



US012045693B2

(12) **United States Patent**
Swan et al.

(10) **Patent No.:** **US 12,045,693 B2**
(45) **Date of Patent:** **Jul. 23, 2024**

(54) **PACKAGING AND DEPLOYING ALGORITHMS FOR FLEXIBLE MACHINE LEARNING**

(71) Applicant: **Amazon Technologies, Inc.**, Seattle, WA (US)

(72) Inventors: **Charles Drummond Swan**, Seattle, WA (US); **Edo Liberty**, New York, NY (US); **Steven Andrew Loeppky**, Seattle, WA (US); **Stefano Stefani**, Issaquah, WA (US); **Alexander Johannes Smola**, Sunnyvale, CA (US); **Swaminathan Sivasubramanian**, Sammamish, WA (US); **Craig Wiley**, Redmond, WA (US); **Richard Shawn Bice**, Sammamish, WA (US); **Thomas Albert Faulhaber, Jr.**, Seattle, WA (US); **Taylor Goodhart**, Issaquah, WA (US)

(73) Assignee: **Amazon Technologies, Inc.**, Seattle, WA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 779 days.

(21) Appl. No.: **16/001,548**

(22) Filed: **Jun. 6, 2018**

(65) **Prior Publication Data**

US 2020/0311617 A1 Oct. 1, 2020

Related U.S. Application Data

(60) Provisional application No. 62/590,242, filed on Nov. 22, 2017.

(51) **Int. Cl.**
G06N 20/00 (2019.01)
G06F 9/455 (2018.01)

(52) **U.S. Cl.**
CPC **G06N 20/00** (2019.01); **G06F 9/45558** (2013.01); **G06F 2009/45595** (2013.01)

(58) **Field of Classification Search**
CPC G06F 9/455-45558; G06F 2009/45562-45595
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

8,521,664 B1 * 8/2013 Lin G06N 5/04 706/12
9,378,044 B1 6/2016 Gaurav et al. (Continued)

FOREIGN PATENT DOCUMENTS

CN 105378667 A 3/2016
EP 3182280 A1 6/2017
WO 2015/085475 A1 6/2015

OTHER PUBLICATIONS

Yash. Single versus Multiple Containers. May 17, 2016. StackOverflow. <https://stackoverflow.com/questions/37268790/single-versus-multiple-containers> (Year: 2016).*

(Continued)

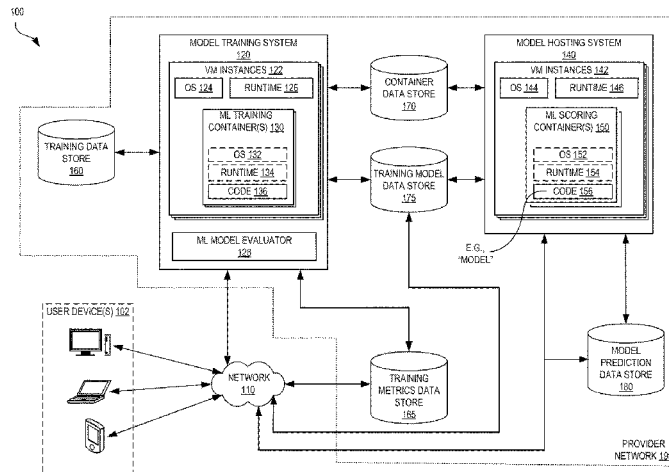
Primary Examiner — Liang Y Li

(74) *Attorney, Agent, or Firm* — NICHOLSON DE VOS WEBSTER & ELLIOTT LLP

(57) **ABSTRACT**

Techniques for using scoring algorithms utilizing containers for flexible machine learning inference are described. In some embodiments, a request to host a machine learning (ML) model within a service provider network on behalf of a user is received, the request identifying an endpoint to perform scoring using the ML model. An endpoint is initialized as a container running on a virtual machine based on a container image and used to score data and return a result of said scoring to a user device.

20 Claims, 13 Drawing Sheets



(56)

References Cited

U.S. PATENT DOCUMENTS

9,454,733	B1	9/2016	Purpura et al.	
11,080,435	B2 *	8/2021	Bourhani	G06F 30/00
2008/0288622	A1	11/2008	Gordon et al.	
2013/0144819	A1 *	6/2013	Lin	G06N 5/02 706/12
2014/0137111	A1 *	5/2014	Dees, Jr.	G06F 9/45558 718/1
2014/0344194	A1	11/2014	Lee et al.	
2015/0379072	A1 *	12/2015	Dirac	G06N 20/00 707/693
2016/0148115	A1	5/2016	Sirosh et al.	
2016/0294870	A1	10/2016	Banerjee et al.	
2016/0359622	A1	12/2016	Bunch	
2017/0017903	A1	1/2017	Gray et al.	
2017/0052807	A1	2/2017	Kristiansson et al.	
2017/0177413	A1	6/2017	Wisniewski et al.	
2017/0178027	A1 *	6/2017	Duggan	G06F 9/543
2017/0213156	A1 *	7/2017	Hammond	G06F 18/2148
2017/0220949	A1	8/2017	Feng	
2017/0223190	A1	8/2017	Mandel et al.	
2017/0249141	A1	8/2017	Parees et al.	
2018/0089592	A1	3/2018	Zeiler	
2018/0095778	A1 *	4/2018	Aydelott	G06Q 10/0637
2018/0300653	A1 *	10/2018	Srinivasan	H04L 67/06
2019/0081955	A1 *	3/2019	Chugtu	H04L 63/0236
2019/0114370	A1	4/2019	Cerino	

OTHER PUBLICATIONS

Stringfellow, Angela. Soap vs. Rest: Differences in Performance, APIs, and More. Dzone. <https://dzone.com/articles/differences-in-performance-apis-amp-more> (Year: 2017).*

Bezgachev, Vitaly. How to deploy Machine Learning models with TensorFlow Parts 1-4. Nov. 12, 2017. TowardsDataScience.com. <https://towardsdatascience.com/how-to-deploy-machine-learning-models-with-tensorflow-part-1-make-your-model-ready-for-serving-776a14ec3198> (Year: 2017).*

Cozannet, Samuel. GPUs & Kubernetes for Deep Learning Parts 1-3. Feb. 15, 2017. Medium.com. <https://medium.com/hackernoon/gpus-kubernetes-for-deep-learning-part-1-3-d8eebe0dd6fe> (Year: 2017).*

Blaise, "How to add a custom domain for a serverless-1.0.0 framework defined/deployed API?", StackOverflow, Sep. 20, 2016. <https://stackoverflow.com/questions/39507004/how-to-add-a-custom-domain-for-a-serverless-1-0-0-framework-defined-deployed-api> (Year: 2016).*

Casalboni, Alex. How we use AWS for Machine Learning and Data Collection. Cloud Academy. May 5, 2016. <https://cloudacademy.com/webinars/cloud-academy-and-aws-how-we-use-aws-machine-learning-and-data-collection-10/> <https://www.youtube.com/watch?v=fdlDn3hr27k&t=2193s> (Year: 2016).*

Avraham, Shif Ben. What is Rest—A Simple Explanation for Beginners. Medium.com. <https://medium.com/extend/what-is-rest-a-simple-explanation-for-beginners-part-1-introduction-b4a072f8740f> (Year: 2017).*

Communication pursuant to Article 94(3) EPC, EP App. No. 18815886.9, Sep. 15, 2020, 11 pages.

International Preliminary Report on Patentability, PCT App. No. PCT/US2018/061869, Jun. 4, 2020, 13 pages.

International Preliminary Report on Patentability, PCT App. No. PCT/US2018/062223, Jun. 4, 2020, 8 pages.

Notice of Allowance, U.S. Appl. No. 15/901,751, Jul. 2, 2020, 9 pages.

International Search Report and Written Opinion issued in related International Patent Application No. PCT/US2018/062223, mailed Feb. 25, 2019, 11 pages.

International Search Report and Written Opinion for related International Patent Application No. PCT/US2018/061869, mailed Mar. 11, 2019, 18 pages.

Final Office Action for related U.S. Appl. No. 15/901,751, mailed Feb. 14, 2020, 27 pages.

Non-Final Office Action for related U.S. Appl. No. 15/901,751, mailed Aug. 12, 2019, 30 pages.

Benedicic et al., "Shifter: Fast and consistent HPC workflows using containers", Containers in HPC Workshop, Jun. 29, 2017, 11 pages.

Summons to Attend Oral Proceedings Pursuant to Rule 115(1) EPC, EP App. No. 18815886.9, Mar. 23, 2021, 13 pages.

Anonymous, "Understand Images, Containers, and Storage Drivers", Docker Inc., XP055444825, Oct. 23, 2016, 22 pages.

Non-Final Office Action, U.S. Appl. No. 17/067,285, Mar. 16, 2022, 31 pages.

Preliminary Opinion, EP App. No. 18815886.9, Nov. 9, 2021, 16 pages.

