are employed to try and meet a certain set of needs*

Practitioners expressed that their work with PPTs began to meet a certain set of privacy needs that arose from the increased use of data within evolving technologies. Traditional privacy protection methods such as legal and contractual controls and methods of data transformation are becoming insufficient in the face of big data and advanced analytics (i.e., AI). They expressed a desire for privacy protection throughout computation and analysis to address the privacy risks presented by the increasing availability of data and increasing capabilities of analytics, with the goal of preventing third parties from accessing any identifiable or private information that is not their own.

Practitioners also expressed a need to provide measurable outcomes of privacy protection to meet evolving legislative and regulatory requirements. For instance, under the HIPAA Privacy Rule, de-identification of data under expert determination requires a determination that "the risk is very small that the information could be used, alone or in combination with other reasonably available information, by an anticipated recipient to identify an individual who is a subject of the information."[37] Under European Union's General Data Protection Regulation (GDPR), measurement of a re-identification likelihood is a requirement for justification that a data set can be considered pseudonymized or anonymized.[55] Many traditional methods for de-identification, such as record-swapping, rounding, and suppression, do not have measurable outcomes. New methods enabled by PPTs, such as differential privacy, try to produce these privacy measures to meet these requirements.

PPTs are also employed by organizations to enhance trust in their data handling and analysis processes. This use of PPTs not only enables these organizations to gather more data for their analyses but also enables democratization of data by allowing for privacy protected methods of sharing data with decision-makers and the public.

#### *Theme 2: Successful use of PPTs requires input from key stakeholder groups*

Practitioners expressed that deployment of PPTs oftentimes requires changes to systems and processes from traditional methods. Understanding the requirements of these systems and processes, and communication of potential changes to stakeholders from the start of these projects, is key for success. They identified the following stakeholder groups which are important for engagement: organizational decision-makers, business leaders, legal and regulatory teams, IT support, and analytics (the end users of the data).

Obtaining a clear understanding of the needs of business and analytics and the legal obligations from the stakeholders is needed to define system requirements. Key stakeholders should be involved as early as possible in the development lifecycle to ensure buy-in, build trust, and ensure that risks are addressed at the outset. Communications and understanding of the needs of each group by the other groups are necessary for a complete definition of the use case and the data lifecycle. These conversations help to establish the data governance setup for the data system. Practitioners expressed that these conversations define the requirements for the PPTs to use within the system, allowing them to find a PPT or combination of PPTs that best meets the requirements of every group and then using the PPT or combination of PPTs where they would best fit. Definition of the system requirements with clear input from all stakeholder groups early on can help expedite the successful development and deployment of PPTs within data systems.

*Theme 3: Bridging the technical knowledge gap between PPT practitioners and nonpractitioners*

For effective communications about PPTs between the different stakeholder groups to take place, PPT practitioners must help bridge the technical knowledge gap. Stakeholders without PPT expertise need to understand the essence of these technologies so that they can be used in the right places. Practitioners expressed a desire for communications materials that can explain what these technologies do in a simple and easy-to-understand manner. They also expressed a need for people who can translate the technical aspects of these technologies for those without technical background. Some initiatives are underway to create these materials and train these personnel, such as NIST’s Differential Privacy Blog Series and NIST’s work with academics to produce PPT curricula, and various companies also are working on materials that explain the technical aspects of the PPTs they are developing.

Bridging the knowledge gap helps establish trust in the capabilities of these technologies, as evidenced through PPRL implementations such as the Biomedical Research Informatics Computing System (BRICS)[56] and the National Covid Cohort Collaborative (N3C),[57] which presented the technical aspects in a way that was understandable by stakeholders, winning their trust and obtaining their support, allowing those systems to be developed and used. PPT practitioners in industry also have experience in creating such explanations for their tools to bridge these gaps, which has helped with adoption of their tools in various systems.

*Theme 4: Concerns about tradeoffs between privacy and utility*

Practitioners expressed concerns that PPTs use could affect the ability of the end user to use the data, both in the methods by which analysts perform their data analyses and in the accuracy of the results. For example, when using encryption-based techniques such as homomorphic encryption, computations on encrypted data are different from computations on cleartext data, requiring new sets of analytic tools. Many times, after applying these techniques, the analysis that can be performed on the data will be limited to specific predefined use cases. This can prevent exploratory analysis, data cleaning, and many forms of data modeling, reducing the utility of the data. During interviews, practitioners expressed that output privacy techniques, such as differential privacy and de-identification, introduce noise into the data, which can have negative impacts on data about small, varied, and vulnerable populations (i.e., tribal nations, substance abuse victims), ranging from lack of accuracy in the results to a complete inability to conduct analyses as the data have been suppressed.

PPT practitioners working with analysts (the end users of the data) can resolve these issues by defining the data utility requirements from the start and then tuning the PPTs employed within a data system to output data that can meet those requirements. The PPT practitioners can also use the collaboration with analytics to educate them about usage of data that have been protected using PPTs, which can lessen impacts on analytics workflows.

#### *Theme 5: Importance of standards for PPTs*

Many practitioners we interviewed expressed a desire for establishing standards for the PPTs in which they have expertise. Establishment of standards for these technologies is good for multiple reasons: standards help with establishing trust in the capabilities of PPTs; they can help to establish clarity on which techniques work best in what situations; they can be a measure of technological maturity; and they can be used for auditing, creating an additional trust mechanism. Standards guide implementation of technologies in a safe manner that has been tested and accepted by experts, helping with future development. The practitioners we interviewed expressed that there is a lot of confusion about the capabilities of different technologies. They believe that setting standards can help provide clarity and separate the "signal" from the "noise." With sufficient standards setting, trust marks (such as the lock icon on secure websites) can be established. These trust marks can help organizations demonstrate their compliance to the state of the art in privacy protection, enable those without technical background in PPTs to distinguish which systems have the right privacy protections built in, and give the public an easy method to determine when their privacy is being protected. The standards-setting process also has the additional benefit of bringing experts together to create communities of practice.

Various industry groups have come together to start forming standards, such as the group of experts from industry, academia, and government who are currently working on standardization of homomorphic encryption at https://homomorphicencryption.org/. NIST as well is doing significant work on PPT standards, establishing communities of experts, and running contests to gather information and determine what technologies work best in which situations. Also, at the International Organization for Standardization (ISO), there are standards that are currently being set on various PPTs.[58]–[64]

### *Theme 6: Development of PPT expertise*

Implementation of PPTs often requires specialized technical talent. Many current PPTs are nascent technologies that often require bespoke tailoring to meet individual project requirements. For example, differential privacy algorithms such as the one used in the 2020 Census release are custom designed to meet data release requirements. Cryptographic techniques such as homomorphic encryption, sMPC, and PPRL work best when expert cryptographers are part of the team to ensure that implementations do not leave exploits. Practitioners within the federal government expressed a need to increase technical ability in cryptography and data privacy to meet implementation needs for these technologies.

A solution to address the need for technical expertise is to use technologies developed and supported by industry partners (i.e., purchasing solutions from a technology vendor). Pros to using vendors include vendors' support for implementation, usage, and issues resolution of their technologies; provisioning of regular software and security updates; and access to a tested solution supported by an experienced team. However, there are also several cons to using vendors, including the risk of vendor locking, where systems become dependent on a particular vendor's solution, which leads to risks of support discontinuation. Using vendors also does not solve the issue of a lack of development of internal capabilities with PPTs. Also, vendor solutions may have a lack of transparency due to black box designs. Practitioners expressed that a more complete solution would be to use vendor technologies as a bridge toward developing internal systems controlled by the government. To do so would still require development of in-house expertise, which would need investments in training and talent growth.

In-house development of PPTs requires developing additional talent in cryptography, computer systems, data management, and data privacy. Talent within these fields is in high demand, as both industry and academia are competing with government for individuals with these specialized skills. PPT training materials are still in the early phases of development; few of the practitioners we interviewed had experience giving formal trainings on PPTs, with most information sharing even among practitioners occurring through informal conversations or presentations at conferences. Some initiatives are underway to create training materials, such as efforts by NIST to create curricula that can be used at the university level and academics beginning to teach courses on topics like differential privacy.

Materials for informing and educating nontechnical stakeholders on PPTs also need further development. Federal PPT practitioners whom we interviewed expressed a desire for officially approved communications materials they can use when speaking to stakeholders on projects, such as guidance materials provided by the Office of Information and Regulatory Affairs. These materials would help ease the difficulty in communicating the capabilities of these tools and increase trust. Nonfederal PPT practitioners have more experience creating and communicating these kinds of materials, though they are usually limited to the specific types of PPTs in which they have expertise. PPT practitioners from across all sectors expressed a desire to help create these kinds of resources.

Several practitioners expressed interest in establishing communities of practice both around their specific PPTs and across PPTs in general. These communities could help in both the creation of communications and training resources and in training new talent within the PPT space, as evidenced by past efforts by industry and by NIST. Practitioners expressed that the NSDS could include one or more such communities of practice, acting as a central resource of expertise that can be drawn on to help guide the implementation and usage of PPTs across the government.

### *Theme 7: Establishing a culture of innovation*

Beyond development of talent, establishing a culture of innovation within the government is a key step toward effective usage of PPTs within government systems. The government tends to choose tested, mature, trusted technologies for deployment within government systems. Public scrutiny and political concerns create aversion to risk because failure could mean great consequences. As such, systems used within the government should have a high level of trust and a high level of service stability, which emerging technologies like many PPTs may lack. Many PPTs require capabilities testing and development of systems that can support their use. These activities inherently have a higher level of risk because they require experimentation and innovation—and may not always lead to success.

Establishment of a culture of innovation would need support from various groups. Leadership would need to support trying new systems and technologies over the old ways of doing things. Legal and regulatory bodies would need to support examinations of whether new systems can meet legal and policy requirements. IT would need to support development of infrastructure to enable testing and usage of these technologies. Finally, the end users of the data would need to support usage of these technologies and of performing analytics on their outputs.

Building a community of innovators with knowledge and talent in PPTs that the government can draw on when developing and testing PPTs can be a good first step in creating the culture of innovation. These groups can partner with innovators in industry and academia to test new technologies in government systems, a good extant example of which is the Emerging Technology Fellowship within Census XD.[65] This proposed community of innovators can become a key part of a shared services environment, acting as one of the shared services to support the development and implementation of PPTs across multiple agencies. Another good step could be establishing an environment for testing ideas (i.e., a sandbox), which could encourage innovation in the development and usage of PPTs. The innovations themselves could start small, beginning with low-hanging fruits involving mature technologies using proven techniques (such as tiered access to data), then move on to experimental techniques after learning lessons from initial successes.

*Keys to success*

After speaking to interviewees and participants in facilitated group discussions about the PPT projects they have worked on, we synthesized the following lessons regarding the key components to success in implementation and usage of PPTs.

1. **Starting small and bringing together a multidisciplinary team.** While it may be tempting to try and use PPTs throughout a technology system, from the experiences of practitioners, the best approach is to start small, creating a clear definition of the question that is being asked, then tackling a solvable issue first to show the viability of the technology before scaling up to more complex problems. Demonstrating success using such an approach can help build trust in skeptical stakeholders. Inclusion of all stakeholders in the implementation team, such as specialists in statistical methodologies, experienced security and privacy engineers (cryptographers), legal experts, IT staff, and analytics (the end users of the data), would help to ensure that the use cases and problems that usage of PPTs is aiming to solve are well-defined, that the perspectives and needs of all parties are understood, and that the solution is able to meet the requirements and needs of all groups.
2. **Building trust with stakeholders.** Stakeholders' understanding and having trust in the privacy protection mechanisms used by PPT can help overcome hurdles during the design and implementation process. A key to building understanding and trust is effective communication of privacy and security mechanisms in a simple manner. Technologies that can meet tested standards set by official standards-setting bodies can also build trust, especially in cases where the technical aspects are more difficult to explain.
3. **A culture of innovation will help to build the talent necessary to test and implement these technologies.** Interviewees highlighted a desire for innovation not only in terms of technology usage but also in legal, regulatory, and policy frameworks to accommodate the usage of PPTs and in innovations in governance and management that could be enabled by using PPTs. These innovations can be sped up through partnerships with academic groups and private companies, allowing them to develop and deploy their solutions within shared services environments. They can also help the government develop the culture of innovation necessary to speed adoption of PPTs, such as by helping to create communities of practice from which the government can draw talent.

### *Differences in Perspective between Federal and Nonfederal PPT Practitioners*

PPT practitioners both within and outside the federal government expressed many of the same themes when discussing their experiences with implementation of PPTs, with differences lying in the nuances of their experiences. The most significant difference was in the stakeholder groups that practitioners emphasized as important to engage throughout

PPT projects; federal practitioners focused on legal, regulatory, and policy stakeholders, whereas nonfederal practitioners also included business leadership, IT, and analytics. This difference arose from practitioner motivations when using PPTs to inform data related research and policy questions; federal practitioners' primary aim is to meet legal and regulatory requirements while sharing or releasing data, whereas nonfederal practitioners place using the data to satisfy specific business and analytics needs first. Another difference between federal and nonfederal PPT practitioners is experience in the types of PPTs with which they have worked. Federal practitioners whom we interviewed tend to have more experience with more mature PPTs (such as PPRL) and less experience with less mature PPTs (such as sMPC). Nonfederal practitioners tend to have used a greater variety of PPTs at varying levels of technological maturity in their work.

In terms of PPT implementation, the nonfederal practitioners we interviewed have more experience implementing a variety of PPTs in both experimental and production-level projects. With that experience, they were able to share more about getting over some of the hurdles presented in the earlier sections. From their experience, the key is to have communications between all different groups to establish requirements of systems from the start. While both federal and nonfederal practitioners expressed a need for translation of the technical aspects of PPTs into forms understandable by nontechnical legal, regulatory, and policy stakeholders, nonfederal practitioners also expressed a need to translate legal, regulatory, and policy requirements back to business, IT, and analytics stakeholder groups, so that requirements can be clearly defined with alignment from all stakeholder groups. After clear definition of those requirements, implementation of PPTs becomes much easier.

### *Navigating Legal, Regulatory and Policy Challenges*

In this section, we present legal, regulatory and policy challenges identified by practitioners whom we interviewed and share their perspectives on key approaches for navigating these challenges.

