Overcoming Data Barriers via Trustworthy Privacy-Enhancing Technologies

Demonstration Report

November 2023

GPAI | THE GLOBAL PARTNERSHIP ON ARTIFICIAL INTELLIGENCE

This report was developed by Experts and Specialists involved in the Global Partnership on Artificial Intelligence's project on 'Privacy-Enhancing Technologies'. The report reflects the personal opinions of the GPAI Experts and External Experts involved and does not necessarily reflect the views of the Experts' organisations, GPAI, or GPAI Members. GPAI is a separate entity from the OECD and accordingly, the opinions expressed and arguments employed therein do not reflect the views of the OECD or its Members.

Acknowledgements

This report was developed in the context of the *Privacy-Enhancing Technologies'*, with the steering of the Project Co-Leads and the guidance of the Project Advisory Group, supported by the GPAI Data Governance Working Group. The GPAI Data Governance Working Group agreed to declassify this report and make it publicly available.

Co-Leads:
Shameek Kundu*, Truera
Kim McGrail*, University of British Columbia

GPAI would like to thank its partners for this project, the Infocomm Media Development Authority of Singapore, as well as Singapore's Ministry of Health, and Digital Trust Centre. Without the efforts of these teams, this technical demonstration could not have happened.

GPAI would also like to acknowledge the tireless efforts of colleagues at the International Centre of Expertise in Montréal on Artificial Intelligence (CEIMIA) and GPAI's Data Governance Working Group. We are grateful, in particular, for the support of **Stephanie King** from CEIMIA, as well as the Data Governance Working Group Co-Chairs, **Maja Bogataj Jančič***, Intellectual Property Institute, and **Jeni Tennison***, Connected by Data.

* Expert
** Observer
† Invited Specialist
‡ Contracted Parties by the CofEs to contribute to projects

Citation

GPAI 2023. *Overcoming Data Barriers Trustworthy Privacy-Enhancing Technologies*, Report, November 2023, Global Partnership on AI.

Introduction

In 2022, the Data Governance Working Group kicked-off a project to explore possible applications of privacy-enhancing technologies (PETs+) in AI-for-social-good contexts. The scoping phase of this project, completed in July 2022 with support from Capgemini, recommended a use case focused on data sharing for the purposes of improving resilience of society to pandemics. The project partnered with Singapore's Infocomm Media Development Authority (IMDA) to conduct a demonstration of the use case. Singapore's Digital Trust Centre (DTC) acted as the delivery partner. With the demonstration project now complete, the IMDA/ DTC team has documented key "lessons learnt" from the project.

The Use Case

The initial assumption was that the PETs+ demonstration project would be executed in conjunction with GPAI's existing Pandemic Resilience project, which falls under the Responsible AI Working Group's work plan. However, challenges around data availability and resourcing necessitated the search for alternatives.

Eventually, a model from Singapore's Ministry of Health (MOH)—built to simulate SARS-CoV-2 infection transmission and control in large-scale multi-day events—was selected for the demonstration. The model was derived from raw contact data representing contact episodes (durations, numbers, key characteristics) between pairs of passengers and crew on a cruise. Combined with a disease transmission model, the pandemic model can simulate transmission onboard the cruise and how different interventions may lead to different outbreak dynamics.

The model was published[1] in 2022 and was built using contact data collected in 2020, with consent from passengers and crew of four cruises. The model is on GitHub[2] but the original confidential dataset is only available to the MOH's Communicable Diseases Division.