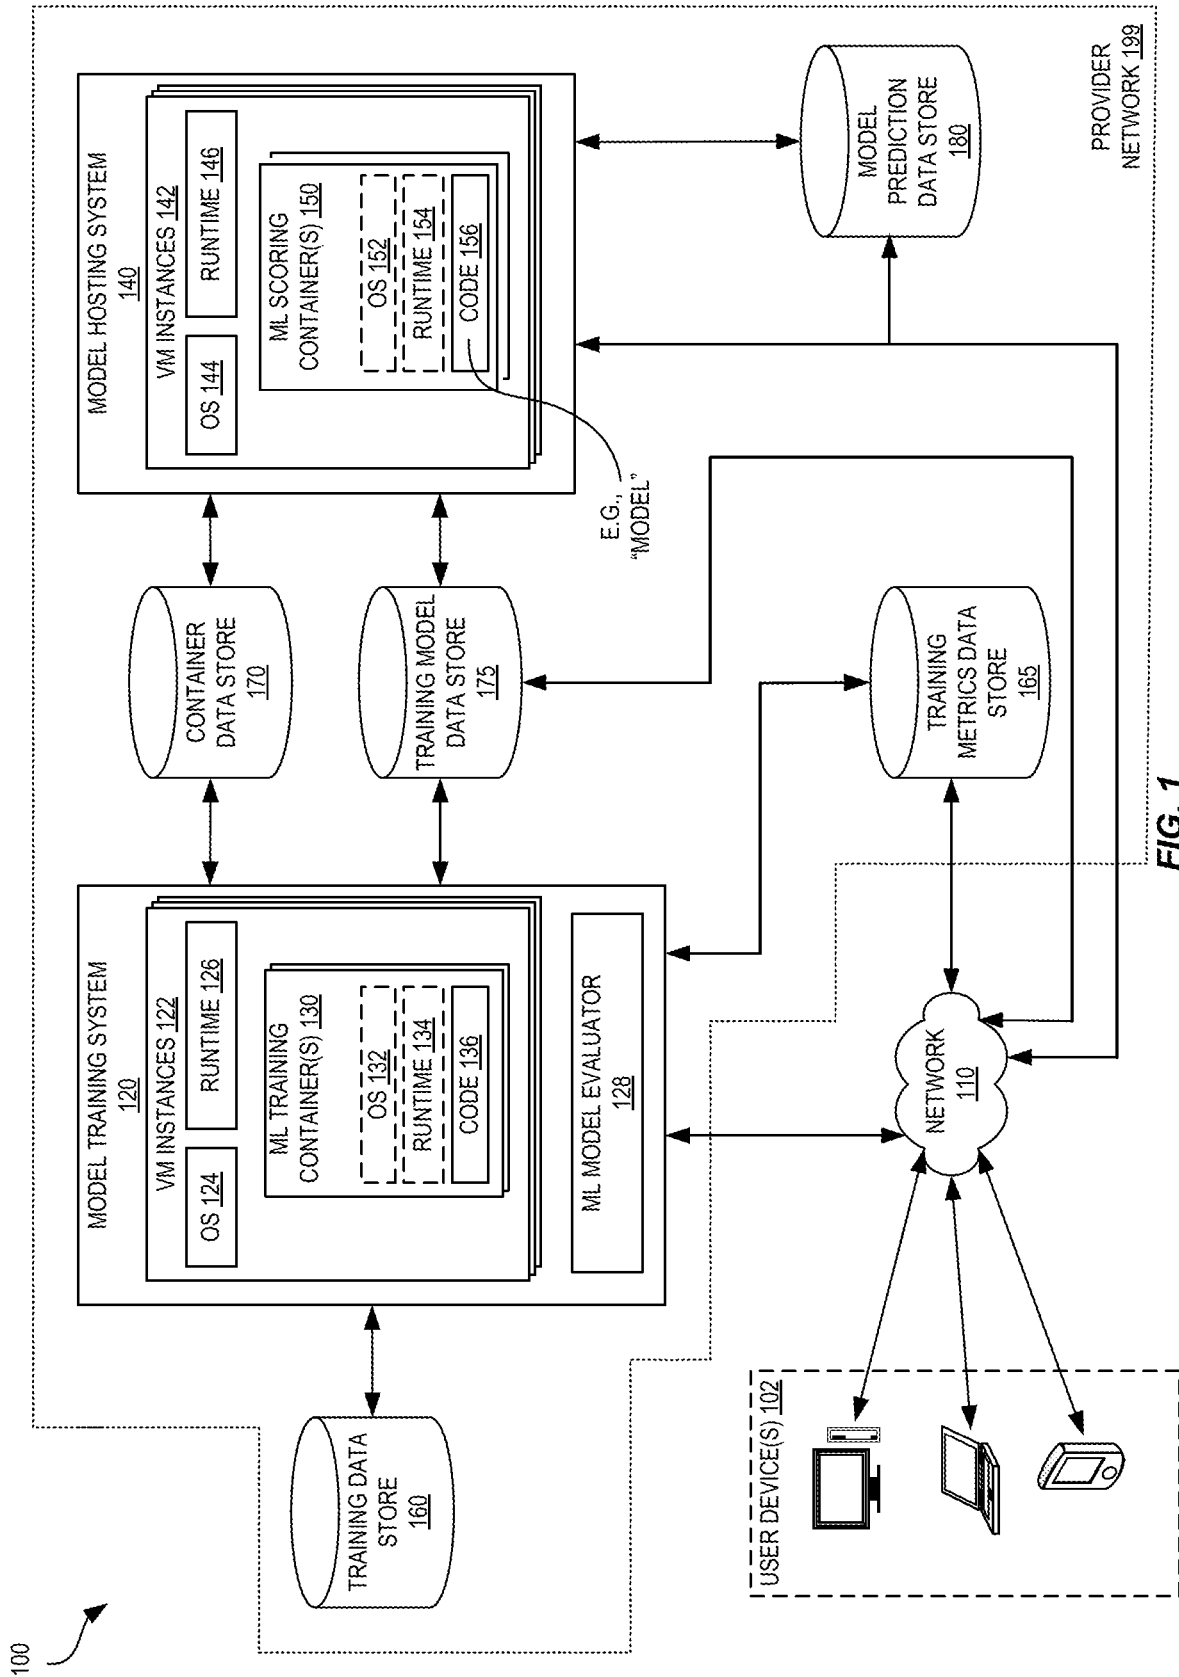
Notice of Allowance, U.S. Appl. No. 17/067,285, Sep. 14, 2022, 9 pages.

Notice on the First Office Action, CN App. No. 201880075354.9, Mar. 31, 2023, 18 pages (8 pages of English Translation and 10 pages of Original Document).

Zhao, et al., "Performance of Container Networking Technologies", HotConNet '17, Proceedings of the Workshop on Hot Topics in Container Networking and Networked Systems, Aug. 25, 2017, pp. 1-6.

Office Action, CN App. No. 201880075354.9, Jan. 19, 2024, 13 pages (6 pages of English Translation and 7 pages of Original Document).

