ILO INSTITUTE

June 2019

Emerging Uses and Potential for Federated Learning in Financial Services

Key Findings:

Federated Learning is an approach to training artificial-intelligence systems from data that remains private. It allows central AI systems to learn from often-personal data without the content of the data being divulged or shared– only the lessons from the structure of the data are taken up.

Federated Learning promises

- to allow competitors to collaborate to improve AI models without divulging data to each other;
- to allow individual companies and other entities to improve their AI models – building out deeper dimensionality – based on individual, personal customer data that the customer does not want to share with the company;
- to support an on-device IoT approach to running models and tools on devices with little or no interaction with a central home system; and

- to significantly lower the energy and compute-power needed, and increase the clock speed of operations, for running distributed client software and models at large scale.

Launched by Google, Federated Learning was first discussed publicly in 2016. Google's widely influential AI tool, Tensor Flow, as of March 2019 now offers a Tensor Flow Federated extension. Google's Android device OS, installed on 2.5 billion devices, is the likely source of private data that early waves of Federated Learning agents will explore, though the model suits the exploration of a large number of relatively small pools of data intended to be kept private.

AML, fraud-detection and risk-scoring AI models will likely be greatly improved by Federated Learning tools, as learning agents explore phones and other devices to identify signature data structures of bad actors, and of exemplary credit-risk profiles. Federated Learning applications cannot – in theory – identify individuals fitting these categories, but they can teach home-station predictive models far deeper correlation identifiers for bad actors and bad actions based on still-private data.

Likely use-cases for Federated Learning in the banking sector include

- shared modeling among competitors based on "round trip" learning and predictive model enhancement;
- enhancing proprietary, not-shared predictive models;
- cost saving by pushing computing onto devices and collecting only structural insights for many operations;
- on-device lending; and
- pricing refinement.

There are no operating pilots of Federated Learning in the financial sector that we can identify today, though some are planned and interest across the sector is building.

WeBank in China – mostly owned by TenCent and WeChat – has built a platform for Federated Learning on the Platform-as-a-Service model, though it is largely untested. Swiss Re has just signed an MOU with WeBank to explore opportunities to use this model for **new-customer targeting and price optimization**.

Venture investing groups at Amex and Citi have begun actively looking for investment opportunities in Federated Learning among start-ups and other fast-moving pioneers.

Huntington Bank is one of the few regional players aware of Federated Learning, and hoping to learn more. Their head of analytics, Jason Black, has a special interest in how Federated Learning might reduce his team's engagement in building their own models, and relying instead on far better pooled-data models and platform-as-a-service offerings. Robert Ragno, head of the engineering team at Google continuing to build out Federated Learning also sees predictive modeling becoming much less important as Federated Learning matures.z

Trojan Horse?

Enthusiasm for Federated Learning runs into concern as users – rather than developers – begin to consider policy and regulatory issues. One executive puts it this way:

> *"We're skeptical about whether the privacy feature here is real. Once you're in my house looking at all my stuff, but not taking any of it home with you, I'm not sure that's OK with me.*

"And technologically I'm afraid that is a just a switch to go from forgetting the personal data to actually remembering it. So we have to ask, could this be a Trojan Horse to get more comfortable launching into the private space on all these devices?"

What Is Federated Learning?

Federated Learning is an approach to training artificial-intelligence systems from data that remains private. It allows central AI systems to learn from often-personal data without the content of the data being divulged or shared – only the lessons from the structure of the data are taken up.

In essence, Federated Learning is a set of tools for accomplishing four tasks.

Federated Learning promises to **allow competitors to collaborate to improve AI models without divulging data** to each other. The data informs the models, but only lessons from the data, not the data itself, is shared.

Federated Learning **promises to allow individual companies and other entities to improve their AI models** – building out deeper dimensionality based on individual, personal customer data that the customer often does not want to share with the company.

Federated Learning should support **an on-device IoT approach to running models** and tools on devices with little or no interaction with a central home system.

Federated Learning is proving out to **significantly lower the energy and compute-power needed, and increase the clock speed of operations,** for running distributed client software and models at large scale.

Federated Learning has been developed and launched principally at Google, with the 2.5 billion in-use Android devices in mind (Google's parent company owns the Android OS). The structure of Federated Learning is purpose-built to examine personal and other data on mobile devices, and take up lessons from that data but not take up the data itself. In theory, this means that a central AI system can learn many lessons from the data on the devices of billions of people, while people's and organizations' proprietary data remains entirely private.

The Federated Data Trojan Horse: One Reason for Resistance

One observer who spoke with ILO shared that "This is exciting. The research organization is very engaged, and we understand that this might be part of the promise of Google remaining a leader in privacy. Maybe so.

> *"But people like me who have some technical fluency but are operating executives by background, **we're skeptical about whether the privacy feature here is real**. Once you're in my house looking at all my stuff, but not taking any of it home with you, I'm not sure that's OK with me.*
>
> *"And technologically I'm afraid that is just a switch to go from forgetting the personal data to actually remembering it. So we have to ask, could this be a Trojan Horse to get more comfortable launching into the private space on all these devices?"*

Google Launches Federated Learning, 2016

Google began publishing technical papers on "Federated Learning" in 2016.

In 2017, Google announced that Federated Learning would be a formal research and product-development area, and in March of 2019 Google Brain announced the availability of TensorFlow Federated (TFF) as a standard tool-set for Federated Learning.