#### *Perspective 1: Clearly articulate the value and benefits of PPTs within data privacy laws and regulations*

Legislation and guidance on the use of PPTs may be addressed explicitly or implicitly in a country’s data privacy laws and regulations. The European Union’s GDPR, Article 25, has several articles that refer to privacy enhancing technology (PET) use. While the United Kingdom has no specific legislation governing PET use, its Data Protection Act (2018), which implements GDPR, includes a requirement for data protection by design and default.[13] In the U.S., the use of personal data is governed by a patchwork of federal, state, and tribal laws and regulations. The breadth of U.S. laws governing the use of personal data creates a complex regulatory landscape that data subjects and data controllers must navigate. These laws and regulations provide a general framework for data collection and use but do not explicitly require or encourage PPTs. Nor does U.S. law speak to the ability of PPTs to meet specific regulatory requirements, creating lack of incentive and gray areas within the law for legal teams considering PPT adoption.[13]

However, the U.S. regulatory landscape for PPTs is evolving. President Biden's 2023 Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence calls for strengthening of privacy preserving research and technologies.[3] The California Consumer Privacy Act (CCPA) permits re-identification of data to validate new privacy preserving techniques.[12] Enhancements to our current regulatory framework which encourage or incentivize the adoption of privacy by design principles, regulatory sandboxes to test PPTs for priority use cases, shifts in perception of privacy along the spectrum of PPT solutions, and streamlined, machine readable data use agreements all hold great promise for clarifying the role and benefits of PPTs.

*Perspective 2: Bridge information gaps between technology and legal teams*

The practitioners we interviewed highlighted the difficulties of getting to yes on PPT deployment and the need for better communication between technical and legal teams. Lawyers may view privacy in a binary fashion, evaluating the extent to which a PPT achieves and maintains a particular privacy threshold. Technologists approach privacy in a more linear fashion, protecting privacy based on a spectrum of risk and making necessary privacy-utility trade-offs. Boundaries need to be clearly communicated to ensure alignment of both viewpoints, as stakeholders try to evaluate the viability of PPTs for a particular use case.

For many nontechnical stakeholder groups, PPTs are new and not fully understood. Technologists will need to clearly communicate to legal teams how and the extent to which the technologies meet business and legal requirements. As expressed by one of the practitioners we interviewed, these conversations are "nascent, bespoke, and very difficult." Collaboration and "translation" between these two stakeholder groups can help to clarify misconceptions about PPT compliance capabilities and provide needed clarity for both sides about what privacy technology can and cannot do for implementers and data owners.

*Perspective 3: Broaden our view of privacy and the role of the privacy professional*

Privacy is often viewed as a compliance and documentation exercise, focused on systems of records notices and privacy impact assessments. These activities, while critical and required under laws such as the Privacy Act of 1974, typically happen after a system has been designed, limiting the opportunity for critical privacy inputs during the system design stage. Practitioners expressed a need to shift thinking about privacy beyond compliance toward a risk management exercise, in which privacy professionals (including lawyers, regulators, and policy experts alongside PPT practitioners) are seen as design partners and participate in all stages of the system development lifecycle.

*Perspective 4: View privacy as highly contextual rather than static and binary*

Under the current regulatory framework, data controllers are responsible for maintaining confidentiality of PII. Data controllers are generally prohibited from sharing PII, with exceptions only for specific cases such as supporting law enforcement. Noncompliance exposes responsible data controllers to increased risk of litigation, financial loss, and reputational damage.[66]. The current regulatory framework and binary perceptions of privacy work well for use cases involving PII data sharing; but in some use cases that involve PPTs, according to practitioners we interviewed, this perception may not account for nuances in the types of information sharing that might occur, examples of which include the following:

- A data controller uses PPRL powered by homomorphic encryption to input patient PII to an interface, but the data are never seen by the data sharing partner. Rather, the data sharing partner receives linked, de-identified records.
- Aggregated insights are gathered from parties that never shared their data.
- A data controller shares attestations about the data, but not the data itself (e.g., a state vital records system returns a yes-or-no validation for a death record query by authorized parties).
- A data controller incorporates privacy budget or limits for statistical risk of re-identification as a risk mitigation or compliance measure.

Practitioners we interviewed acknowledged the ability of PPTs to conduct various types of data sharing, customized to organizational or client needs, but raised the question of how we define privacy given the spectrum of data sharing that is possible without full disclosures of PII through PPTs. This question suggests the need to shift from a binary view of privacy that assumes shared data are either PII or not PII to an approach that assumes varying tiers and types of data sharing with different levels of protection and controls applied to manage risk along a spectrum.

*Perception 5: Streamline data use and data sharing agreements*

Many practitioners we interviewed expressed the need to streamline the process for executing data use agreements (DUAs). DUAs are legal agreements that memorialize terms and conditions for data use, roles and responsibilities of data providers and data users, and enforcement rights of the data provider. They formalize stakeholder consensus on how data will and will not be used, documenting the scope of shared data assets, data access roles, and processing allowed on the data. The downside to DUAs is that they need legal expertise, can be difficult to negotiate, and often take a great amount of time to execute, particularly when there are multiple parties to the agreement. Organizations without sufficient staff and resources to negotiate, draft, and review DUAs may be unable to share or receive mission-critical data.

During interviews and discussion panels, practitioners highlighted the need to shift from a human-driven approach to a more streamlined, automated approach to reduce time and resource burden associated with DUAs. Practitioners also indicated that efficiencies might be possible in use cases where data are shared but not “seen.”

# Proposed Next Steps

## Setting Up a Sandbox Environment for Testing PPTs

Many PPTs are nascent technologies whose capabilities are not well understood, needing testing to ensure that they perform as advertised. These tests could include testing to ensure the privacy and security guarantees are realizable in real use cases; testing to determine the level of customization needed for specific use cases; and testing to determine the best- and worst-case situations for risk management. To test these PPTs, a safe and secure environment would be necessary. This environment should be separated from production systems and have access to synthetic data that mimic real data. Additionally, the tests should be conducted in a manner where the impacts of failure are limited because failures are to be expected in this exploratory phase. Practitioners expressed that a good model for this kind of sandbox environment can be found with the UN PET Lab,[5] which allowed for testing of various PPTs within real-world use cases.

## Establishment of a Community of Practice to Foster PPT Expertise

Many practitioners we interviewed expressed an interest in and a desire for a pool of expertise on PPTs that can be drawn on to help support development and implementation of PPTs within government projects. Some practitioners expressed that such a community of practice could become one of the shared services that are offered by the NSDS. The goals of this community could go beyond just being a shared resource because it can be used to foster new talent through mentorship. This community can involve experts from outside the federal government as well, through hosting events such as group discussions, hackathons, and competitions or by extending invitations to experts to speak about PPTs.

## Exploration of Data Governance Solutions

As mentioned by practitioners, usage of PPTs can come with legal, regulatory, and policy challenges arising from data governance models established before the advent of these technologies. The impact of usage of PPTs on traditional governance models is worth further exploration and research. PPTs can present opportunities for streamlining and automating processes used for enabling access and use of data. Current governance models rely on legal agreements that could take months to process; with the usage of PPTs, there could be technological enforcement of data access and data usage rules. Further exploration of how this can be approached when using PPTs is warranted because there are benefits such as reduced friction and increased timeliness of data delivery.

Note that usage of PPTs in this manner could result in changes to how data systems operate, both from a technological standpoint and from a governance standpoint. How that usage would affect current data governance models is not well understood. The practitioners we interviewed believe that setting down well-defined governance structures prior to using

PPTs helps in defining their requirements and speeds up their implementation and adoption. As such, further exploration of how data governance could change when PPTs are used in certain use cases is a prudent step to take before attempting implementation within data systems.

### Creation of Communications Material to Help Inform Multiple Stakeholder Groups

Many practitioners expressed difficulty in communicating information about PPTs to those without technical background in PPTs, particularly due to a lack of effective materials to aid in the communications. Development of some of these materials, which can then be used more broadly across government departments, can be one of the shared services the NSDS provides. These materials should strive to communicate the capabilities of different PPTs in a nontechnical and approachable manner, enabling understanding by the key stakeholder groups that practitioners identified (e.g., legal, business, IT, and analytics). The development of these materials can be collaborative between in-house PPT experts as well as nonfederal practitioners, using expertise from practitioners within the industry who have successful experience crafting communications about PPT products, and expertise from federal government PPT practitioners who have implemented PPTs within federal technology systems.

# Conclusion

To improve understanding of the current PPT landscape, RTI partnered with FMH to conduct an environmental scan of PPTs currently being developed, tested, and utilized across government, academia, and the private sector through the performance of a literature review and through engaging practitioners of PPTs through interviews and technology demonstrations. The purpose of this study is to gather information to inform future exploration and testing of PPTs as potential shared services in support of an NSDS-D project as authorized under Section 10375 of the CHIPS and Science Act of 2022.

While conducting the environmental scan, we found that the term PPT has been applied to a broad range of technologies. To help us find the types of technologies that can be beneficial for shared services environments, we established a set of inclusion and exclusion criteria (detailed in Appendix A) and developed a taxonomy to classify the technologies we found. Using this taxonomy, we found that the technologies that proved to be the best fit are "hard" data privacy technologies that can protect either input or output privacy. We focus our scan on these technologies to answer the five questions that NCSES posed in the solicitation.

To conduct our scan, we first performed a literature review to identify PPTs frequently referenced in the literature. Next we conducted interviews and facilitated group discussions with practitioners of the technologies identified in the literature so that we can have a holistic view of each technology's benefits and limitations. Through these interviews, we found many common themes expressed by multiple practitioners regarding successful use of PPTs, including bridging knowledge gaps, inclusion of key stakeholder groups, setting standards, development of expertise, and establishing a culture of innovation. They also expressed several common concerns, such as the tradeoff between privacy and utility, the need for communications materials to help educate those without technical background in PPTs, and how to navigate legal, regulatory, and policy challenges. These common themes and concerns became the topics that were discussed within the three facilitated group discussions we held with practitioners, which both gathered valuable insights about these themes and sparked new ideas catalyzed from the group environment.

From the information we gathered, we assessed the technological maturity of the different PPTs based on the factors of standards setting, ease of use, and public trust, which we identified as important for successful implementation through the literature review and interviews with practitioners. We also proposed next steps for furthering the development and usage of PPTs, which include creating a sandbox environment for testing PPTs, setting up a community of practice, exploring the impacts of PPT usage toward data governance, and crafting communications materials for informing those without technical background in PPTs.

# References

[1] “Advancing a Vision for Privacy-Enhancing Technologies | OSTP,” The White House. Accessed: Jan. 04, 2024. [Online]. Available: https://www.whitehouse.gov/ostp/news-updates/2022/06/28/advancing-a-vision-for-privacy-enhancing-technologies/

[2] Fast-Track Action Committee on Advancing Privacy-Preserving Data Sharing and Analytics, “National Strategy to Advance Privacy-Preserving Data Sharing and Analytics,” Mar. 2023.

[3] The White House, “FACT SHEET: President Biden Issues Executive Order on Safe, Secure, and Trustworthy Artificial Intelligence,” The White House. Accessed: Jan. 04, 2024. [Online]. Available: https://www.whitehouse.gov/briefing-room/statements-releases/2023/10/30/fact-sheet-president-biden-issues-executive-order-on-safe-secure-and-trustworthy-artificial-intelligence/

[4] “Advisory Committee on Data for Evidence Building: Year 2 Report”. Advisory Committee on Data for Evidence Building (ACDEB), Oct. 2022. Accessed: Jan. 30, 2024. [Online]. Available: https://www.bea.gov/system/files/2022-10/acdeb-year-2-report.pdf

[5] “UN PET Lab,” United Nations, Press release, Jan. 2022. Accessed: Jan. 30, 2024. [Online]. Available: https://unstats.un.org/bigdata/events/2022/unsc-un-pet-lab/UN%20PET%20Lab%20-%20Press%20Release%20-%2025%20Jan%202022.pdf

[6] AppsFlyer, “Privacy preserving technologies: AppsFlyer mobile glossary,” AppsFlyer. Accessed: Nov. 20, 2023. [Online]. Available: https://www.appsflyer.com/glossary/privacy-preserving-technologies/

[7] S. Blanchard, “Privacy enhancing technologies and how they can help,” Data Protection Network. Accessed: Nov. 20, 2023. [Online]. Available: https://dpnetwork.org.uk/privacy-enhancing-technologies/

[8] M. Cobb, “Privacy-enhancing technology types and use cases,” TechTarget. Accessed: Nov. 20, 2023. [Online]. Available: https://www.techtarget.com/searchsecurity/tip/Privacy-enhancing-technology-types-and-use-cases

[9] F. Ricciato, A. Bujnowska, A. Wirthmann, M. Hahn, and E. Barredo-Capelot, “A reflection on privacy and data confidentiality in Official Statistics,” Aug. 2019.

[10] J. Stutz, “Structured Transparency: Ensuring Input and Output Privacy,” OpenMined Blog: Privacy AI Series. Accessed: Nov. 20, 2023. [Online]. Available: https://blog.openmined.org/structured-transparency-input-output-privacy/

[11] “The United Nations Guide on Privacy-Enhancing Technologies for Official Statistics.” United Nations, 2023. Accessed: Aug. 21, 2023. [Online]. Available: https://unstats.un.org/bigdata/task-teams/privacy/guide/2023_UN%20PET%20Guide.pdf

[12] K. Asrow and S. Samonas, “Privacy Enhancing Technologies: Categories, Use Cases, and Considerations.” Federal Reserve Bank of San Francisco, Jun. 01, 2021. Accessed: Aug. 24, 2023. [Online]. Available: https://www.frbsf.org/economic-research/wp-content/uploads/sites/4/Privacy-Enhancing-Technologies-Categories-Use-Cases-and-Considerations.pdf

[13] "Emerging privacy-enhancing technologies: Current regulatory and policy approaches," OECD, Paris, Mar. 2023. doi: 10.1787/bf121be4-en.

[14] G. M. Garrido, J. Sedlmeir, Ö. Uludağ, I. S. Alaoui, A. Luckow, and F. Matthes, "Revealing the Landscape of Privacy-Enhancing Technologies in the Context of Data Markets for the IoT: A Systematic Literature Review." arXiv, Jul. 12, 2022. Accessed: Aug. 22, 2023. [Online]. Available: http://arxiv.org/abs/2107.11905

[15] D. W. Archer et al., "UN Handbook on Privacy-Preserving Computation Techniques." arXiv, Jan. 15, 2023. doi: 10.48550/arXiv.2301.06167.

[16] "Data Privacy," Boston Women's Workforce Council. Accessed: Nov. 20, 2023. [Online]. Available: https://thebwwc.org/mpc

[17] N. Hart, D. Archer, and E. Dalton, "Privacy-Preserved Data Sharing for Evidence-Based Policy Decisions: A Demonstration Project Using Human Services Administrative Records for Evidence-Building Activities." Rochester, NY, Mar. 28, 2019. doi: 10.2139/ssrn.3808054.