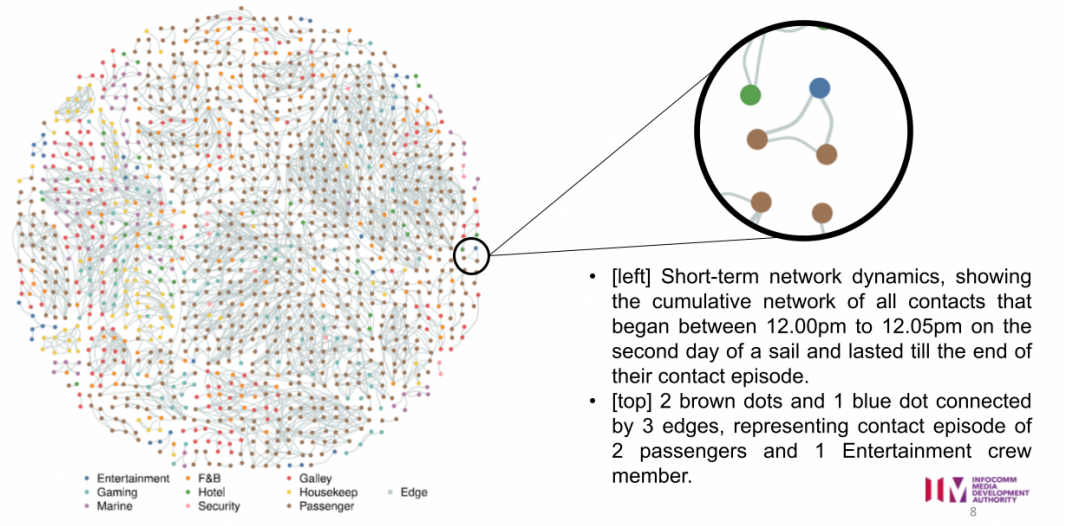

Figure 1: Illustration of MOH 'Cruise Network' Pandemic Model[3]

[1] Pung, Rachael; Firth, Josh A; Spurgin, Lewis G; Singapore CruiseSafe working group; CMMID COVID-19 working group; Lee, Vernon J; Kucharski, Adam J; (2022) Using high-resolution contact networks to evaluate SARS-CoV-2 transmission and control in large-scale multi-day events. Nature Communications, 13 (1). 1956-. ISSN 2041-1723 DOI: https://doi.org/10.1038/s41467-022-29522-y

[2] Rachael Pung Cruise Networks

[3] n1

A PET-Enabled Version of the Use Case

The use of PETs is attractive in the context of such a model, because disclosing contact time stamps could compromise private activities of passengers and crew. PETs could be employed to compute contact duration values without disclosing contact time stamps. They could enable such models to be built on near-real-time data ("now casting"), such as payment transactions and use of public transport, without the data sources having to disclose their original sensitive datasets.

The goal of the demonstration project was to test this hypothesis; specifically, whether PET-protected data could reproduce outcomes of the original pandemic model. Figure 2 provides a high-level overview of the experiment design.

Two PETs were selected as alternative methods for this use case:

- **Homomorphic Encryption (HE)**, which allows computations to be performed on encrypted data (ciphertext) without the need to decrypt the original, sensitive data. This allows a Data Controller to encrypt sensitive data, grant access of ciphertext to an external Data Analyst to run queries, and allow only the output of the queries to be decrypted. In this project, contact durations were computed on pairs of HE-protected time stamps. The contact durations are then used to construct a contact network to simulate disease transmission.
- **Differential Privacy (DP)**, which adds sufficient "noise" to protect a single entry from being reidentified. The Data Controller of a sensitive dataset allocates a certain amount of noise according to target "privacy budget" before granting access of the DP-protected dataset to an external Data Analyst. In this project, DP-protected time stamps induced with "noise" were used to compute contact durations.

