Federated Learning: don’t hammer with a saw

What is Federated Learning?

The traditional approach of sharing insights involves multiple parties contributing with their data to a central repository. Each party then independently analyses the data and extracts insights (figure 1). However, this approach is not suitable when the shared data is private or confidential.

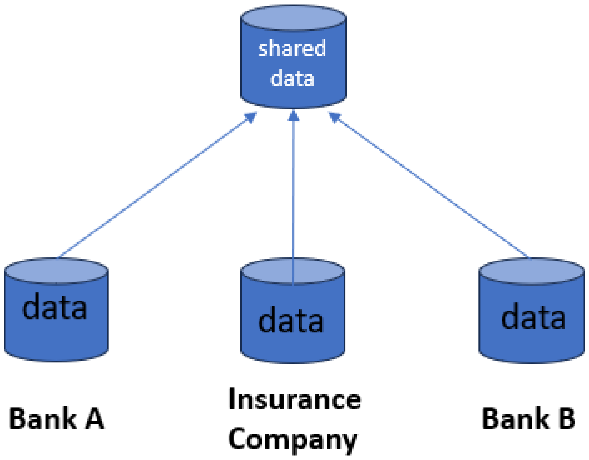

Figure 1. Insight Sharing without Federated Learning Across PCSI Partners

Privacy Enhancing Technologies (PETs), such as Federated Learning and Multi-Party Computation, help with data minimization and compliance with data-sharing regulations. Both are excellent candidates for analysing large amounts of sensitive data distributed over several partners. However, Multi-Party Computation works by sharing encrypted data between parties, while Federated Learning focuses on data aggregation to achieve a sufficient level of confidentiality.

Federated Learning is an efficient and effective technology for training models on distributed data, that protects the privacy and confidentiality of data. The goal of Federated Learning is not to share the data itself but to share insights whilst keeping confidential information secure, on-premises. With Federated Learning, partners will share insights by exchanging not the data itself, but updates to a locally trained machine learning model. Figure 2 depicts the scenario where Federated Learning is used. Instead of sending sensible data to a centralized server, each participant shares an update of the model locally trained on their organizations’ data. This way, the data never leaves the contributor’s premises, but the insights extracted by the local model are shared with the others. Later, the new centralized model is downstream to all partners' local models, which distributes the updated centralized model.

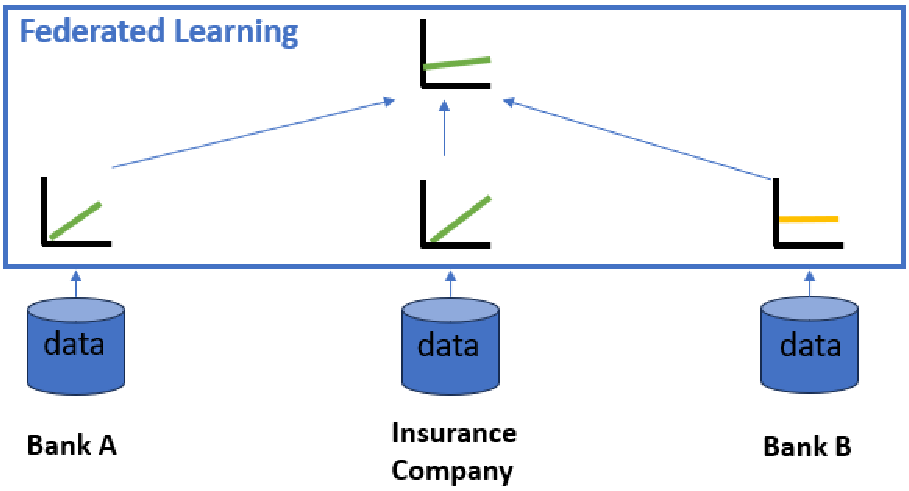

Figure 2. Insight Sharing with Federated Learning Across PCSI Partners

The PCSI use case

All PCSI partners must recognize security threats that endanger their core business. To proactively protect their business, PCSI partners are interested in sharing threat intelligence insights across specific markets (banking, governmental, insurance) to support each other in safeguarding their services against novel threats. However, due to privacy regulations, sensitive data can never leave their premises. Federated Learning could help achieve the goal of sharing threat intelligence insights without sharing actual data. The requirements below represent the needs of the PCSI partners to be satisfied for sharing threat intelligence insights:

- Share threat intelligence information anonymously, with relevant organizations in the same market domain (banking, governmental, insurance, medical, etc.) without sharing the actual data.
- Give each PCSI Partner flexibility in choosing their own data types to share (e.g. data anonymization).
- Timely updates and notifications of potential threats.
- Improve the as-is situation of threat intelligence sharing by removing/minimizing legal compliance challenges (internal/external).

Why Federated Learning?

The choice of Federated Learning was made based on the requirements mentioned above. The technology could help the PCSI partners achieve these goals.

As depicted in figure 2, Federated Learning can help share threat intelligence information without sharing the threat data with the other partners. This way, companies can analyse attack patterns and protect their business services with appropriate measures.