[18] D. Archer, A. O'Hara, R. Issa, and S. Straus, "Sharing Sensitive Department of Education Data Across Organizational Boundaries Using Secure Multiparty Computation," Georgetown University, May 2021. Accessed: Jan. 04, 2024. [Online]. Available: https://drive.google.com/file/d/1CURfl3q8j_NOBiaOuPEleJBZFpwQcwti/view

[19] J. Rogers *et al.*, "VaultDB: A Real-World Pilot of Secure Multi-Party Computation within a Clinical Research Network." arXiv, Jul. 25, 2022. doi: 10.48550/arXiv.2203.00146.

[20] D. Buckley, "Italian National Institute of Statistics and Bank of Italy: Enriching data analysis using privacy-preserving record linkage - UN GWG on Big Data - Privacy Preserving Techniques Wiki - UN Statistics Wiki," UN Statistics Wiki. Accessed: Nov. 20, 2023. [Online]. Available: https://unstats.un.org/wiki/display/UGTTOPPT/5.+Italian+National+Institute+of+Statistics+and+Bank+of+Italy%3A+Enriching+data+analysis+using+privacy-preserving+record+linkage

[21] Statistics Canada, "Statistical Methodology Research and Development Program Achievements, 2021/2022." Accessed: Jan. 04, 2024. [Online]. Available: https://www150.statcan.gc.ca/n1/pub/12-206-x/2022001/01-eng.htm

[22] "Advantages of Homomorphic Encryption - IEEE Digital Privacy." Accessed: Jan. 04, 2024. [Online]. Available: https://digitalprivacy.ieee.org/publications/topics/advantages-of-homomorphic-encryption

[23] "Microsoft SEAL: Fast and Easy-to-Use Homomorphic Encryption Library," Microsoft Research. Accessed: Jan. 04, 2024. [Online]. Available: https://www.microsoft.com/en-us/research/project/microsoft-seal/

[24] "PALISADE Homomorphic Encryption Software Library – An Open-Source Lattice Crypto Software Library." Accessed: Jan. 04, 2024. [Online]. Available: https://palisade-crypto.org/

[25] "TFHE Library," Inpher. Accessed: Jan. 04, 2024. [Online]. Available: https://inpher.io/tfhe-library/

[26] "What is Concrete?" Accessed: Jan. 04, 2024. [Online]. Available: https://docs.zama.ai/concrete

[27] "Nitro Enclaves," Amazon Web Services, Inc. Accessed: Jan. 04, 2024. [Online]. Available: https://aws.amazon.com/ec2/nitro/nitro-enclaves/

[28] A. Hall, "Split Neural Networks on Pysyft and Pytorch," OpenMined Blog. Accessed: Jan. 08, 2024. [Online]. Available: https://blog.openmined.org/split-neural-networks-on-pysyft/

[29] V. Turina, Z. Zhang, F. Esposito, and I. Matta, "Federated or Split? A Performance and Privacy Analysis of Hybrid Split and Federated Learning Architectures," in *2021 IEEE 14th International Conference on Cloud Computing (CLOUD)*, Sep. 2021, pp. 250–260. doi: 10.1109/CLOUD53861.2021.00038.

[30] "Learning with Privacy at Scale," Apple Machine Learning Research. Accessed: Aug. 22, 2023. [Online]. Available: https://machinelearning.apple.com/research/learning-with-privacy-at-scale

[31] B. McMahan and D. Ramage, "Federated Learning: Collaborative Machine Learning without Centralized Training Data," Google Research. Accessed: Nov. 20, 2023. [Online]. Available: https://blog.research.google/2017/04/federated-learning-collaborative.html?abstract_id=3808054

[32] I. Dayan et al., "Federated learning for predicting clinical outcomes in patients with COVID-19," Nat Med, vol. 27, no. 10, Art. no. 10, Oct. 2021, doi: 10.1038/s41591-021-01506-3.

[33] S. Truex *et al.*, "A Hybrid Approach to Privacy-Preserving Federated Learning." arXiv, Aug. 14, 2019. Accessed: Jan. 08, 2024. [Online]. Available: http://arxiv.org/abs/1812.03224

[34] S. Steinert, A. Benaissa, R. Roehm, and M. Hoeh, "A Privacy-Preserving Way to Find the Intersection of Two Datasets," OpenMined Blog. Accessed: Aug. 22, 2023. [Online]. Available: https://blog.openmined.org/private-set-intersection/

[35] "Evaluating the Performance of Privacy Preserving Record Linkage Systems (PPRLS)," Frederick National Laboratory for Cancer Research, Mar. 2023. Accessed: Jan. 30, 2024. [Online]. Available: https://surveillance.cancer.gov/reports/TO-P2-PPRLS-Evaluation-Report.pdf

[36] "Data Linkage and Identity Management - Privacy Protecting Record Linkage (PPRL)," Meeting Summary prepared by HLN Consulting, Mar. 2023. [Online]. Accessed: Jan. 04, 2024. Available: https://www.cdcfoundation.org/CDCFoundationPPRLSummary.pdf?inline

[37] Office for Civil Rights (OCR), "Guidance Regarding Methods for De-identification of Protected Health Information in Accordance with the Health Insurance Portability and Accountability Act (HIPAA) Privacy Rule," U.S. Department of Health and Human Services. Accessed: Nov. 20, 2023. [Online]. Available: https://www.hhs.gov/hipaa/for-professionals/privacy/special-topics/de-identification/index.html

[38] "IAPP Privacy Tech Vendor Report." IAPP, Oct. 2022 Accessed: Jan. 04, 2024. [Online]. Available: https://iapp.org/resources/article/privacy-tech-vendor-report/

[39] B. Lee *et al.*, "Protecting Privacy and Transforming COVID-19 Case Surveillance Datasets for Public Use," *Public Health Rep*, vol. 136, no. 5, pp. 554–561, Sep. 2021, doi: 10.1177/00333549211026817.

[40] K. E. Emam, *Guide to the De-Identification of Personal Health Information*. CRC Press, 2013.

[41] J. Near, D. Darais, and K. Boeckl, "Differential Privacy for Privacy-Preserving Data Analysis: An Introduction to our Blog Series," Cybersecurity Insights: a NIST blog, July 2020. Accessed: Nov. 16, 2023. [Online]. Available: https://www.nist.gov/blogs/cybersecurity-insights/differential-privacy-privacy-preserving-data-analysis-introduction-our

[42] The Population Reference Bureau and the U.S. Census Bureau's 2020 Census Data Products and Dissemination Team, "Why the Census Bureau Chose Differential Privacy," U.S. Census Bureau, Census Brief, Mar. 2023. [Online]. Available: chrome-extension://efaidnbmnnnibpcajpcglclefindmkaj/https://www2.census.gov/library/publications/decennial/2020/census-briefs/c2020br-03.pdf

[43] Guevara, M. "How we're helping developers with differential privacy." Google, Jan. 2021, Accessed: Jan. 04, 2024. [Online]. Available: https://developers.googleblog.com/2021/01/how-were-helping-developers-with-differential-privacy.html

[44] "Collaborative Challenges," *NIST*, Apr. 2023, Accessed: Jan. 04, 2024. [Online]. Available: https://www.nist.gov/itl/applied-cybersecurity/privacy-engineering/collaboration-space/challenges/collaborative-0

[45] "Prize Challenges," *NIST*, May 2023, Accessed: Jan. 04, 2024. [Online]. Available: https://www.nist.gov/itl/applied-cybersecurity/privacy-engineering/collaboration-space/challenges/prize-challenges

[46] U.S. Census Bureau, "Comparing Differential Privacy With Older Disclosure Avoidance Methods," 2021. Accessed: Jan. 04, 2024. [Online]. Available: https://www.census.gov/content/dam/Census/library/factsheets/2021/comparing-differential-privacy-with-older-disclosure-avoidance-methods.pdf

[47] C. Dwork, N. Kohli, and D. Mulligan, "Differential Privacy in Practice: Expose your Epsilons!," *Journal of Privacy and Confidentiality*, vol. 9, no. 2, Art. no. 2, Oct. 2019, doi: 10.29012/jpc.689.

[48] A. F. Barrientos, A. R. Williams, J. Snoke, and C. M. Bowen, "A Feasibility Study of Differentially Private Summary Statistics and Regression Analyses with Evaluations on Administrative and Survey Data." arXiv, Jun. 30, 2023. doi: 10.48550/arXiv.2110.12055.

[49] "Synthetic Data," IEEE Standards Association. Accessed: Nov. 20, 2023. [Online]. Available: https://standards.ieee.org/industry-connections/synthetic-data/

[50] US Census Bureau, "Synthetic SIPP Data," Census.gov. Accessed: Jan. 18, 2024. [Online]. Available: https://www.census.gov/programs-surveys/sipp/guidance/sipp-synthetic-beta-data-product.html

[51] K. Sallier and K. Burnette-Isaacs, "Unlocking the power of data synthesis with the starter guide on synthetic data for official statistics." Statistics Canada. Accessed: Sep. 08, 2023. [Online]. Available: https://www.statcan.gc.ca/en/data-science/network/synthetic-data

[52] A. G. Bates, I. Špakulová, I. Dove, and A. Mealor, "Synthetic data pilot," Office for National Statistics, Working paper series 16, Jan. 2019. Accessed: Jan. 30, 2024. [Online]. Available: https://www.ons.gov.uk/methodology/methodologicalpublications/generalmethodology/onsworkingpaperseries/onsmethodologyworkingpaperseriesnumber16syntheticdatapilot

[53] C. M. Bowen *et al.*, "A Synthetic Supplemental Public Use File of Low-Income Information Return Data: Methodology, Utility, and Privacy Implications," in *Privacy in Statistical Databases*, vol. 12276, J. Domingo-Ferrer and K. Muralidhar, Eds., in Lecture Notes in Computer Science, vol. 12276. , Cham: Springer International Publishing, 2020, pp. 257–270. doi: 10.1007/978-3-030-57521-2_18.

[54] C. Task, K. Bhagat, and G. Howarth, "SDNist: Deidentified Data Report Tool." in SDNist. Apr. 2023. doi: 10.18434/mds2-2943.

[55] "10 Misunderstandings Related to Anonymisation." Agencia Española de Protección de Datos (AEPD), Nov. 16, 2023. Accessed: Nov. 20, 2023. [Online]. Available: https://edps.europa.eu/data-protection/our-work/publications/papers/aepd-edps-joint-paper-10-misunderstandings-related_en

[56] V. Navale et al., "Development of an informatics system for accelerating biomedical research." F1000Research, Jul. 13, 2020. doi: 10.12688/f1000research.19161.2.

[57] N. Consortium, "N3C Privacy-Preserving Record Linkage and Linked Data Governance," Aug. 2021, doi: 10.5281/zenodo.5165212.

[58] "ISO/IEC 20889:2018." ISO, Nov. 2018. Accessed: Nov. 20, 2023. [Online]. Available: https://www.iso.org/standard/69373.html

[59] "ISO/IEC 20889:2018." ISO, Nov. 2018. Accessed: Nov. 20, 2023. [Online]. Available: https://www.iso.org/standard/69373.html

[60] "ISO/IEC 18033-6:2019." ISO, May 2019. Accessed: Nov. 20, 2023. [Online]. Available: https://www.iso.org/standard/67740.html

[61] "ISO/IEC 18033-6:2019." ISO, May 2019. Accessed: Nov. 20, 2023. [Online]. Available: https://www.iso.org/standard/67740.html

[62] "ISO/IEC 27559:2022." ISO, Nov. 2022. Accessed: Nov. 20, 2023. [Online]. Available: https://www.iso.org/standard/71677.html

[63] "ISO/IEC 27559:2022." ISO, Nov. 2022. Accessed: Nov. 20, 2023. [Online]. Available: https://www.iso.org/standard/71677.html

[64] "ISO/IEC 4922-1:2023." ISO, Jul. 2023. Accessed: Nov. 20, 2023. [Online]. Available: https://www.iso.org/standard/80508.html

[65] "xD - U.S. Census Bureau | Announcing the 2023 xD Emerging Technology Fellows." Accessed: Jan. 22, 2024. [Online]. Available: https://www.xd.gov/news/announcing-2023-xd-emerging-technology-fellows/

[66] S. Herath, H. Gelman, and L. McKee, "Privacy Harm and Non-Compliance from a Legal Perspective," *JCERP*, vol. 2023, no. 2, Oct. 2023, doi: 10.32727/8.2023.18.

[67] Y. Ejgenberg, M. Farbstein, M. Levy, and Y. Lindell, "SCAPI: The Secure Computation Application Programming Interface." 2012. Accessed: Jan. 08, 2024. [Online]. Available: https://eprint.iacr.org/2012/629

[68] A. Aly, E. Orsini, D. Rotaru, N. P. Smart, and T. Wood, "Zaphod: Efficiently Combining LSSS and Garbled Circuits in SCALE," in *Proceedings of the 7th ACM Workshop on Encrypted Computing & Applied Homomorphic Cryptography*, London United Kingdom: ACM, Nov. 2019, pp. 33–44. doi: 10.1145/3338469.3358943.

[69] "SCALE-MAMBA Software." Accessed: Jan. 08, 2024. [Online]. Available: https://homes.esat.kuleuven.be/~nsmart/SCALE/

[70] “swanky: A suite of rust libraries for secure computation.” Galois Inc., Dec. 18, 2023. Accessed: Jan. 08, 2024. [Online]. Available: https://github.com/GaloisInc/swanky

[71] L. Braun, D. Demmler, T. Schneider, and O. Tkachenko, “MOTION - A Framework for Mixed-Protocol Multi-Party Computation.” 2020. Accessed: Jan. 08, 2024. [Online]. Available: https://eprint.iacr.org/2020/1137

[72] “JIFF.” multiparty, Jan. 04, 2024. Accessed: Jan. 08, 2024. [Online]. Available: https://github.com/multiparty/jiff

[73] D. Gunning et. al., “CrypTen: A new research tool for secure machine learning with PyTorch.” Meta. Accessed: Jan. 08, 2024. [Online]. Available: https://ai.meta.com/blog/crypten-a-new-research-tool-for-secure-machine-learning-with-pytorch/

[74] “facebookresearch/CrypTen.” Meta Research, Jan. 07, 2024. Accessed: Jan. 08, 2024. [Online]. Available: https://github.com/facebookresearch/CrypTen

[75] K. Chandorikar, “Introduction to Federated Learning and Privacy Preservation using PySyft and PyTorch,” OpenMined Blog. Accessed: Jan. 08, 2024. [Online]. Available: https://blog.openmined.org/federated-learning-additive-secret-sharing-pysyft/

[76] Z. Zanussi, “Supervised Text Classification with Leveled Homomorphic Encryption”. *Proceedings of Statistics Canada Symposium 2021*, Statistics Canada, Oct. 2021. Accessed: Jan. 30, 2024. [Online]. Available: https://www150.statcan.gc.ca/n1/en/pub/11-522-x/2021001/article/00027-eng.pdf?st=IkT0zQ_A

[77] “ISO/IEC WD 18033-8,” ISO. Accessed: Jan. 08, 2024. [Online]. Available: https://www.iso.org/standard/83139.html

[78] M. Paulik *et al.*, “Federated Evaluation and Tuning for On-Device Personalization: System Design & Applications.” arXiv, Feb. 16, 2021. doi: 10.48550/arXiv.2102.08503.