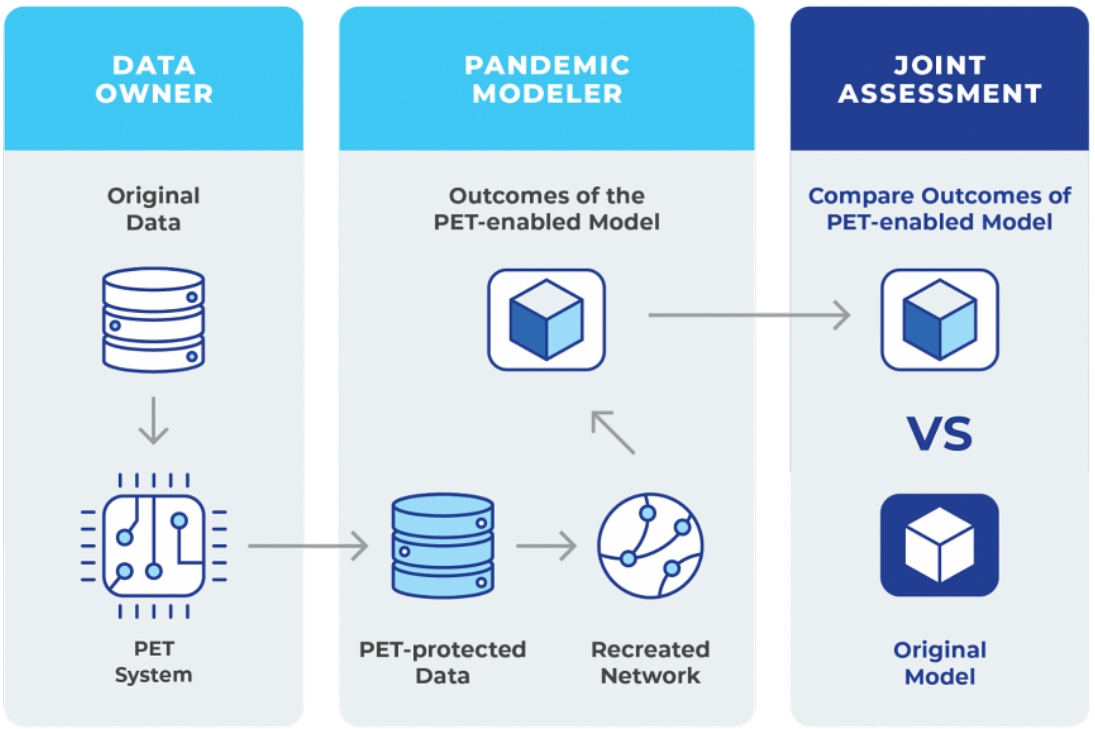

Figure 2: High-Level PET Experiment Design

Results

At an aggregate level, the temporal network created using PET-enabled contact information was near-identical to the one created from the original (sensitive) data. Across a population of 981 passengers and 19 crew, the PET-enabled solution identified:

- 278 instances of “casual contact”, compared to 281 with the original data;
- 211 instances of “close contact”, compared to 202 with the original data; and
- 511 instances of “transient contact”, compared to 515 with the original data.