In the same way, but with a different granularity, Federated Learning can also help share threat intelligence insights between distinct departments of the same company. In this case, legal compliance hinders the insight exchange, and threat intelligence cannot be re-used across departments.

What we learned about Federated Learning in the PCSI project

A Federated Learning model requires all involved parties to adhere to a unified data format and to the same data compliance and security standards. This means that data needs to be saved and handled in the same way across all partners, such that the centralized model can collect and combine all gradients of individual models, without sharing the underlying data. Since the centralized model is sensitive to heterogenous data, which may introduce biases and skew the model, the unified data format is a strong requirement when implementing Federated Learning solutions.

Federated Learning is a great Privacy Enhancing Technology (PET) in scenarios where there are many input sources. Since the centralized model learns from the local models, and sensitive data can be easily deduced from individual model updates, the privacy of data sharing can only be guaranteed if the amount of input sources is sufficiently large.

Lastly, Federated Learning is a method, not a tool. Thus, an infrastructure needs to be built around this method to harness all its capabilities. A centralized platform is needed to collect the local models, train the centralized model, update all local models, and then share insights with contributors by deploying the updated models locally, to all partners. Performing model training steps takes a lot of time and resources, so Federated Learning does not work with time-sensitive data.

Why Federated Learning doesn't work for this use case

For the proposed use case, Federated Learning is not a good fit due to the reasons below.

First, Federated Learning guarantees privacy only when there are many contributors, as a rule of thumb. In this use case, only five parties are present. This would allow malicious attackers to gain knowledge about these companies through Inference Attacks; attacks that reverse-engineer the models' updates, giving insight into the data of a specific contributor.

Second, Federated Learning does not work well with time-sensitive data. As mentioned above, sharing insight between partners involves several steps, from anonymizing part of your data, to locally training your model, to sending an update to the centralized model, and downloading the up-to-date model only after all parties finalized the submission of their local models. Due to complexity of models in the above-mentioned use case, all these steps require days or sometimes weeks, which would not help against time-sensitive attacks such as zero-day exploits.

Finally, PCSI partners need the flexibility to choose data formats appropriate for their internal systems. Offering this flexibility, while ensuring all data formats are compatible with the centralized ML model, would introduce additional data conversion and parsing steps on the PCSI partner's slide, increasing the implementation costs of the Federated Learning solution.

The use case requires confidential data-sharing, fast updates, and flexibility in data types and platform choices. This clashes with the pre-requisites of Federated Learning and thus this technology is not a good fit for the presented use case.

Recommendations

Federated Learning is an excellent way to share insights with many entities while preserving privacy and confidentiality. Whilst it does not fit this PCSI project, it can be useful in other scenarios.

One use case where Federated Learning can assist is monitoring the usage of smartwatches across a large group of users, to obtain insights into users' behaviour. Smartwatches from the same manufacturer will most likely run the same software, and therefore, use the same data management techniques. Thus, it is simple to train local models on all these devices, then combine them into a central server to update the overarching machine-learning model. The high number of independent devices creates a substantial collection of data points for analysis, allowing the manufacturer to gain valuable insights into device usage while safeguarding user data privacy.

Another example is in medicine, where numerous X-rays depicting malformations are constantly reviewed. Analysing and extracting insights from these X-rays could lead to improved computer-aided patient diagnosis. However, since the images represent sensitive data tied to the patient, protected by data-collection regulations, their disclosure is forbidden. Federated Learning would be a great fit for this scenario: there are many parties involved (hospitals), the insights are not time-sensitive, and the sensitive data is never shared with other parties.

Furthermore, Federated Learning can be particularly helpful for big organizations that utilize the same email software (Outlook or Gmail). Usually, employees receive different types of spam emails with varying domains or employing distinct phrasing and keywords. By locally training a machine learning model on individual inboxes, insights regarding spam can be collected and shared with the centralized model without compromising the user's privacy. The spam filter is thus improved and can better aid employees in triaging their emails, minimizing the chance of falling victim to phishing attacks.

Thus, to profit from Federated Learning, the three main key aspects to be satisfied are:

1. **Many data inputs are needed;**
2. **Insights are not time-sensitive, and;**
3. **All data inputs use the same data format.**

If all these points are present, Federated Learning can help give great insight into the data whilst preserving the privacy of the users.

Helpful PET resources

If you would like to learn more about PET technologies and how these can help your organization's needs, TNO has these resources that can support you in the endeavour.

The *"FINALLY, A PRIVACY-FRIENDLY WAY TO HARNESS DATA"* Whitepaper discussing Federated Learning and MPCs, published by TNO.

The PET Explorer tool created by TNO for:

- Description & comparison of PETs;
- Discussion of legal considerations of PETs regarding GDPR and ethics;
- Decision tree and Checklists to help you determine the right PET for your organization's challenge

Authors

Alexandra Garban (TNO), Daan Opheikens (TNO), Leonardo Morelli (TNO), Rob Stübener (Achmea)

Acknowledgements

Special thanks to Bart Kamphorst (TNO), Peter Langenkamp (TNO), Reinder Wolthuis (TNO), Magali Kouwen (Achmea), and Niels Dekkers (Achmea) for their help with content discussions and feedback sessions.