[79] D. Madrigal, A. Manoel, J. Chen, N. Singal, and R. Sim, “Project Florida: Federated Learning Made Easy,” Jul. 2023, Accessed: Jan. 08, 2024. [Online]. Available: https://www.microsoft.com/en-us/research/publication/project-florida-federated-learning-made-easy/

[80] B. Nagy *et al.*, “Privacy-preserving Federated Learning and its application to natural language processing,” *Knowledge-Based Systems*, vol. 268, p. 110475, May 2023, doi: 10.1016/j.knosys.2023.110475.

[81] R. Sukanya, “Remote Data Science Part 2: Introduction to PySyft and PyGrid,” OpenMined Blog. Accessed: Jan. 08, 2024. [Online]. Available: https://blog.openmined.org/remote-data-science-part-2-introduction-to-pysyft-and-pygrid/

[82] “Flower: A Friendly Federated Learning Framework.” Flower. Accessed: Jan. 08, 2024. [Online]. Available: https://flower.dev/

[83] “TensorFlow Federated,” TensorFlow. Accessed: Jan. 08, 2024. [Online]. Available: https://www.tensorflow.org/federated

[84] “IBM Federated Learning,” IBM. Accessed: Jan. 08, 2024. [Online]. Available: https://www.ibm.com/docs/en/watsonx-as-a-service?topic=models-federated-learning

[85] "Welcome to the Open Federated Learning (OpenFL) Documentation! — OpenFL 2023.9 documentation." OpenFL. Accessed: Jan. 08, 2024. [Online]. Available: https://openfl.readthedocs.io/en/latest/

[86] "Azure Machine Learning - ML as a Service | Microsoft Azure." Microsoft. Accessed: Jan. 08, 2024. [Online]. Available: https://azure.microsoft.com/en-us/products/machine-learning

[87] N. Truong, K. Sun, S. Wang, F. Guitton, and Y. Guo, "Privacy preservation in federated learning: An insightful survey from the GDPR perspective," *Computers & Security*, vol. 110, p. 102402, Nov. 2021, doi: 10.1016/j.cose.2021.102402.

[88] A. Lopardo, "What is Federated Learning?," OpenMined Blog. Accessed: Aug. 22, 2023. [Online]. Available: https://blog.openmined.org/what-is-federated-learning/

[89] "IEEE Standards Association," IEEE Standards Association. Accessed: Jan. 08, 2024. [Online]. Available: https://standards.ieee.org

[90] B. Kubo, "Confidential sharing of datasets of two mobile network operators: A case study for tourism statistics from the perspective of a technology provider," United Nations, Jan. 20, 2022. Accessed: Jan. 30, 2023. [Online]. Available: https://unstats.un.org/bigdata/blog/2021/expo2020/sessions/2-3/presentations/6_Baldur_KUBO_Privacy-preserving-location-data-analytics.pdf

[91] F. Ricciato, T. Siil, R. Talviste, B. Kubo, and A. Wirthmann, "A proof-of-concept solution for the secure private processing of longitudinal Mobile Network Operator data in support of official statistics". *Expert Meeting on Statistical Data Confidentiality*, UNECE, Dec. 2021. Accessed: Jan. 30, 2023. [Online]. Available: https://unece.org/sites/default/files/2021-12/SDC2021_Day3_Ricciato_AD.pdf

[92] "IEEE 2952-2023." CTS/ETSC - Emerging Technology Standards Committee. Accessed: Nov. 20, 2023. [Online]. Available: https://standards.ieee.org

[93] "Intel® Software Guard Extensions (Intel® SGX)." Intel. Accessed: Nov. 20, 2023. [Online]. Available: https://www.intel.com/content/www/us/en/architecture-and-technology/software-guard-extensions.html

[94] "Landscape Analysis of Privacy Preserving Patient Record Linkage Software (P3RLS)," National Cancer Institute (NCI), National Institutes of Health (NIH), Department of Health and Human Services (HHS), Final Report Prepared By Synectics for Management Decisions, Inc., Jan. 2020.

[95] A. N. Kho *et al.*, "Design and implementation of a privacy preserving electronic health record linkage tool in Chicago," *J Am Med Inform Assoc*, vol. 22, no. 5, pp. 1072–1080, Sep. 2015, doi: 10.1093/jamia/ocv038.

[96] J. H. Raad, E. Tarlov, A. N. Kho, and D. D. French, "Health Care Utilization Among Homeless Veterans in Chicago," *Military Medicine*, vol. 185, no. 3–4, pp. e335–e339, Mar. 2020, doi: 10.1093/milmed/usz264.

[97] L. B. Mirel, D. M. Resnick, J. Aram, and C. S. Cox, "A methodological assessment of privacy preserving record linkage using survey and administrative data," *Stat J IAOS*, vol. 38, no. 2, pp. 413–421, Jun. 2022, doi: 10.3233/sji-210891.

[98] D. B. Baker *et al.*, "Privacy-Preserving Linkage of Genomic and Clinical Data Sets," *IEEE/ACM Trans Comput Biol Bioinform*, vol. 16, no. 4, pp. 1342–1348, 2019, doi: 10.1109/TCBB.2018.2855125.

[99] "Notice of New Computer Matching Program | Food and Nutrition Service." USDA. Accessed: Jan. 08, 2024. [Online]. Available: https://www.fns.usda.gov/snap/fr-052523

[100] R. Schnell, "Privacy-Preserving Record Linkage in the context of a National Statistics Institute," GOV.UK, Jul. 2021. Accessed: Nov. 15, 2023. [Online]. Available: https://www.gov.uk/government/publications/joined-up-data-in-government-the-future-of-data-linking-methods/privacy-preserving-record-linkage-in-the-context-of-a-national-statistics-institute

[101] R. Schnell and C. Borgs, "Implementing Privacy-preserving National Health Registries," *Proceedings of Statistics Canada Symposium 2018*, Statistics Canada, 2018. Accessed: Jan. 30, 2024. [Online]. Available: https://www.statcan.gc.ca/en/conferences/symposium2018/program/09a3_schnell-eng.pdf

[102] C. Pow *et al.*, "Privacy-Preserving Record Linkage: An international collaboration between Canada, Australia and Wales," *International Journal of Population Data Science*, vol. 1, no. 1, Art. no. 1, Apr. 2017, doi: 10.23889/ijpds.v1i1.101.

[103] R. McMillan, C. Thayer, and B. Watts, "Using Common Hash Values as Linking Keys," Georgia Policy Labs. Accessed: Nov. 20, 2023. [Online]. Available: https://gpl.gsu.edu/publications/using-common-hash-values-as-linking-keys/

[104] "Cryptographic Standards and Guidelines | CSRC | CSRC," CSRC | NIST. Accessed: Nov. 20, 2023. [Online]. Available: https://csrc.nist.gov/Projects/cryptographic-standards-and-guidelines

[105] K. Sallier, "Toward more user-centric data access solutions: Producing synthetic data of high analytical value by data synthesis," *Statistical Journal of the IAOS*, vol. 36, no. 4, pp. 1059–1066, Jan. 2020, doi: 10.3233/SJI-200682.

[106] "RTI SynthPop$^{TM}$ - Synthetic Population Dataset | RTI." RTI. Accessed: Nov. 20, 2023. [Online]. Available: https://www.rti.org/focus-area/rti-synthpoptm

[107] Y. Lu, M. Shen, H. Wang, X. Wang, C. van Rechem, and W. Wei, "Machine Learning for Synthetic Data Generation: A Review." arXiv, Jan. 01, 2024. Accessed: Jan. 08, 2024. [Online]. Available: http://arxiv.org/abs/2302.04062

[108] *Synthetic Data for Official Statistics - A Starter Guide* | *UNECE*. United Nations, 2022. Accessed: Nov. 15, 2023. [Online]. Available: https://unece.org/statistics/publications/synthetic-data-official-statistics-starter-guide

[109] J. Near and D. Darais, "Differentially Private Synthetic Data," NIST. Accessed: Nov. 16, 2023. [Online]. Available: https://www.nist.gov/blogs/cybersecurity-insights/differentially-private-synthetic-data

[110] D. Buckley, "Twitter and OpenMined: Enabling Third-party Audits and Research Reproducibility over Unreleased Digital Assets - UN GWG on Big Data - Privacy Preserving Techniques Wiki - UN Statistics Wiki," UN Statistics Wiki, Feb. 2023. Accessed: Nov. 20, 2023. [Online]. Available: https://unstats.un.org/wiki/display/UGTTOPPT/14.+Twitter+and+OpenMined:+Enabling+Third-party+Audits+and+Research+Reproducibility+over+Unreleased+Digital+Assets

[111] "Jana: Private Data as a Service," Galois Inc. Accessed: Jan. 31, 2023. [Online]. Available: https://galois.com/project/jana-private-data-as-a-service/

[112] "2019 Restricted-Use Linked Mortality Files," National Center for Health Statistics Data Linkage. Accessed: Jan. 31, 2023. [Online]. Available: https://www.cdc.gov/nchs/data-linkage/mortality-restricted.htm

[113] "Office of Shared Solutions and Performance Improvement (OSSPI); Chief Data Officers Council (CDO); Request for Information-Synthetic Data Generation," Federal Register: The Daily Journal of the Federal Government, Jan. 5, 2024. Accessed: Jan. 31, 2023. [Online]. Available: https://www.federalregister.gov/documents/2024/01/05/2024-00036/office-of-shared-solutions-and-performance-improvement-osspi-chief-data-officers-council-cdo-request

# Appendix A: Approach to the Landscape Analysis

As the field of PPTs is both broad, covering many different types of technologies ranging from identity protection and legal compliance to cryptographic methodologies, a detailed methodical approach needed to be taken to broadly understand the types of PPTs that exist, and to focus upon the PPTs that would be useful for the development of an NSDS. As such, the approach taken to conduct this environmental scan involved not only performing a literature review, but also performing outreach to practitioners; conducting interviews and facilitating group discussions panels to obtain first-hand knowledge from practitioners about the PPTs they have worked upon, the challenges they faced when developing, using, and deploying those PPTs, and their successes.

## Establishment of Inclusion/Exclusion Criteria

As not all technologies that are called PPTs could help inform the establishment of an NSDS in terms of the shared services that are offered, the first step in this environmental scan is to find the technologies that could be informative to focus upon. To do so, we first determined inclusion/exclusion criteria to help filter down the breadth of technologies to the most relevant PPTs. From reading the ACDEB Year 2 report, and through discussions with NCSES, the inclusion/exclusion criteria that we established are described in this section.

### *Criteria 1: Technologies should primarily deal with controller-controller and controller-processor relationships.*

The vision of the NSDS is for various Federal Government departments to be able to share data to help perform analyses for evidence building and decision making. Within a shared services environment such as the NSDS, the data have already been gathered with the consent of data subjects and is under the management of data controllers. The PPTs should therefore help manage privacy concerns in the situations where data controllers (in this case the Federal Government departments) need to share data amongst each other to meet their goals, which includes situations where controllers need to utilize the services of data processors to help with data management, storage, and analysis.

### *Criteria 2: Technologies should enable data sharing and analytics.*

As one of the goals of a future NSDS is to help further data sharing for evidence-based decision making, enabling data to be shared, producing evidence, and performing analyses on the evidence are necessary functions that should be enabled by usage of PPTs. The outputs of the PPTs we examine should therefore be data sets that have lower obstacles for data sharing than traditional data sets (through reducing or removing identifiable or private information) and said data must still be able to be analyzed, either through traditional statistical methods or through usage of specialized tools.

### *Criteria 3: Technologies should not rely on sharing unencrypted identifiable or private information to function.*

Traditional methods of data sharing for evidence-based decision making often involve sharing PII. Sending such information could pose privacy and security risks to the rights and freedoms of the benefit receivers, requiring administrative and regulatory overhead to help with risk mitigation. The PPTs we look at should try to mitigate these risks without introduction of significant administrative and regulatory burden by reducing or removing the need for sharing identifiable or private information to conduct these activities.

## Literature Review

To gain a holistic understanding of the landscape of PPTs, we conducted a review of available literature about PPTs, covering their usage across academia, government, and industry over the past 10 years. This review included grey literature, such as technical reports, presentations, and working papers, covering new developments in the field. In performing this review, we used our experience and expertise, including utilization of RTI internal library services and subject-matter experts, and partnered with FMH, an organization with expertise performing environmental scans across industries.

The goals of the literature review were as follows: (1) identify PPTs currently in use and in development; (2) identify key practitioners of PPTs to engage with for interviews and participation in group discussions; and (3) find organizations that are using PPTs and determine which PPTs they are using and how they are using them.

### *How the Review Was Conducted*

The literature review began with reviewing reports that contain surveys of the PPT landscape conducted by other groups, such as the White House [2], the UN [11], the OECD [13], and the Federal Reserve [12]. These reports provided a baseline of knowledge about the PPTs that are of interest to governments and national statistical organizations and gave us an initial list of practitioners to contact. Due to the broad range of technologies that are called PPTs, we also looked for classification systems within the literature to help group similar PPTs, which helped us filter down the range of PPTs to the ones that satisfy the inclusion and exclusion criteria that we established.

Once we found the PPTs that satisfy the inclusion and exclusion criteria, we performed literature searches about those specific PPTs, specifically focusing on papers that describe projects and pilots that utilized those PPTs, finding details about the results of those projects and the lessons learned. During our interviews with PPT practitioners, we also learned about and subsequently reviewed other papers that have relevant use cases. We conducted our literature review to find information that could help answer the five questions posed by NCSES, finding, and synthesizing the common themes and trends that run across the different papers. We present the results of this analysis in Appendix B

In total, we conducted detailed reviews of 21 papers. Summaries of these papers are provided in the annotated bibliography (shared as a separate document).

## Interviews and Facilitated Group Discussions

To go beyond the information that we discovered within the literature review, we also conducted interviews and facilitated group discussions with PPT practitioners, with the goal of getting first-hand accounts of their experiences using the PPTs to answer the questions that NCSES posed in the solicitation. Outreach to PPT practitioners spanned individuals from government, academia, and industry, leveraging networks and contacts known to RTI and FMH, and included international representation as well.

