Privacy-Preserving Collaborative Anomaly Detection to Fight Cybercrime

by Rudolf Mayer (SBA Research)

Anomaly detection is an important part of countering cybercrime, by detecting e.g., fraud or intrusions. Especially with an ever-growing amount of data (such as logs or transactions) being collected, automated analysis of these data for malicious behaviour becomes essential. In several settings, such analysis might be performed by third parties or be collaborative, to learn from more and diverse experiences by different collaborators. Thus, means to access such often confidential data in a privacy-preserving manner are required. Collaborative Learning and synthetic data are two promising approaches to fulfil this purpose.

The demand for and practice of data sharing and exchange between different data collecting parties is increasing, often because different data sets complement each other, or because the processing and analysis of data is outsourced. Many interesting knowledge discovery tasks are dependent on large, high-quality amounts of data, and anomaly detection, e.g., for the purpose of detecting cybercrime, is no exception – especially as fraudulent behaviour changes over time.

However, when data is sensitive, e.g., when it concerns individuals or is business related, there are certain regulatory and other barriers for data sharing and collaboration. Still, collaborative analysis of data can be very beneficial, as different organisations might be affected (targeted) at different stages. Thus, learning from misuse patterns that other parties have already been exposed to, such as network intrusions, financial fraud, malware, or other forms of cybercrime, can be extremely valu-

able. Therefore, means to enable such data exchange collaboration are required. Anonymisation of data is one frequently studied approach, e.g., in the form of k-anonymity or differential privacy. However, k-anonymity has been shown to be still prone to linkage attacks when adversaries have background knowledge and access to other data sources, and differential privacy is not easily applicable to all types of analysis techniques and methods. Therefore, alternative approaches such as synthetic data generation and federated learning have also been explored, and recently specifically evaluated for anomaly detection.

Synthetic data is generally considered as data obtained not from direct measurement. In the context of data analysis efforts, it is often considered to be data generated (or synthesised) from a real dataset that cannot be shared e.g., for privacy considerations. The aim of synthetic data in this context is to generate a dataset containing records that are similar to the original ones, and thus allow a similar analysis as the original data, without actually disclosing real, single data points – which might be data points that contain sensitive information that is not to be disclosed. Synthetic data has successfully been shown to achieve good results, e.g., for classification and regression tasks. However, one could assume that synthetic data might not be easily usable for a task such as anomaly detection, which deals with outliers, while synthetic data generally preserves global characteristics. It thus might tend to represent rather the ones of the legitimate, normal cases, and might fail to generate representative anomalies.

However, in a recent work [1], we evaluated three different approaches to generate synthetic data for supervised, semi-supervised and unsupervised anomaly detection settings, including credit card fraud. While anomaly detection is a hard task especially for the latter two, synthetic data reaches similar effectiveness as the models trained on the original data, and can thus be considered a viable alternative: instead of sharing and centralising real data, different collaborators can exchange synthetic data generated locally.
Another approach for privacy-preserving data analysis is federated learning [2], which is especially useful

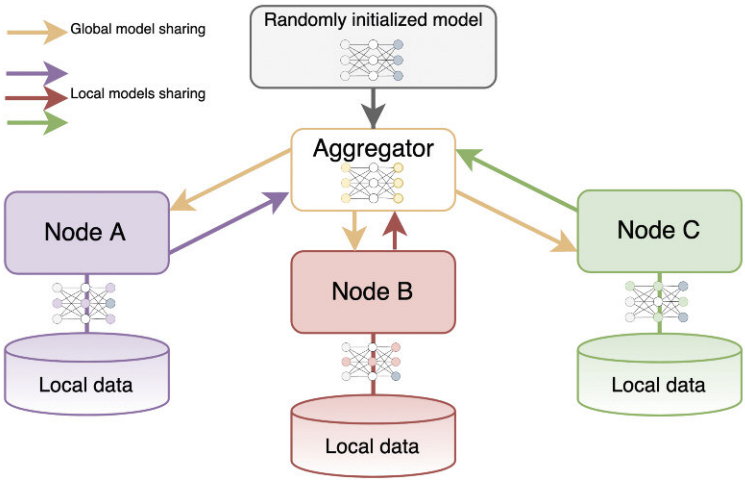

Federated learning with three nodes learning models on their local data, and model averaging by a central aggregator.

in settings where data is collected in several distributed locations. In federated learning, first models are trained locally on each data source, before they are aggregated to a common, global model; this is usually repeated for a few cycles, to allow convergence. Thus, the training data remains at the source, and the only type of information exchanged are the model parameters – which generally represent a strong abstraction of the local model. Aggregation strategies vary for each anomaly detection method. For example, for the supervised task with neural networks, e.g., simple averaging of the learned model weights (or from the gradients to adapt those) is a suitable strategy, while for other methods, different aggregation strategies are required, or ensembles could be built. A recent evaluation shows that especially supervised, but also semi-supervised methods can achieve comparable detection rates [3]. In general, federated learning is heavily influenced by the way data is distributed among the clients, and model aggregation needs to consider if there is an imbalance and skewness in the amount of data and anomalies present at each local site, and more advanced aggregation strategies might need to be developed for these cases.

Protecting the confidentiality of training data is an important aspect in many settings, but recent approaches have shown promising results in several anomaly detection settings, including in cybercrime. While it might not be always possible to replicate results that would be achievable in an idealised approach of centralising all data, collaboration can improve the results clients would achieve if just leveraging their own data.

References:

[1] R. Mayer, M. Hittmeir, A. Ekelhart: "Privacy-preserving anomaly detection using synthetic data", in Proc of the 34th Annual IFIP WG 11.3 Conf. on Data and Applications Security and Privacy (DBSec), Regensburg, Germany, 2020. Springer.

[2] P. Kairouz, H. Brendan McMahan, et al.: "Advances and Open Problems in Federated Learning", Foundations and Trends in Machine Learning, 14(1-2), 2021.

[3] F. Cavallin, R. Mayer: "Anomaly Detection from Distributed Data Sources via Federated Learning". Proceedings of the 36th International Conference on Advanced Information Networking and Applications (AINA), Sydney, Australia, 2022. Springer International Publishing.

Please contact:

Rudolf Mayer
SBA Research, Austria
rmayer@sba-research.org