Google's research organization emphasizes Federated Learning as a model for learning from large numbers of private, connected phones, allowing third parties to train on data from vast networks of individual phones without violating data privacy of the users.

Gboard Proves it Works – with a focus on savings

The first large-scale application of Federated Learning is to run the computation of the Google keyboard Android app called Gboard, allowing autocompletion of words and sentences, and some analytics, entirely on-device with no "calling home" to a central system until and unless the central system calls out to the home base. The on-device Federated Learning tools would then return insights about the structure of the on-device data, without revealing the specifics of searches or other specific data.

The research and development team at Google working on Federated Learning, led by Principal Engineer Robert Ragno, is pleased to see Gboard proving the ability for Federated Learning to work at scale, but they see this as an almost trivial instance of what the approach can do.

The Gboard team is focused on the significant savings in energy use and compute resources now that local devices are doing the work of Gboard and only structural

insights are sent home to central systems – **but the potential for game-changing Federated Learning insights is barely hinted at though this pioneering use.**

Use Cases for Banking and Finance

The promise of Federated Learning is that it allows central AI systems to learn from data that it generally would not have access to. Because the Federated Learning movement begins with Google and moves outward from there, the generally unstated implication of the model is that Google can use Federated Learning to look into all the data on all Android phones and devices, and – while ignoring the specifics of the data – take away lessons from the structure of the data.

The Federated Learning Bullseye for AML

The most promising AML application Federated Learning is likely to be vastly better predictive models for identifying bad actors and rating transaction and counterparty risk.

Here's how it would work:

A Federated Learning agent looks at all of the roughly 2.5 billion Android devices active in the world today, at all or most of the data across each device, but without copying, removing, or event noting what that data says about real users.

The agent can use a pre-existing, standard model to determine that several hundred of the hundreds of millions of users of these devices are clearly prone to fraud, money laundering, or other bad acts.

The identification of these individuals is only based on logic applied to the data on the Android device. This first-level logic is not coming from the device or a federated Learning source; it's the logic already in-house from other, traditional sources of analytic insights. These might include content of emails, social media records, contact with known bad actors, records of engagement with legal authorities, records of engagement with bank AML teams, and other red flags.

Federated Learning tools will NOT record or report who these people are.

However, Federated Learning tools have the ability to look at the millions of bits of data on the Android devices of each of these identified extreme-risk individuals.

The Federated Learning agent will report back to a central AI system on the shape and structure of the bad-actor's device data and compare that with the structure of the data on the devices of hundreds of similarly identified bad actors.

Now the correlation of other data patterns can happen quickly, and the model for identifying extreme risks is vastly improved – and then applied to the already in-house data held by the institution or institutions seeking to enforce AML standards.

If, for example, the hundreds of bad actors are revealed to have certain travel patterns unseen by travel surveillance but present on personal travel data held on phones, that can be fed into the model. If patterns of kinds of books or music, or personal associations, or patterns of transaction sequences with merchants and online resources of all kinds,

correlate with extreme risk, those patterns can now be used as detection tools across the data already in-hand.

A more ambitious and potentially less benign approach would be to have the Federated Learning agent rating individuals for bad-actor traits on-device, and with a user's permission – with an application for funds or a transaction for example – the device owner can consent to having that score revealed.

Far Better Credit and Risk Rating

Applying the same dynamics as above but less at the extreme end of the spectrum, Federated Learning should be able to increase the accuracy of risk rating for individuals and for firms and other institutional counterparties.

The deep structure of the personal or institutional data on a device or devices used by the individual or institution can be matched to fine gradients of real risk, determined by actual banking history.

The model for each increment from terrible to wonderful can be drawn with enormously greater predictive accuracy from all of the data on an individual's device or an institution's devices, and credit and risk rating can be an on-going, on-device utility with little or no central computing or data collection needed, beyond polling the device for a score.

Pooling Data to Improve Shared Models for ML and Fraud

Several analysts expect that a major banking application of Federated Learning will be to allow competing firms to pool data about fraud, bad actors, and related risks, in order to establish and improve shared predictive risk models.

Federated Learning should allow firms to share lessons from data without sharing the data itself, allowing much greater input to these models, and therefore much higher quality.

Currently, this kind of collaboration is limited by the range of data that a bank might have through explicit permissions and transaction history, perhaps overlayed with externally sourced enhancements.

A mature Federated Learning system should be able to examine all or most data on the mobile devices of a vast number of users – increasing the pool of individuals whose data is input for training, and (by an enormous degree) the depth of the data for each individual input for training.

Intel Focusing on Federated Learning for Health

Intel has stood up a research group for applying Federated Learning to biomedical applications, especially diagnostic image reading. They've announced plans to develop a new chip specifically focused on supporting Federated Learning.

The University of Pennsylvania is a key partner working with Intel to build a secure Federated Data platform for medical imaging. No similar formal effort is underway at Intel focused on financial data.

MIT Media Lab – Fighting the Trojan Horse

MIT's Media Lab is trying to address the Trojan Horse privacy potential of Federated Learning through the development of "Split Learning," best understood as a deeper variation of Federated Learning with more coherence on the privacy side, and a more robust technical privacy layer that can be set to make peering-in to PII technically impossible, or close to it.

Interestingly, MIT's approach is coming specifically from their "Camera Culture Group" in the Media Lab, suggesting that visual analytics will be a strong point of focus for Split Learning.