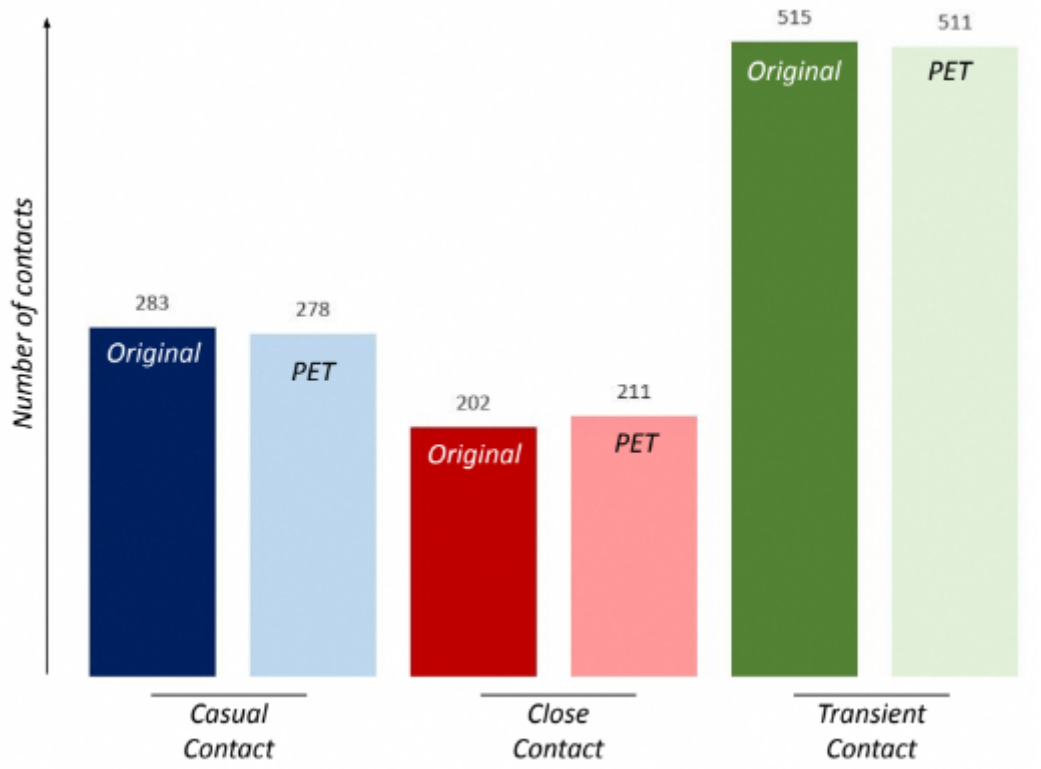

Figure 3: Comparison of Contact Instances between Original and PET-Enabled[4] Results

Lessons Learned

The IMDA/ DTC team highlighted four key lessons from the exercise:

1. **PET systems must be designed to balance the security and privacy concerns of data owners with the usefulness concerns of the modellers.**

   In this case:

   a. Data Owner Considerations: The key concern was for the original dataset to not be accessed or used by any external party. This was not just applicable to the obvious data fields such as those containing personal data, but also to the seemingly benign data field of ‘timestamp’ - which could also be sensitive given the context of the role/ responsibilities of the data owner.

[4] Using Homomorphic Encryption

b. Usefulness for Modellers - Selecting the "right kind" of PET: Multiple PETs can enable the flow of insights without disclosing original sensitive data. However, because of the nature of this particular use case, PETs that bring the code/query to the data environment and retrieve only the output of the data (e.g. Federated Learning, Multi-party computing) were also not considered to be suitable.

2. **Efforts to specify and minimise the data fields involved can improve the performance of the PET-enabled pandemic model.**

   For example:

   a. Choose only the most pertinent sensitive data field(s) (i.e. timestamp) to reduce storage requirements;
   b. Reduce the length of the data field representation to improve speed (here, we went from 7 characters to 4 characters); and
   c. Run the PET system on a fraction of the dataset to estimate computational overheads before procuring infrastructure for the entire dataset.

3. **Education and early involvement of stakeholders in charge of data ownership, governance and compliance helps design a more realistic PET system.**

   In this use case, a simple '*Who-Sees-What*' matrix, used to educate and agree with the data owner how the PET would flow original versus PET-protected data across the involved parties, was extremely valuable to get stakeholder buy-in.

Table 1: Who-Sees-What Matrix Example

Who sees what?	Data Owner	Data Analyst	Project Manager
Original Data	YES	NO	NO
PET-Protected Data	YES	YES	NO
Model Outcomes	YES	YES	YES

4. **There are clear trade-offs between HE and DP.**

   HE offered a lower relative error (i.e. closer to the original values), but required significantly more computation and storage infrastructure. It also entailed incremental governance overhead, including additional encryption/decryption key manager. This may limit scalability of a PET system with multiple data owners/users.

For those looking to use PETs in other similar pilots, the IMDA/ DTC team provides some additional practical advice:

i. Focus on use cases where a data barrier is clear and severely impedes socio-economic outcomes. This would warrant spending resources in a pilot.
ii. Scope pilots such that the data science element is kept simple. This would discipline the pilot to focus on unveiling PETs' technical and regulatory bounds.
iii. Like any other technical safeguard, PET-based solutions also need to be kept honest. Be prepared to update current governance measures or craft new ones.

Next Steps

The Singapore MOH will evaluate the feasibility of using HE for future use cases based on considerations such as software setup requirements, governance on storing and processing encrypted data, and readiness of external transaction data controllers. The DTC in turn will evaluate how to ease the HE software deployment by considering options such as packaging HE into libraries and clearly defined APIs, and developing an HE-infrastructure to support common services/computations at scale.

At a GPAI level, efforts will be made to disseminate the results from the demonstration project. If the community identifies additional pilot use cases for privacy-enhancing technologies that apply AI-for-social-good contexts, the IMDA, CEIMIA, and DTC are prepared to support: please contact info@ceimia.org for more information.