We also invited practitioners to take part in facilitated group discussions to bring them together with NCSES and have a conversation regarding common themes that arose during the interviews. These included topics such as administrative and regulatory hurdles, bridging information gaps between technical experts and nontechnical administrators and decision-makers, standards setting, and training and talent growth. Through these discussions, we were not only able to further gather information about PPTs, but we were also able to foster a sense of community between practitioners from different sectors, supporting collaboration in the field of PPTs.

We present in this section the details about how we conducted the interviews and facilitated group discussions, including the structure of the events, the questions that we asked, the selection criteria we used for outreach, and summaries of the information gathered from the facilitated group discussions.

### *Selection Criteria for Interviewees*

We designed the selection criteria chosen to find PPT practitioners whom we wanted to interview to find leaders within the space who have real knowledge and experience applying PPTs to real-world use cases. We believe that the practitioners selected under these criteria are best suited to providing information that can inform how PPTs can be used within a future NSDS, as their experiences would be able to tell us about the real situations where the PPTs could work best and the challenges and hurdles that could arise. The criteria are as follows:

1. Practitioners chosen for interviews should have implemented their PPT within a real-world use case, with preference given to practitioners whose use cases are government-related evidence-building activities. Note that this does not mean the use case has to be from the US government, though we did specifically target US Federal Government employees to ensure that we covered use cases within the US Federal Government.
2. Practitioners should be leaders within their PPT field as demonstrated through satisfaction of at least one of the following criteria:

a. leadership on projects or technical leadership on implementation,
b. primary authorship on papers,
c. professional, academic or government conference presentation on PPTs,
d. a minimum of five years developing or implementing PPT technologies, and/or
e. privacy academic or professional credentials.

### *Interview Structure and Questions Asked*

The interviews were structured around 15 questions that we derived from the five questions that NCSES posed in the solicitation. We asked obtained the informed consent to be interviewed from each interviewee, as well as their consent to be recorded. The questions we asked, in the order that we asked them, are listed below.

1. What is the name of the PPT that you have used?
2. Was this PPT developed by yourself/your team, or by someone else? (If someone else, who developed it?)
3. How did you come to use this PPT?
4. What projects or pilots have you used this PPT for? Please describe the project.
5. What privacy concern did the technique address?
6. What lessons have you learned while using this PPT?
7. How well did the PPT work in your situation?
8. In your opinion, what are the use cases where this PPT might work best?
9. Are there situations where this PPT might not work well?
10. What steps are required to implement this PPT?
11. What challenges and barriers exist for usage of this PPT?
12. What advantages and disadvantages does this PPT have over other types of PPTs that you know?
13. Have you conducted training/communication on usage of this PPT? If so, what was the training/communication and what recommendations do you have from your experience?
14. Are there any new or emerging PPTs that you are developing or aware of?
    a. If yes- what do you know about this new or emerging PPT and its potential uses?
15. Are there any other practitioners of PPTs that you know of whom you could refer to us for an interview?

In total, we interviewed 10 Federal PPT practitioners and 12 nonfederal PPT practitioners over the course of this project.

### *Facilitated Group Discussions*

After conducting interviews with practitioners, we conducted three Facilitated Group Discussions. The first two were with federal PPT Practitioners, the other with nonfederal PPT Practitioners. In this section, we present the topics discussed, and a summary of the conversation.

### *Facilitated Group Discussion 1: Federal PPT Practitioners*

The first facilitated group discussion was conducted on October 19, 2023, bringing together PPT practitioners who worked within Federal Government. This discussion focused on topics that were discussed during interviews with other federal PPT practitioners, such as spurring PPT use within the government, legal, regulatory, and policy hurdles, communication of the technical aspects of PPTs for a non-technical audience, and current initiatives within the government that are using or investigating PPTs. The discussion also touched on the practitioners hopes for a future NSDS, which include housing a community of experts that can help with development and implementation of PPTs within shared services environments across the Federal Government, creating shared templates for non-disclosure agreements, authority to use and operate, and streamlined data sharing agreements, and educating key stakeholder groups through highlighting successful use cases, publishing outcomes, and bringing products and services for use.

### *Facilitated Group Discussion 2: Federal PPT Practitioners*

The second facilitated group discussion, conducted on November 29, 2023, was a continuation of the conversations had with PPT practitioners from the interviews and the first group discussion, diving deeper into the common themes that we found across the engagements with those practitioners. These themes included communications and trust-building with stakeholders, navigating legal, regulatory, and policy challenges, implementation of PPTs, setting up a culture of innovation, and identifying future opportunities.

Attendees had much to say about each of the themes, highlighting topics such as using pilots to demonstrate capabilities of PPTs to build trust, starting small by attempting to solve tractable problems using PPTs, and testing PPTs to identify solutions that work. They discussed how there is a lack of good communications materials for educating those without technical background in PPTs, and how many aspects of PPTs are lacking in clear meaningful definitions. Setting of formal standards could help with development of these materials and creating those definitions, however many PPTs are still early in the standards-setting process, with difficulties in determining where to start setting standards. They said that resolution of these issues would be key to building trust and speeding adoption, and that if these issues are not addressed, they could lead to issues with regulations and policies down the line that frustrate PPT adoption, such as what has been seen in Europe with GDPR data transport regulations.

The attendees also expressed a commitment towards further development of talent within the Federal Government and building a culture of innovation around PPTs. Examples they gave of activities to promote talent growth and innovation include prize challenges, the Collaborative Research Cycle,[44],[45] and participation within international groups such as the UN PET lab.[5] They highlighted that training new talent is key to future success, and they expressed a desire for new models for education to help spur innovation in PPTs. They expressed hope that the NSDS could become a platform for education and innovation for PPTs, providing not only a community of practice, but also educational materials for those without technical background in PPTs, and a sandbox for piloting PPTs within the Federal Government.

### *Facilitated Group Discussion 3: Nonfederal PPT Practitioners*

The third facilitated group discussion, conducted on January 12, 2024, consisted of conversations with nonfederal PPT practitioners about the common themes that arose in the interviews we conducted. These are the same themes that were discussed with federal practitioners, though the responses non- federal practitioners gave had nuanced differences from those given by federal practitioners. Some of these differences arose from the different experiences these practitioners have had in their work researching and implementing PPTs outside of the government. The practitioners included representatives from industry, with technology vendors and consultants as well as industry users of PPTs, academics researching PPTs, and advisors and practitioners from foreign governments.

Regarding the different themes, nonfederal practitioners echoed many of the same observations noted by federal practitioners , with some added details. For communication and trust building, they mentioned how understanding the requirements of the data systems in which PPTs are to be deployed  was key to managing the balance between privacy and utility, and that storytelling can enhance communication of this balance to make it understandable by stakeholders who do not have PPT expertise. There also need to be better tools to facilitate these conversations, hearkening back to the same need mentioned by federal PPT practitioners. Many times, these conversations are best had not through explaining how the PPT works, but by providing evidence that the PPT has worked for other organizations, or by showing how the PPT is able to meet standards set by standards-setting bodies. Stakeholders also view risk in different ways, with many within the legal space viewing privacy and risk in a binary fashion (i.e., protected vs. non-protected, any risk is negative). This contrasts with what PPTs can do, as usage does not eliminate risk, but manages it along a spectrum. They mentioned how some regulations set up clear frameworks for measuring privacy risk, such as the HIPAA Privacy Rule, however, most current laws leave that ambiguous. This, combined with risk aversion from most regulatory agencies, has made adoption of PPTs challenging.

In terms of implementation, the nonfederal practitioners stressed how industry vendors can provide access to both broad and specific expertise, allowing for faster solutioning thanks to their experience. They also mentioned how through use of vendors there can be shifting of liability, which is important to some organizations. Academic practitioners stressed that while industry has the advantage when it comes to building and supporting PPT technologies, they can provide education to develop talent and spaces for innovative research. All practitioners stressed starting small when it comes to implementation, defining the problem well, creating a prototype that demonstrates viability, and educating stakeholders on how it works, before attempting adoption. They also mentioned that engagement of all the key stakeholder groups, business leadership, legal and policy, IT, and the end users of the data, and all groups coming to a mutual understanding of each other's needs, is key for expediting implementation.

Practitioners mentioned that despite the talent that exists within academia and industry, missing still are integrators who can put the pieces together. They are optimistic that training to meet needs will be developed as demand for PPTs increases. They mentioned that sandbox environments, both regulatory sandboxes such as what was set up by the United Kingdom Information Commissioners Office, and sandboxes for technology pilots such as the UN PET Lab, have been beneficial for both development of talent and development of PPTs in general.

# Appendix B: Detailed information about each PPT

Definitions of each of the listed PPTs, along with information that helps to answer the questions that NCSES has posed, and the technological maturity assessment, are given in the following sections.

## Secure Multiparty Computation (sMPC)

| | |
|---|---|
| **Definition** | Computational techniques involving data from multiple parties while preventing any party from learning about data that is not theirs beyond the results of the computation. |
| **What projects or pilots has this PPT been used for?** | ▪ Privacy Preserved Data Sharing for Evidence-Based Policy Decisions: A Demonstration Project Using Human Services Administrative Records for Evidence-Building Activities [17]<br>▪ Italian National Institute of Statistics and Bank of Italy: Enriching data analysis using privacy preserving record linkage [20]<br>▪ Boston Women's Workforce Council gender and racial wage gap study [16]<br>▪ Sharing Sensitive Department of Education Data Across Organizational Boundaries Using Secure Multiparty Computation [18]<br>▪ STATISTICS CANADA: MEASURING THE COVERAGE OF A DATA SOURCE USING A PRIVATE SET INTERSECTION [21]<br>▪ VaultDB: A Real-World Pilot of Secure Multi-Party Computation within a Clinical Research Network [19] |
| **What lessons were learned while using this PPT?** | ▪ Can result in significant computational overhead as complexity of computations increases. Efficient computation can only be performed on simple operations, such as sums or averages. Design of more complex computations would require technical and cryptographic expertise. [11]<br>▪ Does not necessarily need to be implemented in a cryptographic manner, as trusted execution environments can also enable sMPC. [17]<br>▪ Private Set Intersection (PSI) techniques can eliminate the need for a trusted third party when conducting privacy preserving record linkage. [34] |
| **What are the use cases where this PPT works best?** | ▪ Cases where there is a need to ensure Input Privacy amongst all parties.<br>▪ Ones that require mostly local operations on the shares with not many interactions among the parties.<br>▪ Distributed voting, private bidding, and auctions, sharing of signature or decryption functions, private set intersection and private information retrieval |
| **In what situations does this PPT not work?** | ▪ Datasets larger than a few thousand records will be slow to process. Floating point operations are much less easily represented and can require orders of magnitude more resources. Computations that rely on generative functions such as random number generation are also typically slow. [11] |

| | |
|---|---|
| **What challenges and barriers exist for usage of this PPT?** | ▪ sMPC technology performance depends heavily on the functions to be securely computed. For general computation such as the calculations needed to process typical relational database query operators, recent results show a slowdown up to 10,000 times. [18]<br>▪ sMPC protocols may need to be customized to the use case.<br>▪ Typically requires expert cryptographers to be implemented well, limiting the number of providers who can support development of sMPC PPTs and increasing the cost of development, deployment, and maintenance. |
| **What steps are needed to implement this PPT?** | ▪ Multiple frameworks exist for sMPC, such as:<br>– SCAPI (from Bar-Ilan University) - an API over various sMPC primitives [67]<br>– SCALE-MAMBA (from KU Leuven) - a complete sMPC system [68], [69]<br>– swanky (from Galois Inc.) - a set of Rust libraries for secure sMPC with garbled circuit, oblivious transfer, private set intersection protocol [70]<br>– Jana (from Galois Inc.) - a private data as a service model funded by the Defense Advanced Research Projects Agency's Brandeis program [17], [111]<br>– Motion (from TU Darmstadt, Aarhus University and the University of Hamburg) - a mixed protocol sMPC framework [71]<br>– JIFF (from Boston University) - a library allowing users to build applications JavaScript on top of sMPC protocols [72]<br>– CrypTen (from Facebook) - secure training and inference of machine learning models using sMPC [73], [74]<br>– Pysyft (from Openmined) – open source python library for sMPC [75] |
| **What advantages and disadvantages does this PPT have over other types of PPTs?** | ▪ In contrast to homomorphic encryption, which currently only supports polynomial functions, general sMPC offers a broader set of possible operations. [11]<br>▪ Does not necessarily need to rely on cryptographic methods, as hardware-based trusted execution environments can allow for sMPC as well, which alleviates many of the computational issues that arise from the cryptographic methods, and allows for running regular code instead of functions based on encrypted logic circuits. [17], [18] |

### *Technological Maturity Assessment*

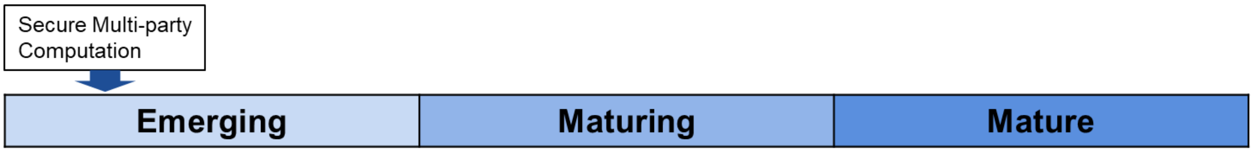

| Level of Standards Setting | Ease of Use | Public Trust |
|---|---|---|
| ▪ **Existence of formal standards**<br>– **ISO/IEC 4922-1:2023** [64]<br>– Terminology standard – Experts have agreed about how to talk about this PPT. | ▪ **Commercial tool availability**<br>– Most tools are prototypes or proof of concepts.<br>▪ **Expertise required to use the PPT.**<br>– Requires experienced cryptographers familiar | ▪ **Level of public understanding on how the PPT operates.**<br>– Low – significant technical knowledge about cryptography is needed for understanding how the |

| Level of Standards Setting | Ease of Use | Public Trust |
|---|---|---|
| - **Where the PPT sits within the standards-setting process**<br>– Early stages – experts have just begun to gather for discussions and workshops.<br>- **The parts of the PPT that still require setting of standards.**<br>– Many parts including acceptable cryptographic techniques, systems architecture, security model, etc. | with sMPC techniques for implementation.<br>- **Amount of customization and optimization needed.**<br>– Significant customization needed for each use case for optimal results | privacy mechanisms work.<br>- **Amount of public knowledge and scrutiny about the PPT**<br>– Low - currently most of these techniques remain in the realm of academic research.<br>- **Level of difficulty as to informing the public about how the PPT works.**<br>– High – few materials for ease of explanation to those without technical knowledge |