* cited by examiner



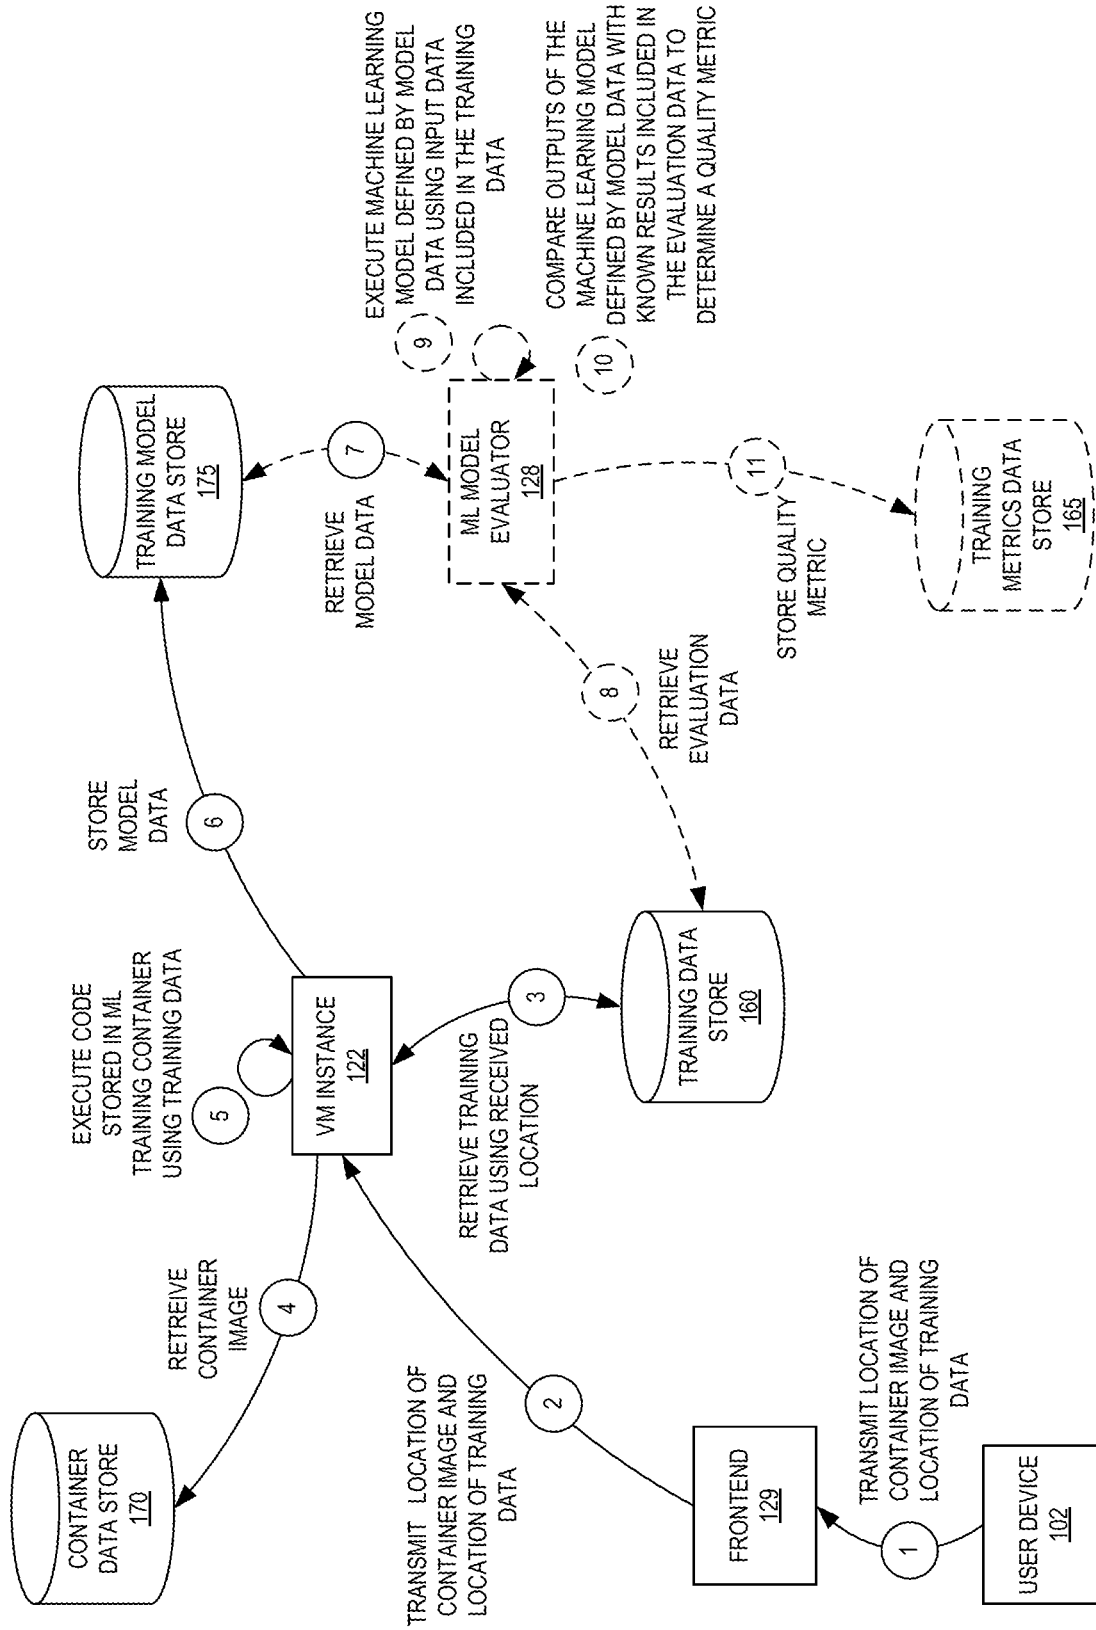


FIG. 2

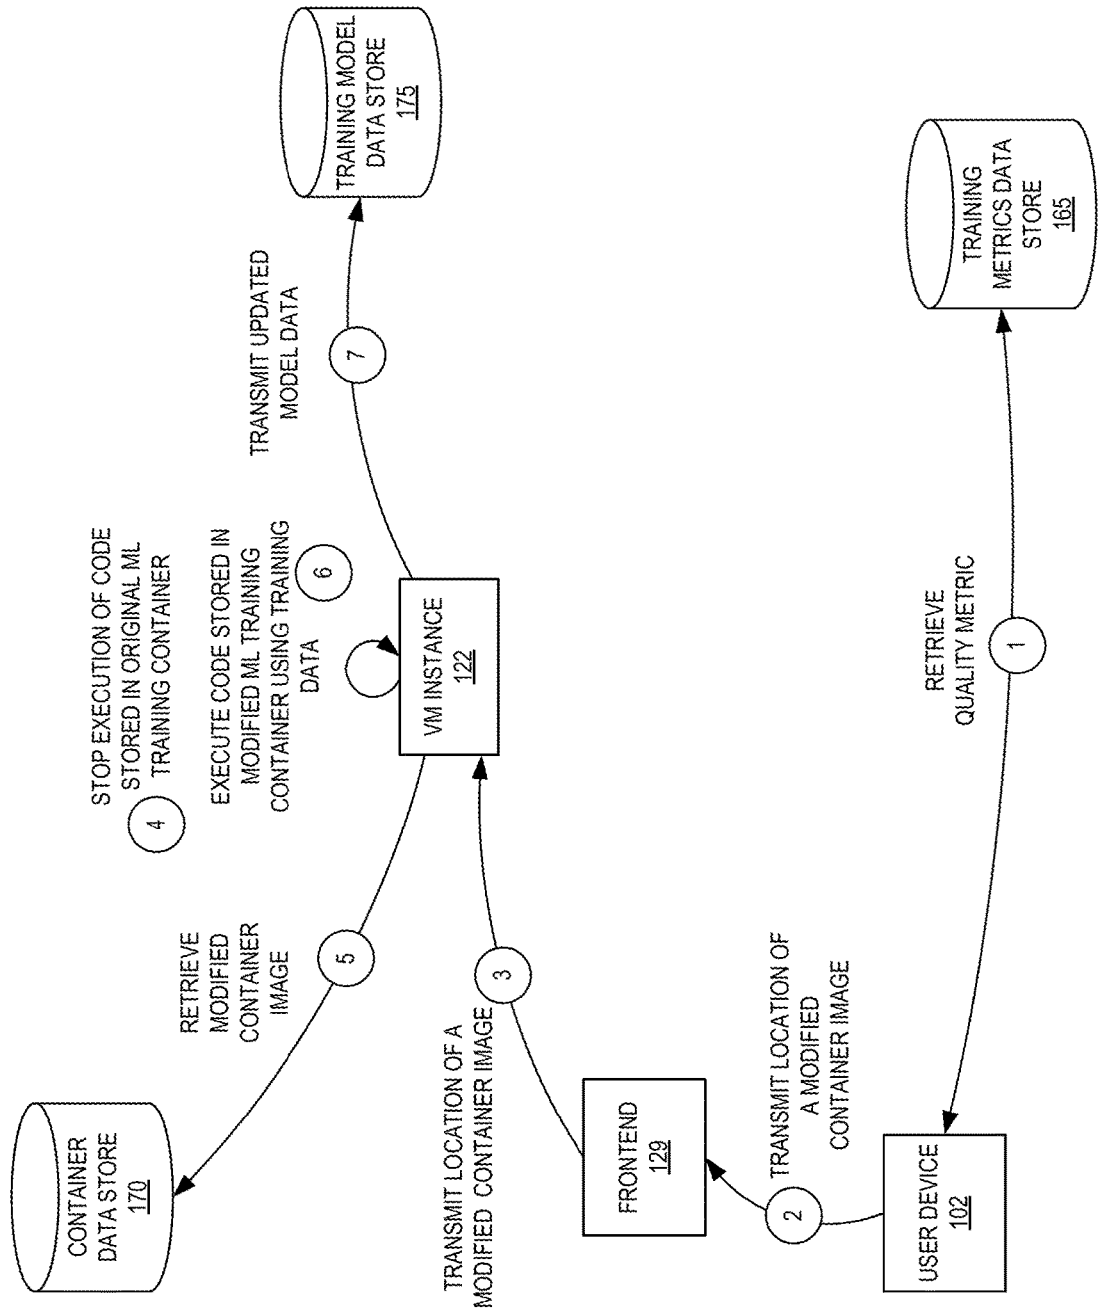


FIG. 3

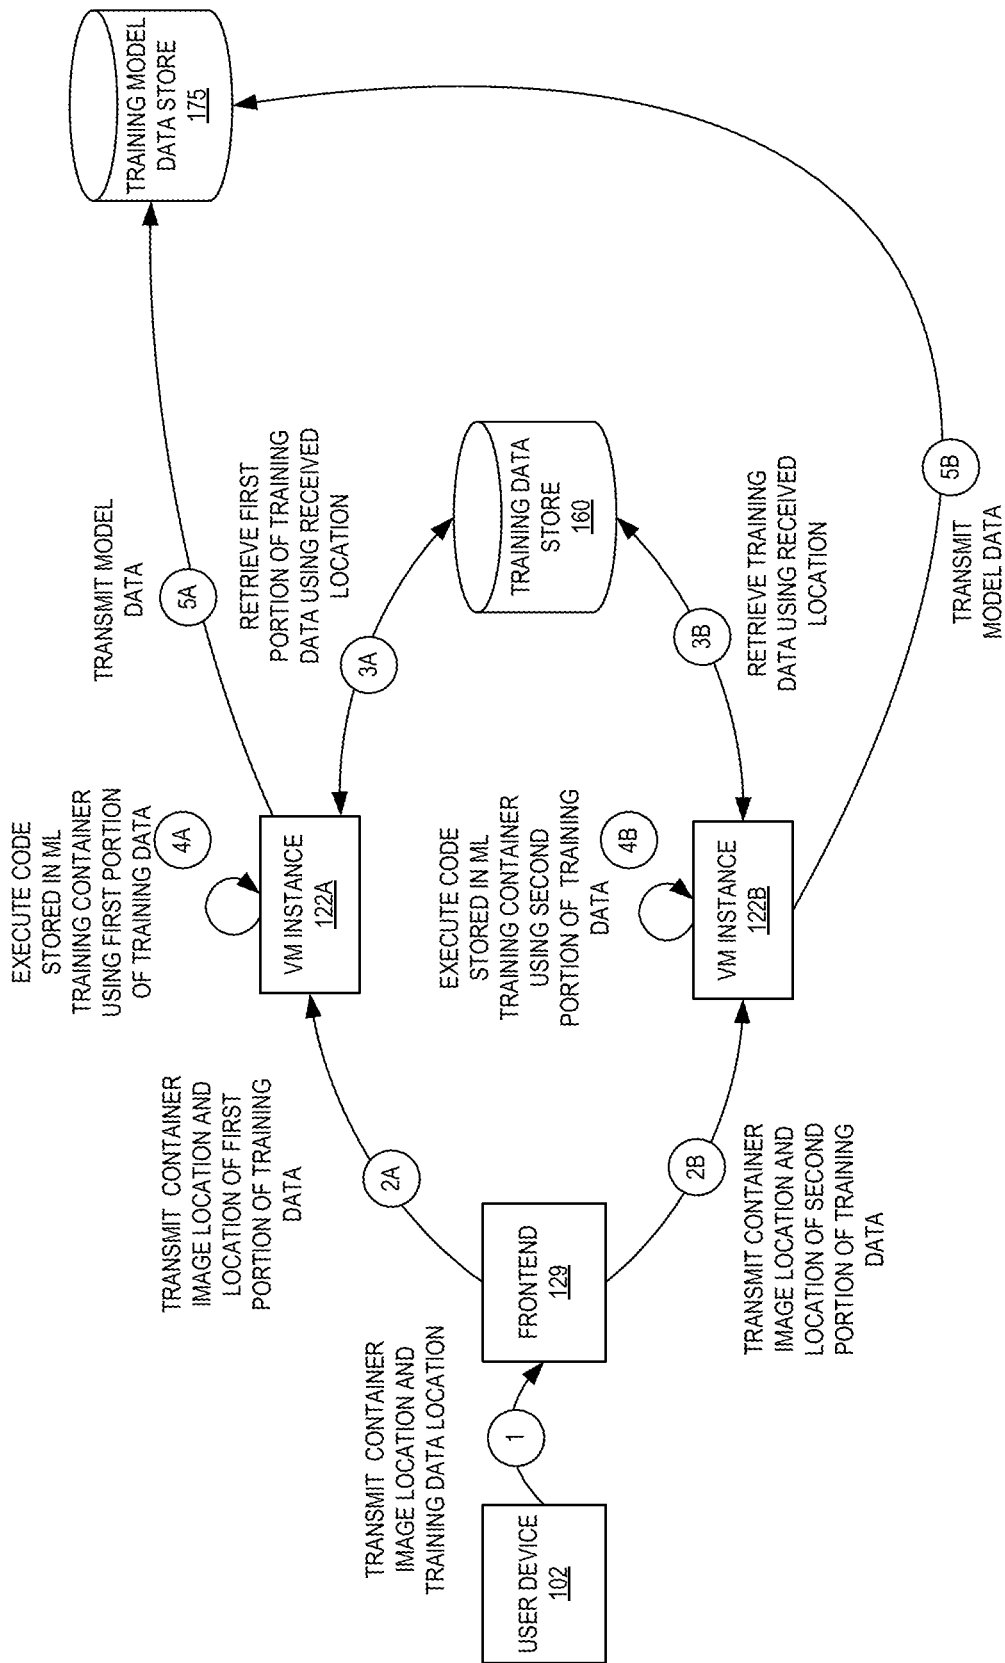


FIG. 4

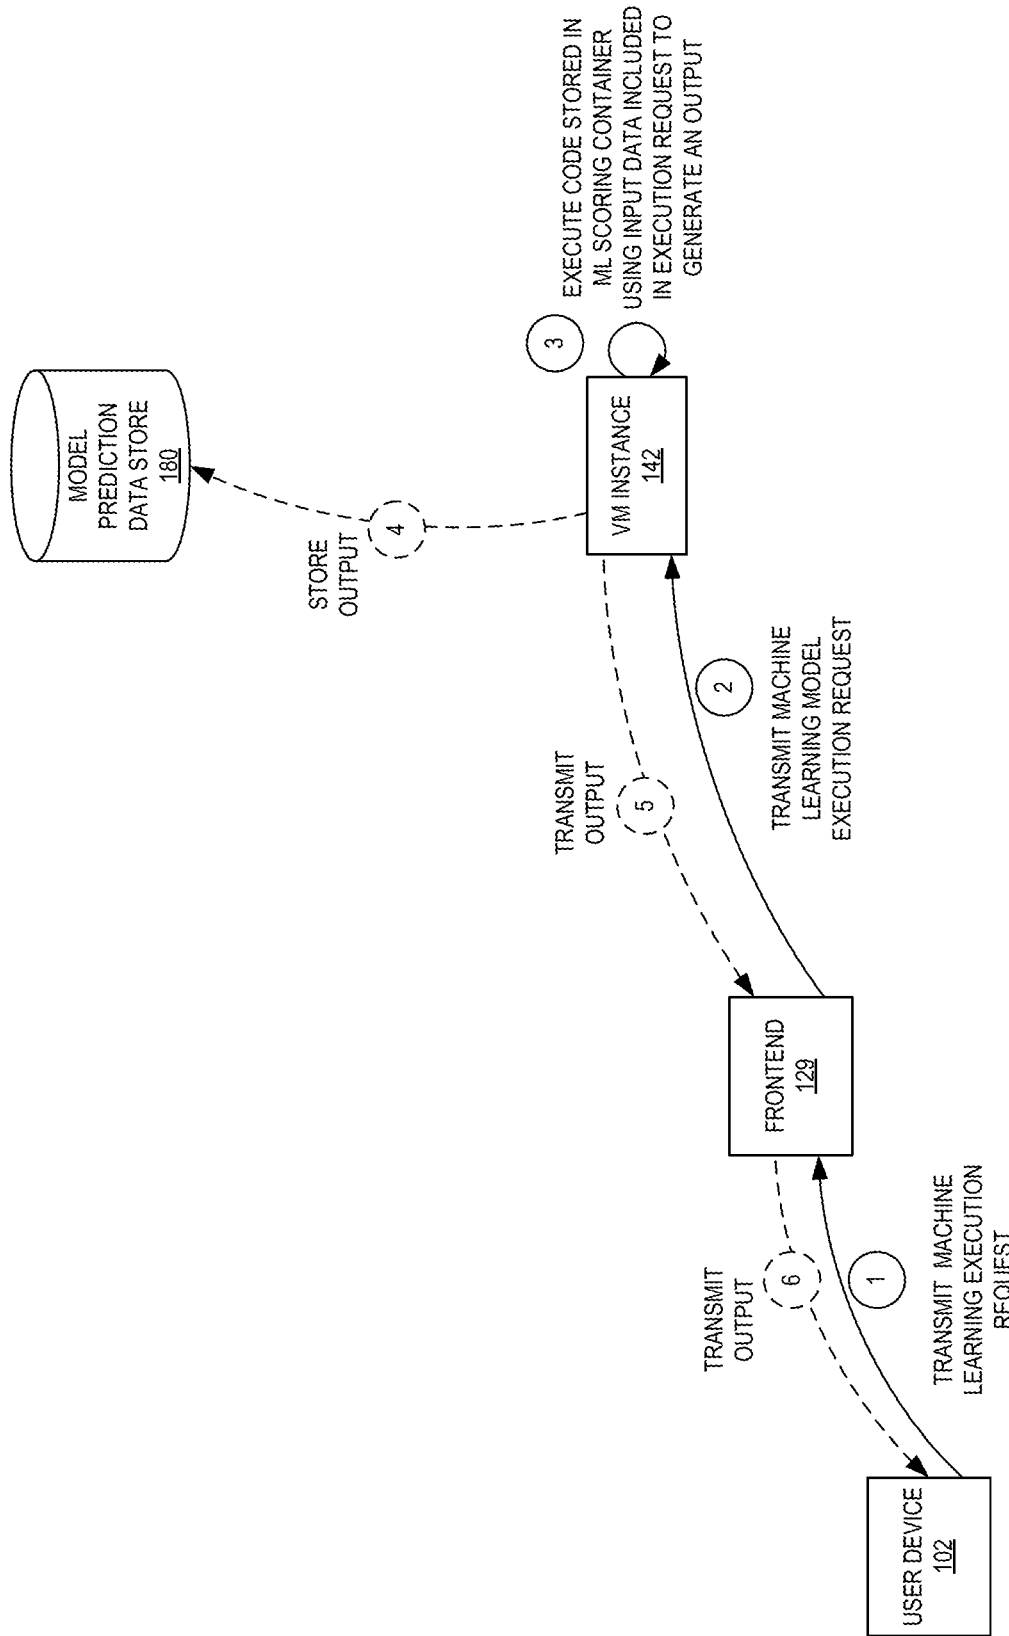


FIG. 5B

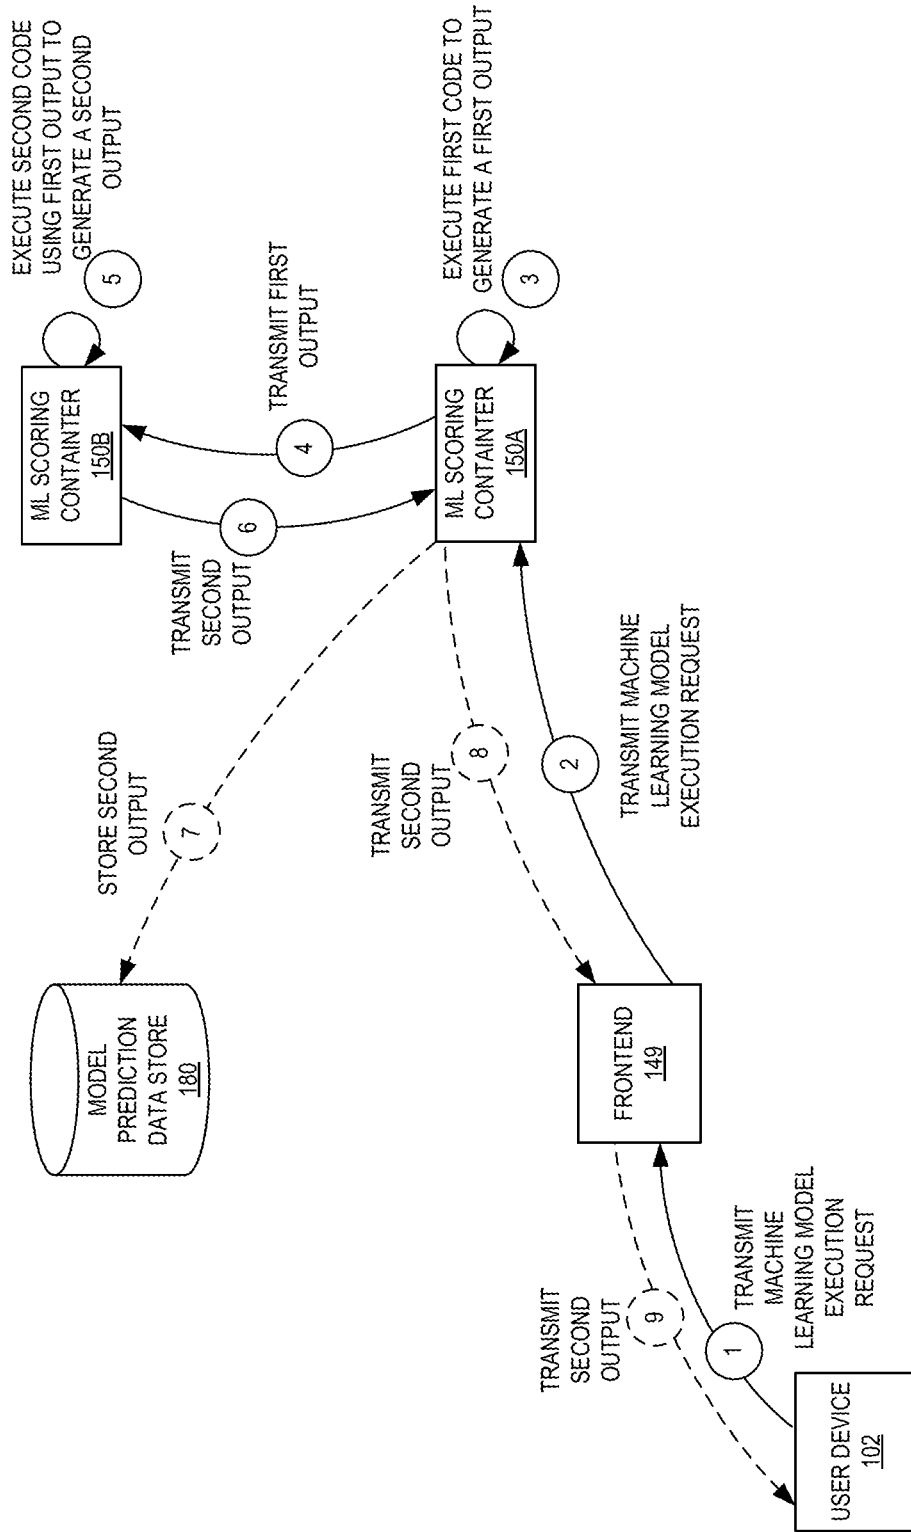


FIG. 6

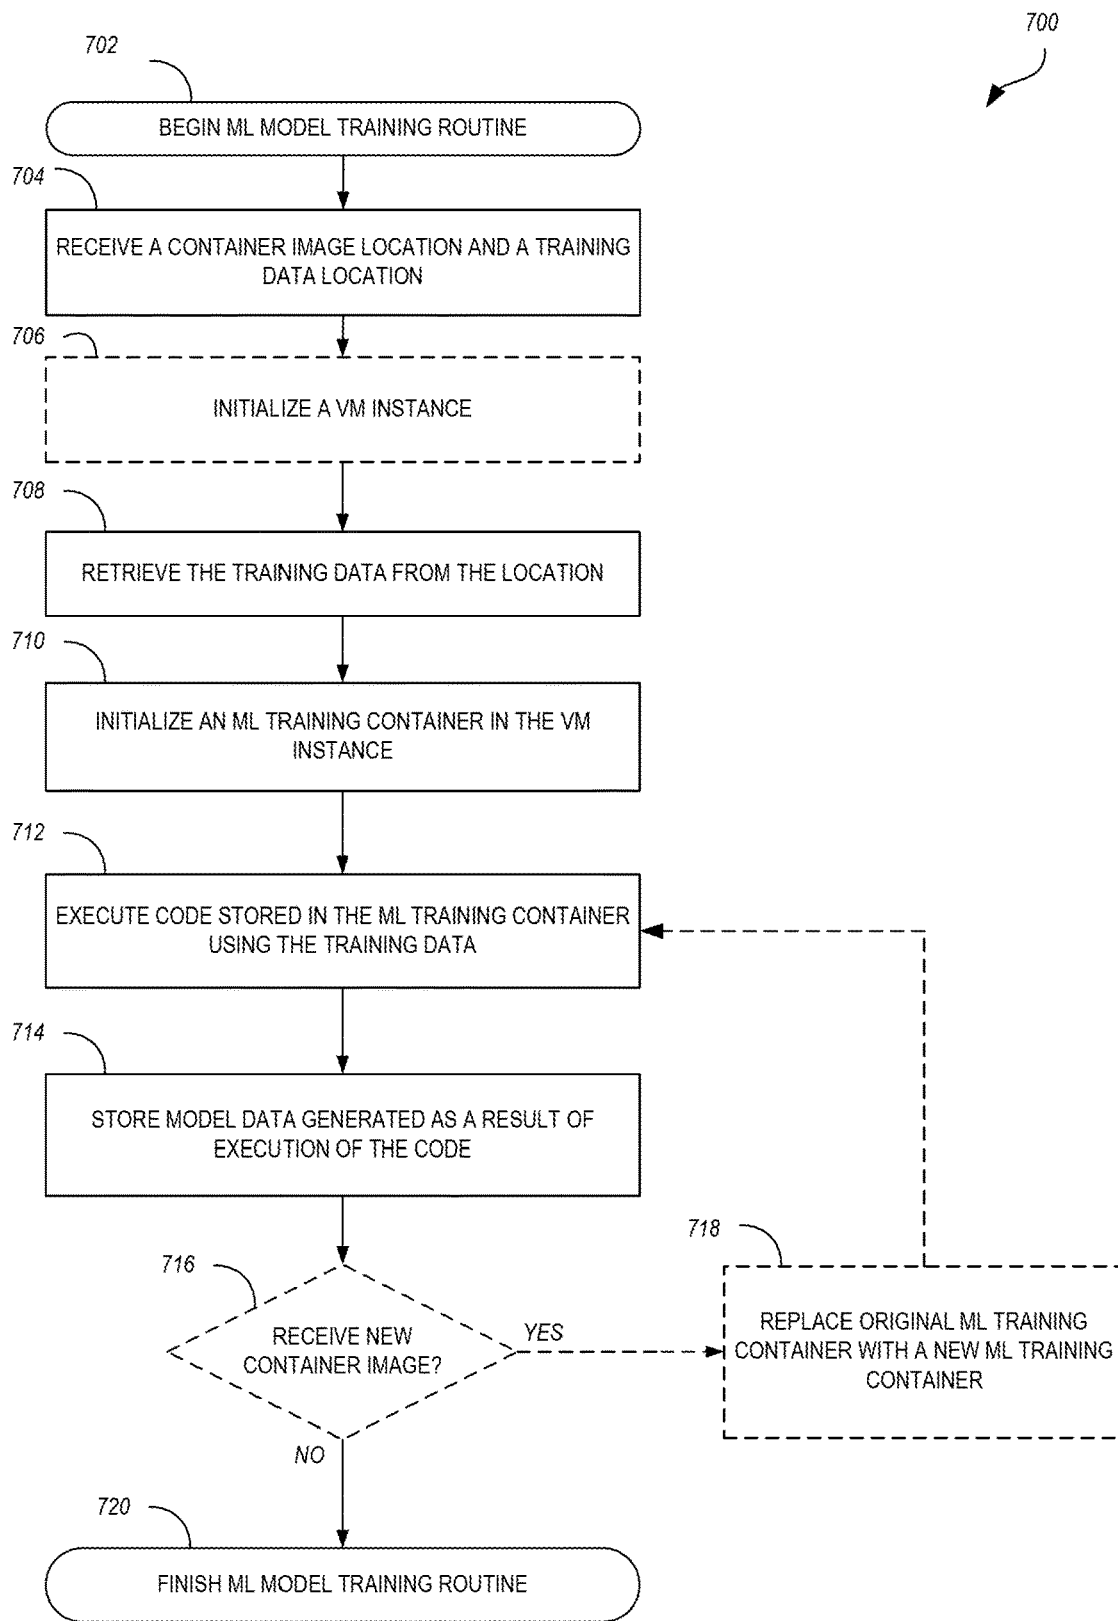


FIG. 7

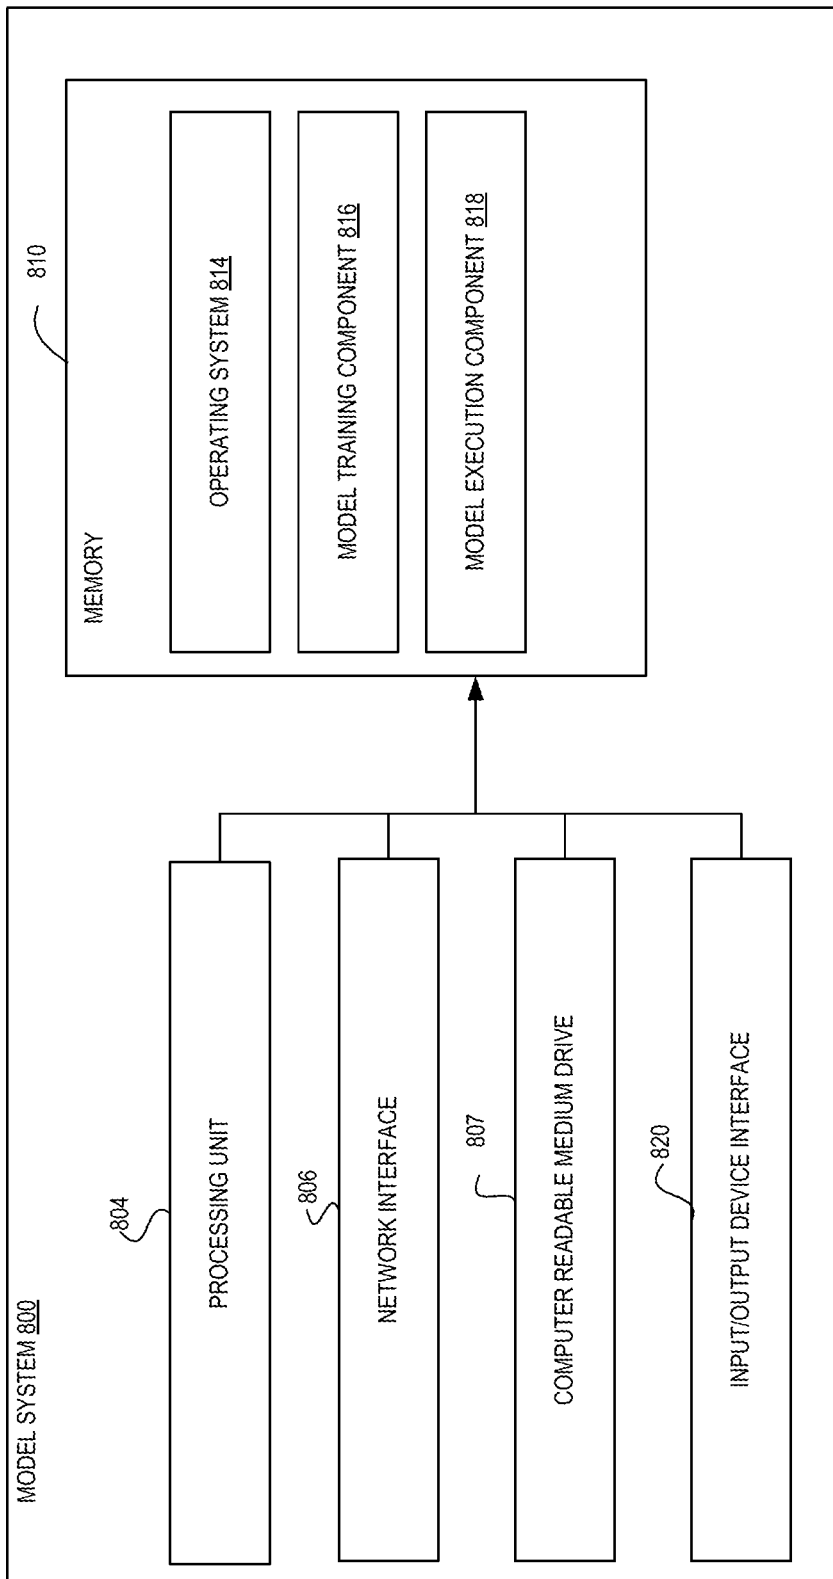


FIG. 8

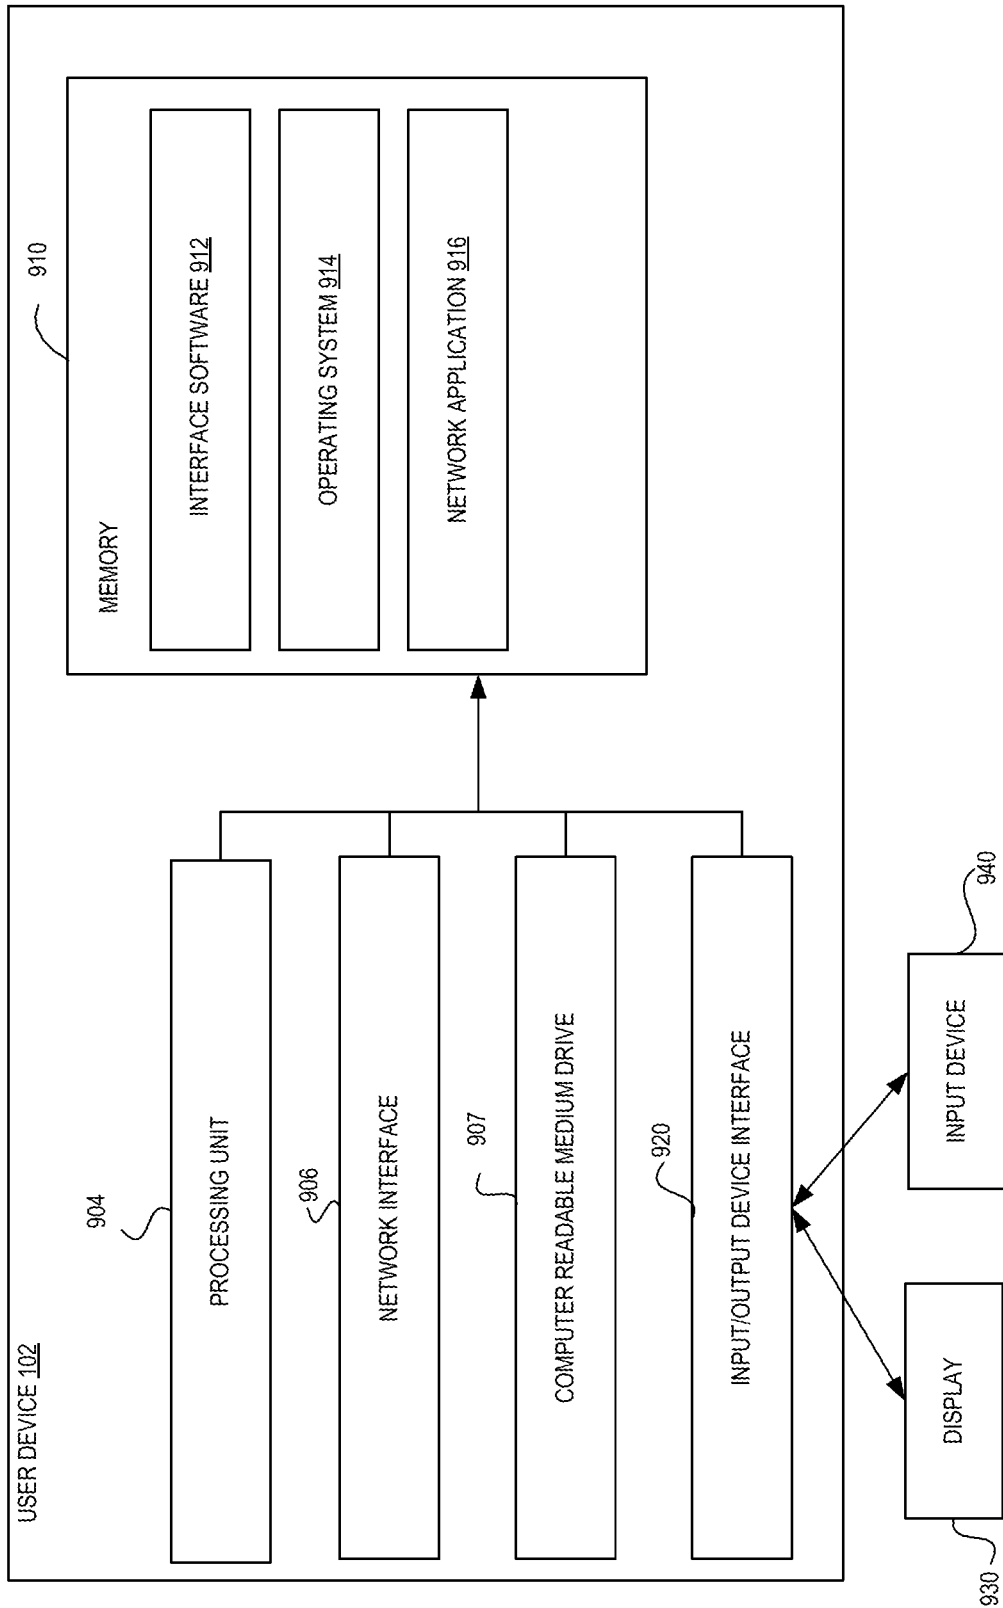


FIG. 9

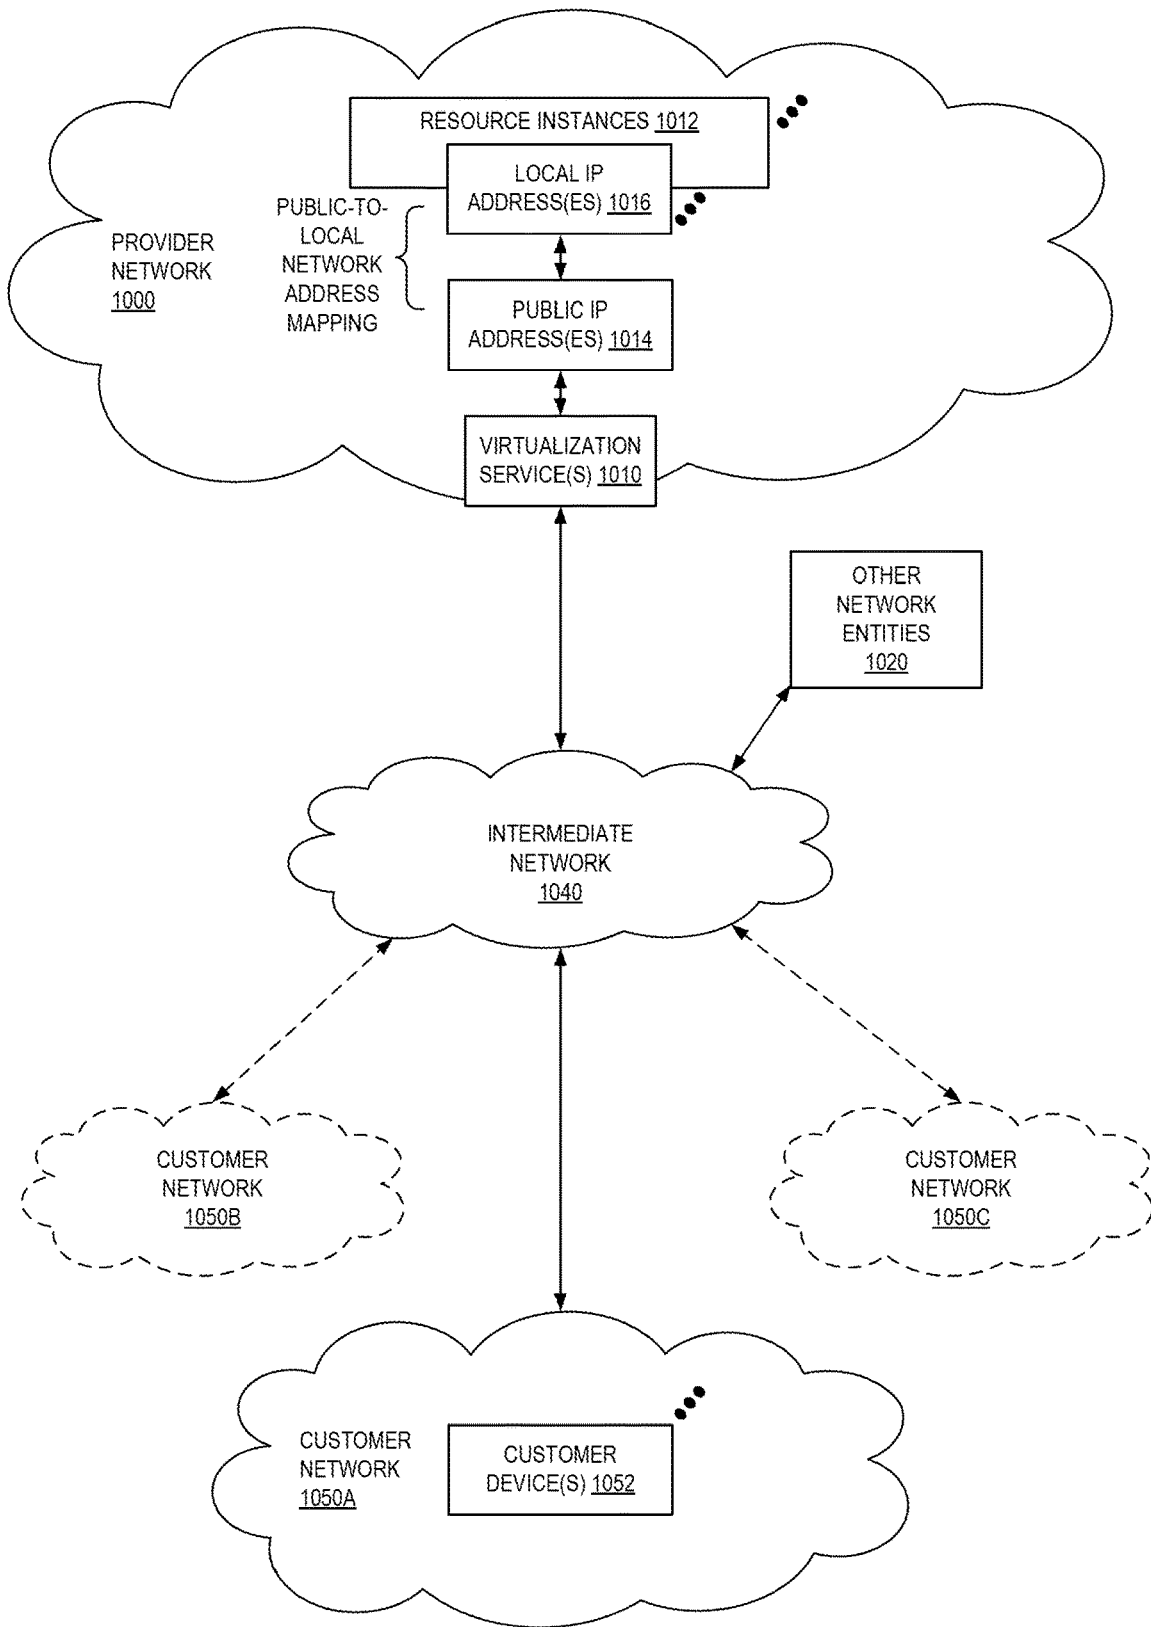


FIG. 10

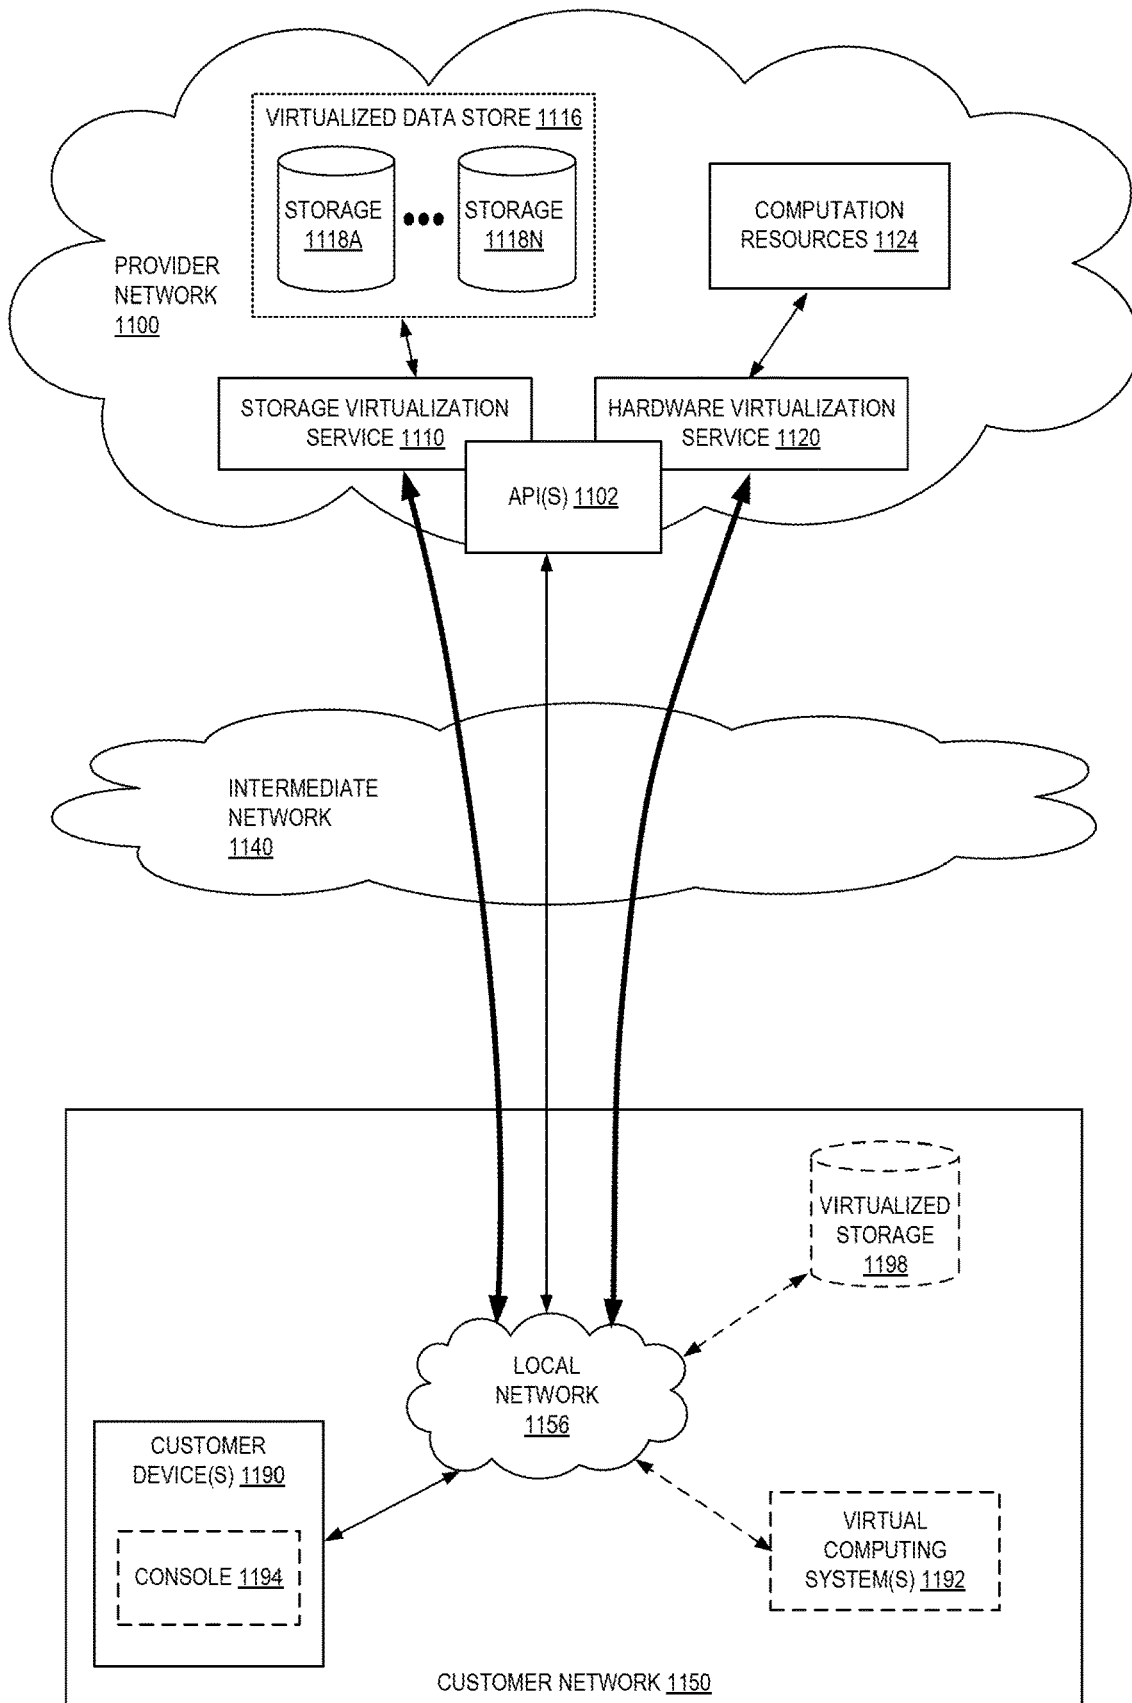


FIG. 11

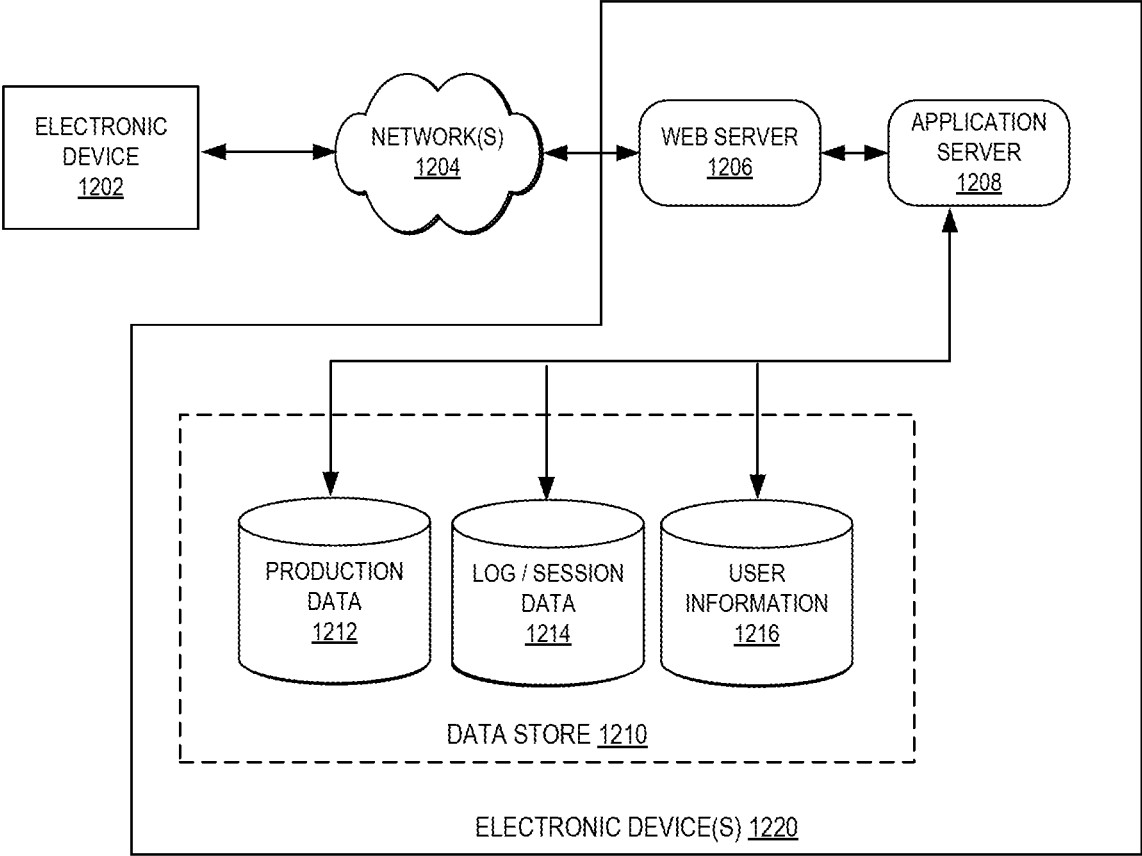


FIG. 12

1

PACKAGING AND DEPLOYING ALGORITHMS FOR FLEXIBLE MACHINE LEARNING

CROSS-REFERENCE TO RELATED APPLICATIONS

This application claims the benefit of U.S. Provisional Application No. 62/590,242, filed Nov. 22, 2017, which is hereby incorporated by reference.

BACKGROUND

Computing devices can utilize communication networks to exchange data. Companies and organizations operate computer networks that interconnect a number of computing devices to support operations or to provide services to third parties. The computing systems can be located in a single geographic location or located in multiple, distinct geographic locations (e.g., interconnected