April 2023

Secure and Private Access to Machine Learning Models via Cloud-based Execution Platform

Jérôme Glisse

Secure and Private Access to Machine Learning Models via Cloud-based Execution Platform

ABSTRACT

Machine learning (ML) models for various purposes may be made available via cloud-based execution platforms. The data owner, the model provider, and the execution platform where a model is deployed may be different entities. In this configuration, it is necessary to establish trust between the model provider, the execution platform, and the user to enable secure and private use of machine learning models. This disclosure describes an execution platform that can be verified and authenticated by a data owner as a model provider using techniques such as confidential computing and/or trusted execution environment. The execution platform enables the model provider to make their models available for use by data owners. Data owners can provide encrypted data to the execution platform and specify a ML model to be applied to the data. The execution platform implements an instance of the ML model and enables secure access to the user data without the model provider being able to access the data. Thus, users can obtain the benefit of using a ML model of their choice via the execution platform and model providers can provide models in a secure marketplace, while preserving the privacy and security for both entities.

KEYWORDS

- Confidential computing
- Trusted execution environment
- Secure enclave
- Machine learning marketplace
- Data privacy
- Model confidentiality

BACKGROUND

Machine learning models can be used for a variety of tasks such as email spam filtering, image classification, etc. To perform tasks such as spam filtering or image classification, the emails or images are provided to the model. Models for various purposes may be made available by different entities. Many machine learning tasks are performed on cloud-based execution platforms. In current applications of machine learning models, the data owner (user), the model provider, and the execution platform where a model is deployed (e.g., a cloud computing provider) may thus be different entities. In this configuration, it is necessary to establish trust between the model provider, the execution platform, and the user to enable the user to utilize machine learning models.

DESCRIPTION

This disclosure describes techniques to allow an entity (e.g., a model provider) to provide another, separate entity (e.g., an execution platform) with access to a machine learning model, without the model provider getting access to user data to be processed by the model, and without the data owner getting access to the model.

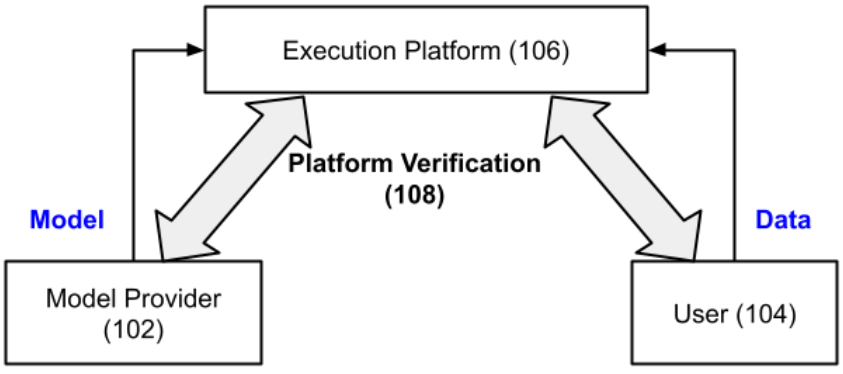

Fig. 1: Secure and Private Access to Machine Learning Models via Execution Platform

Fig. 1 illustrates an example method to automatically provide access to machine learning models via an execution platform (e.g., hosted by a cloud computing provider) to end users while preserving user privacy. A model provider (102) makes one or more machine learning models available for users. The model(s) have capabilities that can be used by a user (104) that wants to use the capabilities to process their data, e.g., filter spam from incoming email, assign labels to images, etc. In this context, either or both - the model provider and the user - may wish to keep their information private from the other entity. The model provider wants to make the model available to process user data without the user getting access to the model itself. Similarly, the user wants to leverage the capabilities of various models without providing their data to the model provider.

In the example of Fig. 1, a secure execution platform (106) enables this. The secure execution platform hosts the models and provides the computing infrastructure for the model to receive the user data as input and provide results as output to the user.

The model provider and the user each authenticate and verify (108) the execution platform. Authentication and verification can be done using industry standard mechanisms such as confidential computing (CC) and/or a trusted execution environment (TEE). The execution platform can provide a secure execution enclave in which only the end user is able to access their data.

Upon successful verification of the execution platform, the model of interest to the user is requested from the model provider, for execution in the secure execution enclave. The user data is also made available in the secure execution enclave. The model is executed on the user data to provide results (e.g., whether an email in the user data is likely spam, labels for images in the user data, or any other inference task). The results are accessible only to the user and not

to the model provider. The model is usable by the user via the execution platform but is not revealed to the user.

The execution platform implements technical mechanisms that enforce isolation between the user and the model provider. Such mechanisms can include a combination of hardware mechanisms such as a memory protection unit, memory encryption, etc. and software, e.g., scheduler, firmware that enforces security policy, etc.

The described techniques can be employed by a cloud computing provider or any other hosting provider to enable secure access to machine learning models, while preserving the confidentiality of the model as well as that of the user data processed by the model. In situations where the model provider and the cloud computing provider are the same, the verification by the model provider can be omitted. The techniques enable a model provider to make the functionality of their machine learning models/ algorithms available to customers while ensuring model security and data privacy. A cloud computing provider can utilize the techniques to provide a secure marketplace for machine learning algorithms/models.

Further, vendors that provide machine learning accelerators (e.g., hardware units optimized for machine learning) can provide access to their technology using the techniques described herein. Software-as-a-Service (SaaS) providers can leverage the described techniques to utilize machine learning capabilities that are available in their customer's infrastructure (e.g., private servers) .

Examples of use

- A model provider provides a spam detection model and makes it available on an execution platform provided by a cloud computing provider. A customer (e.g., a

corporate entity) provides a stream of their email to an instance of the model running on the execution platform to perform spam detection.

- A model provider provides a model to classify images. The model provider is the same entity as the cloud provider that hosts the execution platform. A customer uses the execution platform to obtain classifications for their images. The customer encrypts their images prior to submitting the images for classification. The images are classified by the model, after being decrypted using decryption keys that are available only via the trusted execution platform that has been verified by the customer. In this case, even though the model provider and the execution platform are provided by the same entity, the user has a guarantee that the entity does not have access to their data.

While the foregoing description relates to secure and private access to machine learning models via cloud-based execution platform, the described techniques can be utilized for any application where security and privacy of the user and of a technology provider are important. For example, the techniques can also be used to provide secure and private access to custom hardware/software (e.g., proprietary to the cloud platform) for video transcoding or video conferencing.

CONCLUSION

This disclosure describes an execution platform that can be verified and authenticated by a data owner as a model provider using techniques such as confidential computing and/or trusted execution environment. The execution platform enables the model provider to make their models available for use by data owners. Data owners can provide encrypted data to the execution platform and specify a ML model to be applied to the data. The execution platform implements an instance of the ML model and enables secure access to the user data without the

model provider being able to access the data. Thus, users can obtain the benefit of using a ML model of their choice via the execution platform and model providers can provide models in a secure marketplace, while preserving the privacy and security for both entities.