## Privacy Preserving Record Linkage (PPRL)

| | |
|---|---|
| **Definition** | Enables two or more parties, which both have a set of data, to compare these data sets without giving up on their individual data privacy. These parties compute the intersection of their data by encrypting identifiers that are used for linkage and linking upon the encrypted codes. |
| **What projects or pilots has this PPT been used for?** | - NIH BRICS System [56]<br>- N3C PPRL System [57]<br>- Landscape Analysis of Privacy Preserving Patient Record Linkage Software (P3RLS) [94]<br>- Design and implementation of a privacy preserving electronic health record linkage tool in Chicago [95]<br>- Health Care Utilization Among Homeless Veterans in Chicago [96]<br>- A methodological assessment of privacy preserving record linkage using survey and administrative data [97]<br>- Privacy Preserved Data Sharing for Evidence-Based Policy Decisions: A Demonstration Project Using Human Services Administrative Records for Evidence-Building Activities [17]<br>- Privacy Preserving Linkage of Genomic and Clinical Data Sets [98]<br>- Evaluation of Privacy Preserving Record Linkage Solutions to Broaden Linkage Capabilities in Support of Patient-Centered Outcomes Research Objectives [35]<br>- SNAP Computer Matching Program [99]<br>- Privacy Preserving Record Linkage in the context of a National Statistics Institute [100]<br>- Implementing Privacy preserving National Health Registries [101]<br>- Privacy Preserving Record Linkage: An international collaboration between Canada, Australia and Wales [102] |

| | |
|---|---|
| | ▪ USING COMMON HASH VALUES AS LINKING KEYS; A Solution for Identifying Linkage Keys (SILK) [103] |
| **What lessons were learned while using this PPT?** | ▪ Enables sharing of anonymized data, which can simplify policy compliance as no identifiable information needs to be shared between different parties to enable data linkage.<br>▪ Combinations of data from different sources can fill in gaps in the data, providing more complete views of the data subject.<br>▪ PPRL may enable agencies to share data reliably across programs and jurisdictions where sharing identifiable data are not allowed. This especially applies to sensitive data types, such as data about vulnerable populations.<br>▪ Can potentially achieve low false positive matching rates and provide high quality linkages, though that is dependent on the cleanliness of the underlying identifiable data.<br>▪ The linkage process can be privacy preserving in that no identifiable information is shared between the parties conducting linkage, however the resulting linked dataset may not have the same privacy protections as the individual de-identified data sets as linkage adds information that could increase the risk of re-identification.<br>▪ Many PPRL solutions rely on having an honest broker either store the keys that are used to encrypt PII to generate tokenized IDs, or lookup tables that match between different sets of tokenized IDs. There are a family of protocols that do not rely on honest brokers to perform these kinds of matches known and Private Set Intersection (PSI), which rely on purely cryptographic methods to enable the identification of matched records, without ever needing to share those records with another party. Methods that rely upon honest brokers are more mature and easily implemented, however the honest broker adds a central point of failure to the system. PSI methods can remove that central point of failure; however, they are harder to implement, requiring cryptographic expertise to be performed properly. |
| **What are the use cases where this PPT works best?** | ▪ Deduplication of records and linking data sources between data from different controllers.<br>▪ Systems where having an outside vendor provide the PPRL solution does not pose a problem. |
| **In what situations does this PPT not work?** | ▪ When there is not sufficient identifiable information in the original data to perform matching, PPRL will not be able to enable accurate linkages between records.<br>▪ When the original data being used within the PPRL system is not clean, it can result in false matches or matches being unable to be performed.<br>▪ When outside vendors cannot be used, there may be further barriers in designing a PPRL system from scratch. |
| **What challenges and barriers exist for usage of this PPT?** | ▪ Many administrative and policy barriers exist, such as establishment of data use and data sharing agreements, ethics reviews, and opposition from administrators who may not understand the security models that enable PPRL.<br>▪ Reliance upon PPRL vendors may be problematic for the long-term sustainability of a PPRL system.<br>▪ Increased risks of re-identification of data can arise from data linkages. |
| **What steps are needed to implement this PPT?** | ▪ Evaluation of the PPRL method to ensure that privacy and security concerns that arise from data sharing are addressed (i.e., the method does not inadvertently leak PII at some point in the process, and that it is able to satisfy the data sharing needs without sharing PII) |

| | |
|---|---|
| | ▪ Input data into the PPRL system should be cleaned and should follow a single standard, otherwise false matches and mismatches may occur.<br>▪ Selection of a vendor with a proven PPRL method or incorporate expertise from experienced architects of PPRL systems.<br>▪ Selection of an honest broker with clear rules and governance over what data the honest broker is allowed to see and what they are allowed to do.<br>▪ Determine and obtain agreement with all parties about what the outputs of the PPRL process will be (aggregate statistics, linked microdata, or identification of matched records)<br>▪ Provide clear documentation and explanations of the privacy and security model to administrators, regulators, and leadership to obtain their support.<br>▪ After linkage, if the linked data are to be made available, then a privacy evaluation should be performed upon the linked data to ensure that it has the same privacy protections as the original unlinked data. |
| **What advantages and disadvantages does this PPT have over other types of PPTs?** | ▪ Advantages:<br>▪ PPRL using hash matching has many successful use cases and mature proven technology.<br>▪ Easy to implement from a technical perspective as there are many tools available, both commercial and open source.<br>▪ PPRL systems do not necessarily require sharing any data with a third-party honest broker, as PSI methods can reliably result in identification between two intersections of data without any third-party.<br>▪ Disadvantages<br>▪ Reliance upon an honest broker produces a single source of failure in many PPRL systems.<br>▪ Linked data may not have the same privacy guarantees as the unlinked de-identified data.<br>▪ Using vendor tools could put PPRL systems at risk of vendor locking. Careful systems design is needed to avoid that from happening. |

### *Technological Maturity Assessment: PPRL based on Private Set Intersection*

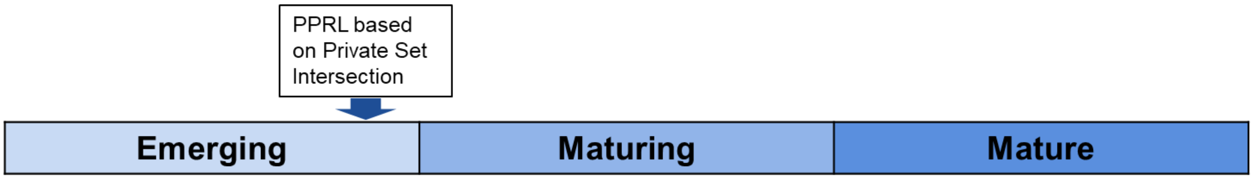
PPRL based on Private Set Intersection

Emerging

Maturing

Mature

| Level of Standards Setting | Ease of Use | Public Trust |
|---|---|---|
| ▪ **Existence of formal standards**<br>– ISO/IEC 18033-6:2019 [60]<br>– Encryption methods specified in the Homomorphic Encryption standard can be used for private set intersection. | ▪ **Commercial tool availability**<br>– Some commercial tools are available for private set intersection-based record linkage.<br>▪ **Expertise required to use the PPT.**<br>– Requires experienced cryptographers familiar | ▪ **Level of public understanding on how the PPT operates.**<br>– Low – significant technical knowledge about cryptography is needed for understanding how the privacy mechanisms work. |

| Level of Standards Setting | Ease of Use | Public Trust |
|---|---|---|
| ▪ **Where the PPT sits within the standards-setting process**<br>– Early stages – no formal standards-setting specific for PSI, however it is a part of standards for sMPC and Homomorphic Encryption<br>▪ **The parts of the PPT that still require setting standards.**<br>– Protocols to enable private set intersection | with private set intersection protocols for implementation.<br>▪ **Amount of customization and optimization needed.**<br>– Some customization needed for each use case for optimal results | ▪ **Amount of public knowledge and scrutiny about the PPT**<br>– Low - currently most of these techniques remain in the realm of academic research.<br>▪ **Level of difficulty as to informing the public about how the PPT works.**<br>– High – few materials for ease of explanation to those without technical knowledge |

### *Technological Maturity Assessment: PPRL Based on Tokenization*

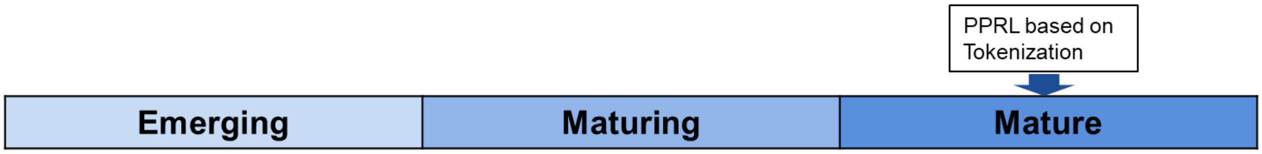

PPRL based on Tokenization

Emerging | Maturing | Mature

| Level of Standards Setting | Ease of Use | Public Trust |
|---|---|---|
| ▪ **Existence of formal standards**<br>– NIST Cryptographic Standards and Guidelines [104]<br>▪ **Where the PPT sits within the standards-setting process**<br>– NIST has produced an authoritative list of cryptographic algorithms for hashing and encryption that are secure to use, which can be employed within PPRL.<br>▪ **The parts of the PPT that still require setting standards.**<br>– No formal standard on PPRL architecture exists, however industry practices are fairly uniform | ▪ **Commercial tool availability**<br>– Many commercial and open-source tools for implementation of different PPRL methods<br>▪ **Expertise required to use the PPT.**<br>– Some expertise on record linkage is needed to identify the best fields to link upon<br>– Some expertise on use of hashing and encryption algorithms is necessary.<br>▪ **Amount of customization and optimization needed.**<br>– For rules-based de-identification, very little to none<br>– Expert determination based de-identification can have a high degree of customization | ▪ **Level of public understanding on how the PPT operates.**<br>– High – NIST approved hashing and encryption algorithms have undergone public review and commentary and has made it into laws and regulations.<br>▪ **Amount of public knowledge and scrutiny about the PPT**<br>– High – These algorithms are well-known and employed in many technologies and services.<br>▪ **Level of difficulty as to informing the public about how the PPT works.**<br>– Medium – The details of how these algorithms function can be very difficult to explain to non-cryptographers, but there is a high level of trust in these algorithms and in the process that developed them, so detailed technical explanations are often not necessary |

## Homomorphic Encryption

| | |
|---|---|
| **Definition** | Cryptographic techniques that allow for computation over encrypted data, so that no party other than the party providing the data learns anything about the data. Outputs from computations are encrypted as well so that only the party providing the data can decrypt and view them. |
| **What projects or pilots has this PPT been used for?** | ▪ STATISTICS CANADA: TRAINING A MACHINE LEARNING MODEL FOR PRIVATE TEXT CLASSIFICATION USING LEVELED HOMOMORPHIC ENCRYPTION [76] |
| **What lessons were learned while using this PPT?** | ▪ Can result in significant computational overhead as complexity of computations increases. Efficient computation can only be performed on simple operations, such as sums or averages. Design of more complex computations would require technical and cryptographic expertise. [22] |
| **What are the use cases where this PPT works best?** | ▪ Performing analytics on data at-rest within secure storage environments, where decryption of the data is to be avoided to prevent exposure to third parties (such as cloud providers.<br>▪ Image analysis upon encrypted patient medical imaging data. [11] |
| **In what situations does this PPT not work?** | ▪ Not practical in situations where you are not performing relatively simple computations on small amounts of encrypted data. |
| **What challenges and barriers exist for usage of this PPT?** | ▪ Can result in a high computational overhead and large expansion of data representation. Can be many orders of magnitude slower than plaintext calculations.<br>▪ Homomorphic encryption is a low level cryptographic primitive and building secure protocols from it is difficult without the help of a cryptography expert. Without expert guidance there can be security gaps in the final system. [11] |
| **What steps are needed to implement this PPT?** | ▪ There are multiple open-source homomorphic encryption libraries that implement different Homomorphic Encryption techniques. These include:<br>– Microsoft SEAL - implementing both BFV and CKKS schemes [23]<br>– PALISADE - supporting a range of different schemes and variants thereof including not BFV, BGV, CKKS, Levelled Somewhat Homomorphic Encryption and others [24]<br>– TFHE (from Inpher) - a implementation of TFHE - Fast Fully Homomorphic Encryption over the Torus [25]<br>– Concrete (from Zama.ai) - implementing a variant of TFHE [26] |
| **What advantages and disadvantages does this PPT have over other types of PPTs?** | ▪ Computations are performed on encrypted data so input privacy can be assured. Using Homomorphic Encryption increases the security and privacy levels while allowing data storage providers (cloud providers) to be the compute party.<br>▪ Relies upon cryptographic methods that require specialized expertise to implement.<br>▪ Computationally intensive when compared with non-cryptographic methods. |

### *Technological Maturity Assessment*

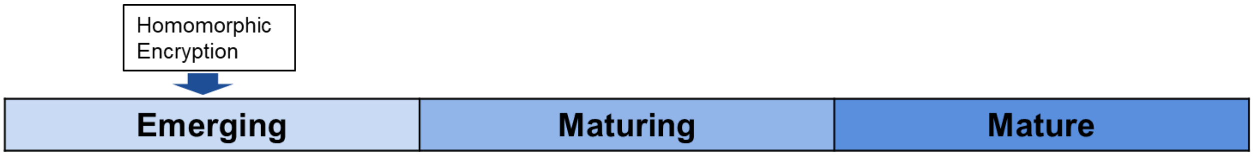

| Level of Standards Setting | Ease of Use | Public Trust |
|---|---|---|
| ▪ **Existence of formal standards**<br>– [**ISO/IEC 18033-6:2019**](#) [60]<br>– Standard specifies two acceptable Homomorphic Encryption processes and how they can be implemented.<br>– [**ISO/IEC WD 18033-8**](#) [77]<br>– Standard for fully homomorphic encryption is still in proposal phase.<br>▪ **Where the PPT sits within the standards-setting process**<br>– Early stages – certain parts of Homomorphic Encryption have been standardized, but Fully Homomorphic Encryption is in the earliest phases.<br>▪ **The parts of the PPT that still require setting standards.**<br>– All aspects of Fully Homomorphic Encryption, including the terminology, taxonomy, security model, assumptions, formats, architecture, etc. | ▪ **Commercial tool availability**<br>– Most tools are prototypes or proof of concepts.<br>▪ **Expertise required to use the PPT.**<br>– Requires experienced cryptographers familiar with Homomorphic Encryption techniques for implementation.<br>▪ **Amount of customization and optimization needed.**<br>– Significant customization needed for each use case for optimal results | ▪ **Level of public understanding on how the PPT operates.**<br>– Low – significant technical knowledge about cryptography is needed for understanding how the privacy mechanisms work.<br>▪ **Amount of public knowledge and scrutiny about the PPT**<br>– Low - currently most of these techniques remain in the realm of academic research.<br>▪ **Level of difficulty as to informing the public about how the PPT works.**<br>– High – few materials for ease of explanation to those without technical knowledge |