via private or public communication networks). Specifically, data centers or data processing centers, herein generally referred to as a “data center,” may include a number of interconnected computing systems to provide computing resources to users of the data center. To facilitate increased utilization of data center resources, virtualization technologies allow a single physical computing device to host one or more instances of virtual machines that appear and operate as independent computing devices to users of a data center.

BRIEF DESCRIPTION OF DRAWINGS

Various embodiments in accordance with the present disclosure will be described with reference to the drawings, in which:

FIG. 1 is a block diagram of an illustrative operating environment in which machine learning models are trained and hosted, in some embodiments.

FIG. 2 is a block diagram of the operating environment of FIG. 1 illustrating the operations performed by the components of the operating environment to train a machine learning model, according to some embodiments.

FIG. 3 is a block diagram of the operating environment of FIG. 1 illustrating the operations performed by the components of the operating environment to modifying machine learning model training, according to some embodiments.

FIG. 4 is a block diagram of the operating environment of FIG. 1 illustrating the operations performed by the components of the operating environment to parallelize the machine learning model training process, according to some embodiments.

FIG. 5A is a block diagram of the operating environment of FIG. 1 illustrating the operations performed by the components of the operating environment to deploy a trained machine learning model, according to some embodiments.

FIG. 5B is a block diagram of the operating environment of FIG. 1 illustrating the operations performed by the components of the operating environment to execute a trained machine learning model, according to some embodiments.

FIG. 6 is a block diagram of the operating environment of FIG. 1 illustrating the operations performed by the components of the operating environment to execute related machine learning models, according to some embodiments.

2

FIG. 7 is a flow diagram depicting a machine learning model training routine illustratively implemented by a model training system, according to some embodiments.