## Federated Learning

| | |
|---|---|
| **Definition** | Training of machine learning models by sending copies of a model to each place data resides and performing training on-site, eliminating the necessity of moving large amounts of data to a central location. The central server only receives updates to the model from each location, which are then aggregated to make the global model. |
| **What projects or pilots has** | ▪ [(Google) Federated Learning: Collaborative Machine Learning without Centralized Training Data ](#)[31] |

| | |
|---|---|
| **this PPT been used for?** | ▪ (Apple) Federated Evaluation and Tuning for On-Device Personalization: System Design & Applications [78]<br>▪ (Microsoft) Project Florida: Federated Learning Made Easy [79]<br>▪ Federated learning for predicting clinical outcomes in patients with COVID-19 [32] |
| **What lessons were learned while using this PPT?** | ▪ By itself, federated learning is not necessarily privacy preserving, as there is the potential of reverse engineering input data from local models. [80] However, federated learning can be combined with other methods, such as differential privacy or homomorphic encryption to ensure that privacy is protected. [33] Such combined methods involve adding noise to weights, or adding noise to aggregates, or using split learning, where only a portion of weights are sent and updated, or through encrypting the weights. These methods do not result in global models that are less performant than models that use the fully trained local models. |
| **What are the use cases where this PPT works best?** | ▪ Training neural networks on a set of distributed edge devices, such as mobile phones<br>▪ Leveraging sensitive smaller individual data sets that are stored locally on a network of entities or organizations with limited resources to collaboratively train a machine learning solution in a variety of domains, such as healthcare, finance, logistics, etc. |
| **In what situations does this PPT not work?** | ▪ Situations where data are centralized and not distributed across multiple mutually distrusting parties |
| **What challenges and barriers exist for usage of this PPT?** | ▪ Requires combination with other PPTs to protect privacy.<br>▪ Requires data owners to perform computations on the device that holds data. For some devices with limited computation capacity this may not be possible or economic. |
| **What steps are needed to implement this PPT?** | ▪ Requires agreement from all participating parties upon a neural network architecture to be used for learning.<br>▪ May require data sharing agreements to send updated models back to the central environment.<br>▪ There are a range open source enabling distributed learning:<br>– Syft + Grid (from OpenMined) - Syft combined federated learning, differential homomorphic encryption, and multi-party computation to enable private distributed learning. Grid provides an API to deploy Syft [81]<br>– Flower - a flexible framework for federated learning compatible with many ML frameworks (PyTorch, TensorFlow, MXNet and others) [82]<br>– TensorFlow Federated - Python library supported and used by Google [83]<br>– IBM Federated Learning - Python framework supporting a range of models including neural networks (in Keras, TensorFlow and PyTorch), linear regressions, decision trees [84].<br>– OpenFL - another Python library for federated learning from Intel [85]<br>– AzureML – federated learning in Microsoft Azure [86] |
| **What advantages and disadvantages does this PPT have over** | Advantages:<br>▪ Researchers can train models using private and sensitive data without having to handle any data - the data remains on the device and only learned model updates are transferred. |

| | |
|---|---|
| **other types of PPTs?** | ▪ Can be made compliant with data protection regulations like GDPR [87].<br>Disadvantages:<br>▪ The cost for implementing federated learning is higher than centralized data storage and processing, especially during the early phases of R&D. [88]<br>▪ Implementing FL is not enough to guarantee privacy, as private and sensitive information may be inferred from model updates. Adding other PPTs can help to mitigate this weakness. |

### *Technological Maturity Assessment*

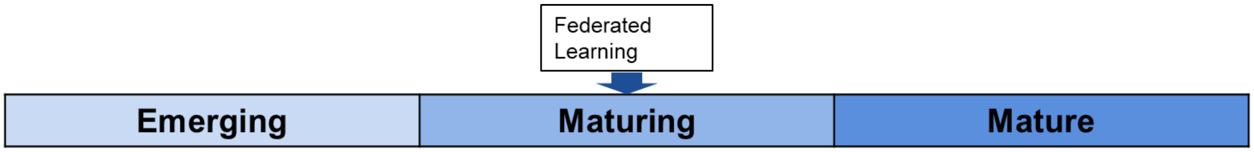

| Level of Standards Setting | Ease of Use | Public Trust |
|---|---|---|
| ▪ **Existence of formal standards**<br>– IEEE 3652.1-2020 [89]<br>– No formal standards exist for privacy preserving Federated Learning<br>▪ **Where the PPT sits within the standards-setting process**<br>– Middle stages – the architectural standard for Federated Learning has been set, however privacy considerations are not a part of the standard.<br>▪ **The parts of the PPT that still require setting standards.**<br>– Privacy preserving Federated Learning | ▪ **Commercial tool availability**<br>– Many commercial and open-source tools are available for implementing Federated Learning systems.<br>▪ **Expertise required to use the PPT.**<br>– Requires some expertise with machine learning to design the machine learning model, and some expertise with both hardware and software platforms to enable local learning and model updates.<br>▪ **Amount of customization and optimization needed.**<br>– Requires significant customization to the use case | ▪ **Level of public understanding on how the PPT operates.**<br>– Medium – The public understands the basics; however, the nuances of each system are not well explained, leading to limited understanding.<br>▪ **Amount of public knowledge and scrutiny about the PPT**<br>– Medium – Federated Learning has been implemented in various systems used by the public, bringing scrutiny.<br>▪ **Level of difficulty as to informing the public about how the PPT works.**<br>– Medium – companies have created communications to inform the public about their usage of Federated Learning, but the details of the systems are not well communicated. |

## Trusted Execution Environments (TEEs)

| | |
|---|---|
| **Definition** | A feature of modern CPU hardware that allows for execution of code in a way that mitigates input privacy, code privacy, and code assurance, by creating an execution environment that is separate from the rest of the computer system. Software enabled TEEs also exist, such as AWS Nitro Enclaves, and are available from many cloud providers. |
| **What projects or pilots has this PPT been used for?** | ▪ INDONESIA MINISTRY OF TOURISM: CONFIDENTIALLY SHARING DATASETS BETWEEN TWO MOBILE NETWORK OPERATORS VIA A TRUSTED EXECUTION ENVIRONMENT [90]<br>▪ Privacy Preserved Data Sharing for Evidence-Based Policy Decisions: A Demonstration Project Using Human Services Administrative Records for Evidence-Building Activities [17]<br>▪ UNITED NATIONS PET LAB: INTERNATIONAL TRADE [5]<br>▪ A proof-of-concept solution for the secure private processing of longitudinal Mobile Network Operator data in support of official statistics (ESTAT 2019.0232) [91] |
| **What lessons were learned while using this PPT?** | ▪ Scalability is typically not an issue as with cryptographic PPTs as software within TEEs can be written with regular code instead of relying on cryptographic operators. The software that is run is reviewed and approved by all parties before execution (to prevent any possible privacy breaches that can arise from running the software).<br>▪ Has a much more intuitive security model when compared with cryptographic solutions, making it easier to understand and accept. |
| **What are the use cases where this PPT works best?** | ▪ Large data sets with more complex operations required to be performed.<br>▪ Situations where either trusted TEE hardware can be provisioned, or trust can be placed in a software provider to provide a software-enabled TEE environment. |
| **In what situations does this PPT not work?** | ▪ Situations where parties do not have access to trusted hardware environments, or do not trust a software provider to provision a TEE environment |
| **What challenges and barriers exist for usage of this PPT?** | ▪ Hardware-based TEEs may suffer from reduced performance due to the limited amount of memory that they can access without resorting to encryption and decryption as data migrates on and off the CPU.<br>▪ TEE architectures are proprietary to CPU vendors, and so code developed for one hardware-based TEE is not necessarily likely to work on another vendor’s TEE, or possibly even on the TEE of a different generation within the same CPU family.<br>▪ Note that software-based TEEs do not suffer from the issues that arise from hardware TEEs, as memory can be dynamically allocated, and code can be ported from container to container.<br>▪ TEEs do not permit software to easily call operating system services, so a key part of a TEE is a software library that provides commonly used system calls inside the TEE (all necessary software needs to be provisioned beforehand within the TEE). |
| **What steps are needed to** | ▪ TEE technology typically requires hardware support inside a CPU. That support includes the use of dedicated on-chip memory in which to store data used frequently during the computation; specific features in virtual memory management that prevent other processes on the CPU from |

| | |
|---|---|
| **implement this PPT?** | accessing the memory space used by the TEE; hardware encryption support to encrypt and decrypt any data that must be moved out of the CPU and into system main memory; precautionary limits on advanced CPU features such as speculative execution or branch prediction, and so on.<br>▪ Today, secure enclaves are offered on all major cloud providers such as Google Cloud Platform, Azure, AWS, and IBM. However, if you are not planning to use a cloud environment, a purchase of enclave supported servers may be required. |
| **What advantages and disadvantages does this PPT have over other types of PPTs?** | ▪ Advantages:<br>▪ A characteristic of TEE solutions is that they prevent anyone – even users with control privileges on the host where the TEE runs – from learning anything about the code, data, or execution of that code inside the TEE.<br>▪ Can execute regular code, so does not have the scalability issues, or require the technical expertise of cryptographic solutions.<br>▪ Disadvantages<br>▪ The security model of TEEs is ultimately tied to trust in the physical hardware and/or hypervisor design. In the case of hypervisor-based TEEs, trust is in both the software security (i.e., that there are no bugs by the provider of the TEE) and that the hypervisor owner, typically a cloud provider, will not maliciously attack the system.<br>▪ TEEs are also a what-you-see-is-what-you-get (WYSIWYG) model. As arbitrary code can run in the TEE, parties using TEEs must understand and agree on exactly what is acceptable. This may include the specific versioning of libraries and frameworks, the sources they are from, and other security considerations.<br>▪ TEEs have no guarantees against timing-based attacks, and as such users should be exceptionally careful not to run code that signals specific proprietary sensitive input based on the number of branches created (how many times a loop runs or similar). |

### *Technological Maturity Assessment*

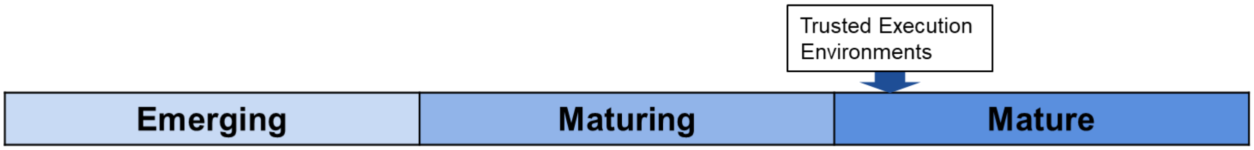

| Level of Standards Setting | Ease of Use | Public Trust |
|---|---|---|
| ▪ **Existence of formal standards**<br>– **IEEE 2952-2023** [92]<br>– **Intel® Software Guard Extensions (Intel® SGX)** [93]<br>▪ **Where the PPT sits within the standards-setting process**<br>– Standardization is complete, with formal standards that define requirements, and | ▪ **Commercial tool availability**<br>– Many hardware and software-based implementations of this technology<br>– Major cloud providers offer secure computing enclaves.<br>▪ **Expertise required to use the PPT.**<br>– Hardware-based TEEs require some expertise | ▪ **Level of public understanding on how the PPT operates.**<br>– High – Though the specific security implementations may require technical knowledge, the basic ideas of how a TEE operates can be explained in a simple manner. |

| Level of Standards Setting | Ease of Use | Public Trust |
|---|---|---|
| mature industry best practices.<br>- **The parts of the PPT that still require setting standards.**<br>– Regular updates to security profiles are still needed | with IT and networking to set up.<br>– Software-based TEEs can be provided by cloud service providers and require litter expertise to set up.<br>– Security reviews by qualified InfoSec personnel are needed.<br>- **Amount of customization and optimization needed.**<br>– TEEs can be used off the shelf or customized to needs, depending on the use case | - **Amount of public knowledge and scrutiny about the PPT**<br>– High – TEEs have been in use for over a decade in various projects that have high public visibility.<br>- **Level of difficulty as to informing the public about how the PPT works.**<br>– Low – Lots of material for public consumption exists about TEEs from both hardware and software providers |

## De-Identification

| | |
|---|---|
| **Definition** | Methods of transforming data sets to remove identifying information. These include statistical methods such as K-anonymity (Transforms a given set of k records in such a way that in the published version, each individual is indistinguishable from the others), and rules-based methods such as HIPAA Safe Harbor (removal of 18 types of identifying information). |
| **What projects or pilots has this PPT been used for?** | Most used method of protecting output privacy.<br>US census prior to 2020<br>HIPAA Privacy Rule [37] |
| **What lessons were learned while using this PPT?** | De-identification is most commonly used because it is widely understood and relatively simple to implement as compared to other PPTs.<br>De-identification is not foolproof, depends on the context of the data release, legal and contractual controls, and evaluation of the information available to recipients to protect individual privacy.<br>Cannot remove all risk of re-identification of data subjects, so de-identification is about managing risk levels to make them below an acceptable threshold [40]<br>There are a variety of methods to measure risk of re-identification, however methods of transforming data to reduce the risk of re-identification come to either masking/suppression (data deletion), generalization (increasing granularity), or noise addition.<br>What constitutes as identifying information depends on how the data was gathered and the context of the data release. |
| **What are the use cases where this PPT works best?** | Can be applied to nearly all data sets to protect output privacy. Does not protect input privacy. |

| | |
|---|---|
| **In what situations does this PPT not work?** | De-identification of unstructured data (images, videos, free text) is more difficult and involved than de-identification of structured data.<br>Not useful for situations where input privacy is a concern |
| **What challenges and barriers exist for usage of this PPT?** | Vulnerable to reidentification attack if additional public information is available.<br>Vulnerable to reconstruction or linkage attacks that can lead to the re-identification of a data subject if an adversary has relevant auxiliary information.<br>As more information becomes publicly available, the amount of identifying information has increased to the point where de-identification may not be possible under many circumstances |
| **What steps are needed to implement this PPT?** | HIPAA Safe Harbor calls for removal of 18 types of identifiers.<br>Other method is expert determination, where an expert using scientific and statistical methods measures the risk of re-identification upon a dataset and applies transformations to reduce the risk beneath an acceptable threshold, documenting the procedures to act as evidence [37]. |
| **What advantages and disadvantages does this PPT have over other types of PPTs?** | Advantages:<br>Written into regulation and legislation, giving this method official legal backing.<br>Relatively simple to implement.<br>Techniques are mature and well-understood.<br>Disadvantages:<br>Datasets could become vulnerable to re-identification as more data and better techniques become available.<br>Requires an expert to determine what information is identifying that warrants removal.<br>Not compositional, so multiple releases of the same data under these techniques can result in a catastrophic loss of privacy. |

### *Technological Maturity Assessment*

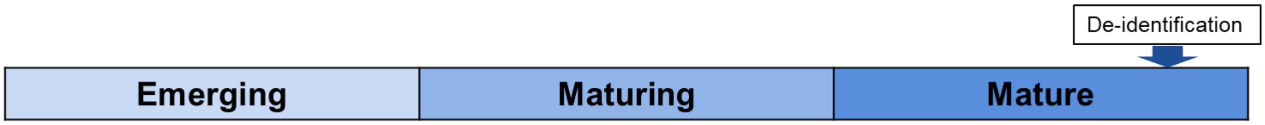

| Level of Standards Setting | Ease of Use | Public Trust |
|---|---|---|
| ▪ **Existence of formal standards**<br>– **HIPAA Privacy Rule De-Identification Standard** [37]<br>– **EU Guidance for Anonymization under GDPR** [55] | ▪ **Commercial tool availability**<br>– Many commercial and open-source tools for implementation of various de-identification methods<br>▪ **Expertise required to use the PPT.** | ▪ **Level of public understanding on how the PPT operates.**<br>– High – De-identification has undergone significant public review and commentary and has made it into laws and regulations. |

| Level of Standards Setting | Ease of Use | Public Trust |
|---|---|---|
| – **ISO/IEC 20889:2018** [59]<br>– **ISO/IEC 27559:2022** [63]<br>▪ **Where the PPT sits within the standards-setting process**<br>– Standards for de-identification have been defined and adopted under legal and regulatory frameworks.<br>▪ **The parts of the PPT that still require setting standards.**<br>– Worldwide adoption of a single accepted de-identification framework is underway | – Rules-based de-identification (such as under HIPAA Safe Harbor) does not require significant expertise.<br>– Expert determination requires a qualified expert with specialized skills and knowledge to perform an evaluation and de-identification.<br>▪ **Amount of customization and optimization needed.**<br>– For rules-based de-identification, very little to none<br>– Expert determination based de-identification can have a high degree of customization | ▪ **Amount of public knowledge and scrutiny about the PPT**<br>– High – De-identification techniques have been in use for decades and are the primary method for managing data privacy risk used by statistical agencies, governments, and the private sector.<br>▪ **Level of difficulty as to informing the public about how the PPT works.**<br>– Low – Significant materials exist to describe de-identification in an approachable manner. Rules-based de-identification is easy to explain. Expert determination has a higher level of difficulty. |

## Synthetic Data

| | |
|---|---|
| **Definition** | Creating a dataset containing brand new records using statistical or machine learning techniques that has similar aggregate statistical properties as the original dataset but has individual records that are significantly different from the original data. |
| **What projects or pilots has this PPT been used for?** | ▪ Survey of Income and Program Participation (SIPP) Synthetic Beta [50]<br>▪ Office for National Statistics: Trialing the use of synthetic data at the United Kingdom's national statistics institute [52]<br>▪ STATISTICS CANADA: TRIALLING THE USE OF SYNTHETIC DATA [105]<br>▪ RTI Synthpop [106]<br>▪ A Synthetic Supplemental Public-Use File of Low-Income Information Return Data: Methodology, Utility, and Privacy Implications (Urban Institute) [53]<br>▪ NCHS Public Use Mortality Files [112] |
| **What lessons were learned while using this PPT?** | ▪ Using a machine learning model to generate brand new records (the "model-based approach") has emerged as the preferred alternative for generating synthetic data. [107]<br>▪ There are challenges associated with developing synthetic for generic analytics. Synthetic data sets can only preserve some relationships found within the real data. [105]<br>▪ Combination of synthetic data with other privacy techniques such as differential privacy can provide a measure of the upper limit of the number |

| | |
|---|---|
| | of synthetic data sets that can be released before privacy could be compromised. |
| **What are the use cases where this PPT works best?** | ▪ Synthetic data sets can be used to give fine grain understanding of the original data without the risks and compliance hurdles.<br>▪ Validating a proof-of-concept or evaluating third-party solutions.<br>▪ Expand training data sets for AI systems that typically benefit from large training sets.<br>▪ Transforming old data into synthetic data is a way to keep the benefit of using the data for potential future studies while following data retention requirements. |
| **In what situations does this PPT not work?** | ▪ Synthetic data are not an option when one wants to ask questions in the future which are beyond the scope of the requirements when initially creating the data sets, as the synthetic data algorithm cannot guarantee that the specific characteristics required to answer such future questions will be preserved by the generating model. |
| **What challenges and barriers exist for usage of this PPT?** | ▪ Synthetic data generators may remember some personal information, especially when the original data are sparse, which is likely in high dimensional data sets such as images, text, or series of events, and the model has a large learning capacity, which is the case of most neural-network-based generative models.<br>▪ Very flexible models can "overfit," leading to potentially sensitive information influencing the synthetic data generation and hence to reidentification of certain samples.<br>▪ Synthetic data can be faithful for a limited number of predefined objectives but cannot be universally faithful. |
| **What steps are needed to implement this PPT?** | ▪ Synthetic data generators use deep learning techniques to learn the distributions and relationships within the original input data. Then, they use generative models to sample from learned distributions to produce new data. Techniques include the use of copulas, generative-adversarial networks (GANs) and variational auto-encoders (VAEs), amongst others. [108]<br>▪ Can be combined with other techniques such as differential privacy to create synthetic data sets that have measurable privacy guarantees. [109] |
| **What advantages and disadvantages does this PPT have over other types of PPTs?** | Advantages:<br>▪ Training a synthetic data generative model is a one-off exercise. Depending on the dataset size and the type of learning procedure it can incur a significant cost, but once the data aregenerated, using synthetic data for analysis or model training is identical to using the real data from the user's perspective.<br>Disadvantages:<br>▪ The main cost of using synthetic data is loss of utility. Unless the synthetic data are the original data, some queries on the synthetic data will differ from queries on the original dataset. |

***Technological Maturity Assessment***

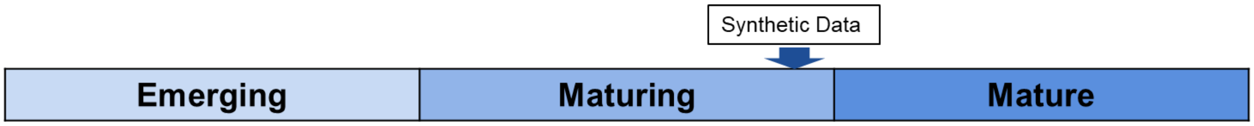

| Level of Standards Setting | Ease of Use | Public Trust |
|---|---|---|
| ▪ **Existence of formal standards**<br>– IEEE IC21-013-01 [49]<br>– No formal standards exist for privacy preserving synthetic data.<br>▪ **Where the PPT sits within the standards-setting process**<br>– Early stages – discussions are only beginning about setting standards for privacy preserving synthetic data.<br>▪ **The parts of the PPT that still require setting standards.**<br>– Privacy measures for synthetic data<br>• NIST tool **SDNist** does begin to establish a privacy metric standard [54] | ▪ **Commercial tool availability**<br>– Numerous commercial tools exist for generating synthetic data.<br>– Numerous open-source tools exist for generating synthetic data.<br>▪ **Expertise required to use the PPT.**<br>– Many commercial tools require little to no expertise, as the models and tools are streamlined for ease of use to the end users.<br>– Open-source tools require some technical expertise to operate.<br>▪ **Amount of customization and optimization needed.**<br>– Open-source tools require some optimization and customization.<br>– Commercial tools have few requirements from end users | ▪ **Level of public understanding on how the PPT operates.**<br>– Medium – Public has a good idea as to what synthetic data are, but there are some misconceptions about synthetic data and privacy protection that need correction.<br>▪ **Amount of public knowledge and scrutiny about the PPT**<br>– Medium – Many new synthetic data implementations have caught public awareness and scrutiny.<br>▪ **Level of difficulty as to informing the public about how the PPT works.**<br>– Medium – Synthetic data is not hard to explain to the public, however there are still misconceptions about synthetic data that need correction |

## Differential Privacy

| | |
|---|---|
| **Definition** | A mathematically formal definition of privacy based on the idea of the "differencing attack," to prevent the results of data queries from isolating the information about any single individual in the data. Differencing attacks are prevented through the addition of randomness to the results of queries based on a privacy budget (defined by the parameter epsilon). |
| **What projects or pilots has this PPT been used for?** | • 2020 US Census Aggregate Public Data Release [42]<br>• Twitter and OpenMined: Enabling Third-party Audits and Research Reproducibility over Unreleased Digital Assets [110]<br>• (Apple) Learning with Privacy at Scale [30] |

| | |
|---|---|
| | - A Feasibility Study of Differentially Private Summary Statistics and Regression Analyses with Evaluations on Administrative and Survey Data [48] |
| **What lessons were learned while using this PPT?** | - No clear consensus on how to choose epsilon, nor agreement on how to approach this and other key implementation decisions. There is little collaboration, information sharing, or publishing to advance critical reflection. Given the importance of these details there is a need for shared learning amongst the differential privacy community. [47]<br>- It is worth noting that almost all of the implementations surveyed by Prof. Cynthia Dwork err on the side of utility over privacy. [47] |
| **What are the use cases where this PPT works best?** | - The Differential Privacy formalism can be applied to any computation from a single database query to all the iterative steps needed to train a machine learning model. |
| **In what situations does this PPT not work?** | - Differential Privacy only addresses the privacy of an output of a flow of information (Output Privacy). It does not solve the privacy risks when managing input data between where it is collected, stored, and eventually processed (Input Privacy).<br>- Differential privacy is the wrong tool to use to study outliers, as it hides their presence or absence. It is also not the right tool for analyzing small data sets. This is because depending on the choice of epsilon, differential privacy may hide important differences in small populations or subpopulations of interest. |
| **What challenges and barriers exist for usage of this PPT?** | - Differential Privacy is still limited to simpler data types; it is challenging to manage tradeoffs between privacy, accuracy, or utility of data as data complexity increases.<br>- There is not yet a generalizable method of how to best set the privacy parameter to control the strength of the privacy guarantee while optimizing for accurate analytic results. Although parameter values epsilon and delta have a very precise statistical interpretation, there is no general application-agnostic recipe for choosing appropriate values of these parameters.<br>- Communication of what the privacy budget means to those without technical background in Differential Privacy can be challenging.<br>- They may require customized algorithms to be developed that are suited to the data type and planned data use.<br>- While certain algorithms may be differentially private, they may not be able to satisfy other privacy metrics, resulting in data that may not have as strong privacy protections as supposed. |
| **What steps are needed to implement this PPT?** | - Differential Privacy is achieved by the introduction of random noise as the privacy mechanism. Random noise can be introduced in a variety of fashions, both at the individual record level (e.g., through sampling) and at the aggregate level (e.g., through perturbation).<br>- Differential Privacy requires setting a privacy budget to limit harm. Each query upon the dataset uses up a part of the privacy budget. Once this budget has been exhausted, the dataset should not be queried again. |
| **What advantages and disadvantages does this PPT have over** | Advantages:<br>- Assumes all information is identifying information, eliminating the challenging (and sometimes impossible) task of accounting for all identifying elements of the data. |

| | |
|---|---|
| **other types of PPTs?** | • Resistant to privacy attacks based on auxiliary information, so it can effectively prevent the linking attacks that are possible on de-identified data.<br>• It is compositional in that the privacy loss of running two differentially private analyses on the same data can be determined by simply adding up the individual privacy losses for the two analyses.<br>• Traditional anonymization techniques need to make strong assumptions about the auxiliary information accessible to the recipient in their threat model. Differential privacy makes no such assumption and can theoretically protect against much stronger attackers.<br>Disadvantages:<br>• The main cost of using differential privacy is a loss in terms of output accuracy with respect to solutions for the same problem that do not provide output privacy. This cost depends on the level of privacy required (more privacy incurs more loss in accuracy), the number of individuals in the dataset (increasing the amount of data available reduces the accuracy loss), the number of queries to be made on the data, and the range of possible values for each individual. |

### *Technological Maturity Assessment*

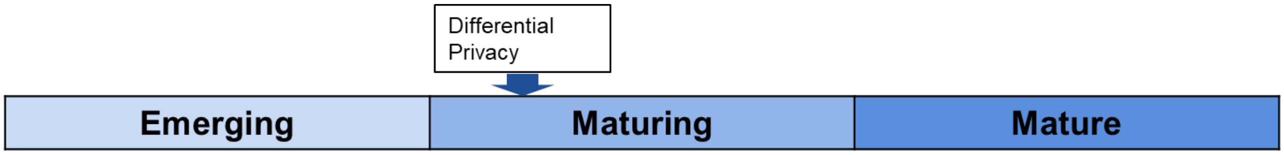

Differential Privacy

Emerging

Maturing

Mature

| Level of Standards Setting | Ease of Use | Public Trust |
|---|---|---|
| ▪ **Existence of formal standards**<br>– **ISO/IEC 20889:2018** [58]<br>– **ISO/IEC 27559:2022** [63]<br>– Terminology standard has been set, and there exists a framework standard for conformance.<br>▪ **Where the PPT sits within the standards-setting process**<br>– Middle stages – some standards have been set, but there are numerous portions that require standardization.<br>▪ **The parts of the PPT that still require setting standards.**<br>– Acceptable Differential Privacy algorithms | ▪ **Commercial tool availability**<br>– Some commercial tools are available implementing certain Differential Privacy algorithms.<br>▪ **Expertise required to use the PPT.**<br>– Requires some expertise to design algorithms that are differentially private and mathematically prove differential privacy.<br>▪ **Amount of customization and optimization needed.**<br>– Requires significant customization to the use case to ensure differential privacy | ▪ **Level of public understanding on how the PPT operates.**<br>– Low – Differential Privacy is complicated and requires significant technical knowledge to understand, and different algorithms have different Differential Privacy mechanisms that can get highly technical.<br>▪ **Amount of public knowledge and scrutiny about the PPT**<br>– Medium – implementation of Differential Privacy by companies like Apple have brought public scrutiny.<br>▪ **Level of difficulty as to informing the public about how the PPT works.**<br>– Medium – some materials have been created that |

| Level of Standards Setting | Ease of Use | Public Trust |
|---|---|---|
| – Setting the privacy budget | | make Differential Privacy approachable (NIST Differential Privacy blog series, Apple’s marketing), however the specific technical implementations are still difficult to explain |