FIG. 8 depicts some embodiments of an architecture of an illustrative model system, such as the model training system and the model hosting system, that train and/or host machine learning models in accordance with the present application.

FIG. 9 depicts some embodiments of an architecture of an illustrative end user device that can receive data, prepare data, transmit training requests to the model training system, and transmit deployment and/or execution requests to the model hosting system in accordance with the present application.

FIG. 10 illustrates an example provider network environment according to some embodiments.

FIG. 11 is a block diagram of an example provider network that provides a storage virtualization service and a hardware virtualization service to customers according to some embodiments.

FIG. 12 illustrates an example of an environment for implementing aspects in accordance with various embodiments.

DETAILED DESCRIPTION

Various embodiments of methods, apparatus, systems, and non-transitory computer-readable storage media for packaging and deploying algorithms using containers for flexible machine learning. In some embodiments, users can create or utilize relatively simple containers adhering to a specification of a provider network, where the containers include code for how a machine learning model is to be trained and/or executed. The provider network can automatically train a model and/or host a model using the containers. The containers can use a wide variety of algorithms and use a variety of types of languages, libraries, data types, etc. Accordingly, users can simply perform machine learning training and hosting with extremely minimal knowledge of how the overall training and/or hosting is actually performed.

As described above, embodiments enable a single physical computing device (or multiple physical computing devices) to host one or more instances of virtual machines that appear and operate as independent computing devices to users. In some embodiments, a service provider can leverage virtualization technologies to provide a network-accessible machine learning service, such as the network-accessible machine learning model training and hosting system described herein. For example, the service provider can operate one or more physical computing devices accessible to user devices via a network. These physical computing device(s) can host virtual machine instances that are configured to train and/or execute machine learning models in response to commands received from user devices.

The embodiments described herein provide several technical benefits over conventional computing systems configured to train machine learning models. For example, training machine learning models can result in the usage of a large amount of processing power because machine learning models can be very complex and the amount of data used to train the models can be very large (e.g., in the gigabytes, terabytes, petabytes, etc.). Thus, some users acquire physically large conventional computing machines to perform the training. Users, however, may customize these conventional computing machines with specific software to execute the desired model training. On the other hand, embodiments described herein provide an environment in which users do

not have to generate and implement a large amount of customized code. Rather, users can simply provide just enough information to define a type of machine learning model to train, and the embodiments described herein can automatically initialize virtual machine instances, initialize containers, and/or perform other operations to implement a model training service.

On the other hand, embodiments described herein are configured to distribute the training across different physical computing devices in some embodiments. Thus, the time to train a model can be significantly reduced.

Valuable time can be lost if the resulting trained model turns out to be inaccurate. On the other hand, embodiments described herein can periodically evaluate models during the training process and output metrics corresponding to the evaluation. Thus, users can review the metrics to determine if, for example, a machine learning model being trained is inaccurate and whether it may be beneficial for the training job to be stopped.

Users can experience significant machine learning model training delays if a conventional computing machine is already in the process of training another model. On the other hand, embodiments described herein dynamically allocate computing resources to perform model training based on user demand in some embodiments. Thus, if a single user or multiple users desire to train multiple machine learning models during an overlapping time period, the trainings can be performed simultaneously.

These conventional services, however, are generally restricted to a single type of machine learning model and only allow prescribed data input formats. Users, on the other hand, may desire to train and use many different types of machine learning models that can receive different types of input data formats. Unlike these conventional services, embodiments described herein provide a flexible execution environment in which machine learning models can be trained and executed irrespective of the type of machine learning model, the programming language in which the machine learning model is defined, the data input format of the machine learning model, and/or the data output format of the machine learning model.

Example Machine Learning Model Training and Hosting Environment

FIG. 1 is a block diagram of an illustrative operating environment 100 in which machine learning models are trained and hosted, in some embodiments. The operating environment 100 includes end user devices 102, a model training system 120, a model hosting system 140, a training data store 160, a training metrics data store 165, a container data store 170, a training model data store 175, and a model prediction data store 180.

Example Model Training System

In some embodiments, users, by way of user devices 102, interact with the model training system 120 to provide data that causes the model training system 120 to train one or more machine learning models. A machine learning (ML) model, generally, may be thought of as one or more equations that are “trained” using a set of data. In some embodiments, the model training system 120 provides ML functionalities as a Web service, and thus messaging between user devices 102 and the model training system 120 (or provider network 199), and/or between components of the model training system 120 (or provider network 199), may

utilize HyperText Transfer Protocol (HTTP) messages to transfer data in a machine-readable file format, such as eXtensible Markup Language (XML) or JavaScript Object Notation (JSON).

The user devices 102 can interact with the model training system 120 via frontend 129 of the model training system 120. For example, a user device 102 can provide a training request to the frontend 129 that includes a container image (or multiple container images, or an identifier of one or multiple locations where container images are stored), an indicator of input data (e.g., an address or location of input data), one or more hyperparameter values (e.g., values indicating how the algorithm will operate, how many algorithms to run in parallel, how many clusters into which to separate data, etc.), and/or information describing the computing machine on which to train a machine learning model (e.g., a graphical processing unit (GPU) instance type, a central processing unit (CPU) instance type, an amount of memory to allocate, a type of virtual machine instance to use for training, etc.).

In some embodiments, the container image can include one or more layers, where each layer represents an executable instruction. Some or all of the executable instructions together represent an algorithm that defines a machine learning model. The executable instructions (e.g., the algorithm) can be written in any programming language (e.g., Python, Ruby, C++, Java, etc.). In some embodiments, the algorithm is pre-generated and obtained by a user, via the user device 102, from an algorithm repository (e.g., a network-accessible marketplace, a data store provided by a machine learning training service, etc.). In some embodiments, the algorithm is completely user-generated or partially user-generated (e.g., user-provided code modifies or configures existing algorithmic code).

In some embodiments, instead of providing a container image (or identifier thereof) in the training request, the user device 102 may provide, in the training request, an algorithm written in any programming language. The model training system 120 then packages the algorithm into a container (optionally with other code, such as a “base” ML algorithm supplemented with user-provided code) that is eventually loaded into a virtual machine instance 122 for training a machine learning model, as described in greater detail below. For example, a user, via a user device 102, may develop an algorithm/code using an application (e.g., an interactive web-based programming environment) and cause the algorithm/code to be provided—perhaps as part of a training request (or referenced in a training request)—to the model training system 120, where this algorithm/code may be containerized on its own or used together with an existing container having a machine learning framework, for example.

In some embodiments, instead of providing a container image in the training request, the user device 102 provides, in the training request, an indicator of a container image (e.g., an indication of an address or a location at which a container image is stored). For example, the container image can be stored in a container data store 170, and this container image may have been previously created/uploaded by the user. The model training system 120 can retrieve the container image from the indicated location and create a container using the retrieved container image. The container is then loaded into a virtual machine instance 122 for training a machine learning model, as described in greater detail below.

The model training system 120 can use the information provided by the user device 102 to train a machine learning

model in one or more pre-established virtual machine instances **122** in some embodiments. In particular, the model training system **120** includes a single physical computing device or multiple physical computing devices that are interconnected using one or more computing networks (not shown), where the physical computing device(s) host one or more virtual machine instances **122**. The model training system **120** can handle the acquisition and configuration of compute capacity (e.g., containers, instances, etc., which are described in greater detail below) based on the information describing the computing machine on which to train a machine learning model provided by the user device **102**. The model training system **120** can then train machine learning models using the compute capacity, as is described in greater detail below. The model training system **120** can automatically scale up and down based on the volume of training requests received from user devices **102** via front-end **129**, thereby relieving the user from the burden of having to worry about over-utilization (e.g., acquiring too little computing resources and suffering performance issues) or under-utilization (e.g., acquiring more computing resources than necessary to train the machine learning models, and thus overpaying).

In some embodiments, the virtual machine instances **122** are utilized to execute tasks. For example, such tasks can include training a machine learning model. As shown in FIG. 1, each virtual machine instance **122** includes an operating system (OS) **124**, a language runtime **126**, and one or more machine learning (ML) training containers **130**. Generally, the ML training containers **130** are logical units created within a virtual machine instance using the resources available on that instance, and can be utilized to isolate execution of a task from other processes (e.g., task executions) occurring in the instance. In some embodiments, the ML training containers **130** are formed from one or more container images and a top container layer. Each container image may further include one or more image layers, where each image layer represents an executable instruction. As described above, some or all of the executable instructions together represent an algorithm that defines a machine learning model. Changes made to the ML training containers **130** (e.g., creation of new files, modification of existing files, deletion of files, etc.) are stored in the top container layer. If a ML training container **130** is deleted, the top container layer is also deleted. However, the container image(s) that form a portion of the deleted ML training container **130** can remain unchanged. The ML training containers **130** can be implemented, for example, as Linux containers (LXC), Docker containers, and the like.

The ML training containers **130** may include individual copies of an OS **132** (e.g., portions of an OS, while OS kernel code may not be included within a container but instead be “shared” amongst containers), runtime **134**, and code **136** in some embodiments. The OS **132** and/or the runtime **134** can be defined by one or more executable instructions that form at least a portion of a container image that is used to form the ML training container **130** (e.g., the executable instruction(s) in the container image that define the operating system and/or runtime to run in the container formed from the container image). The code **136** includes one or more executable instructions that form at least a portion of a container image that is used to form the ML training container **130**. For example, the code **136** includes the executable instructions in the container image that represent an algorithm that defines a machine learning model. The OS **132** and/or runtime **134** are configured to execute the code **136** in response to an instruction to begin

machine learning model training. Execution of the code **136** results in the generation of model data, as described in greater detail below.

In some embodiments, the code **136** includes executable instructions that represent algorithms that define different machine learning models. For example, the code **136** includes one set of executable instructions that represent a first algorithm that defines a first machine learning model and a second set of executable instructions that represent a second algorithm that defines a second machine learning model. In some embodiments, the virtual machine instance **122** executes the code **136** and trains all of the machine learning models. In some embodiments, the virtual machine instance **122** executes the code **136**, selecting one of the machine learning models to train. For example, the virtual machine instance **122** can identify a type of training data indicated by the training request and select a machine learning model to train (e.g., execute the executable instructions that represent an algorithm that defines the selected machine learning model) that corresponds with the identified type of training data.

In some embodiments, the OS **132** and the runtime **134** are the same as the OS **124** and runtime **126** utilized by the virtual machine instance **122**. In some embodiments, the OS **132** and/or the runtime **134** are different than the OS **124** and/or runtime **126** utilized by the virtual machine instance **122**.

In some embodiments, the model training system **120** uses one or more container images included in a training request (or a container image retrieved from the container data store **170** in response to a received training request) to create and initialize a ML training container **130** in a virtual machine instance **122**. For example, the model training system **120** creates a ML training container **130** that includes the container image(s) and/or a top container layer.

Prior to beginning the training process, in some embodiments, the model training system **120** retrieves training data from the location indicated in the training request. For example, the location indicated in the training request can be a location in the training data store **160**. Thus, the model training system **120** retrieves the training data from the indicated location in the training data store **160**. In some embodiments, the model training system **120** does not retrieve the training data prior to beginning the training process. Rather, the model training system **120** streams the training data from the indicated location during the training process. For example, the model training system **120** can initially retrieve a portion of the training data and provide the retrieved portion to the virtual machine instance **122** training the machine learning model. Once the virtual machine instance **122** has applied and used the retrieved portion or once the virtual machine instance **122** is about to use all of the retrieved portion (e.g., a buffer storing the retrieved portion is nearly empty), then the model training system **120** can retrieve a second portion of the training data and provide the second retrieved portion to the virtual machine instance **122**, and so on.

To perform the machine learning model training, the virtual machine instance **122** executes code **136** stored in the ML training container **130** in some embodiments. For example, the code **136** includes some or all of the executable instructions that form the container image of the ML training container **130** initialized therein. Thus, the virtual machine instance **122** executes some or all of the executable instructions that form the container image of the ML training container **130** initialized therein to train a machine learning model. The virtual machine instance **122** executes some or

all of the executable instructions according to the hyperparameter values included in the training request. As an illustrative example, the virtual machine instance **122** trains a machine learning model by identifying values for certain parameters (e.g., coefficients, weights, centroids, etc.). The identified values depend on hyperparameters that define how the training is performed. Thus, the virtual machine instance **122** can execute the executable instructions to initiate a machine learning model training process, where the training process is run using the hyperparameter values included in the training request. Execution of the executable instructions can include the virtual machine instance **122** applying the training data retrieved by the model training system **120** as input parameters to some or all of the instructions being executed.

In some embodiments, executing the executable instructions causes the virtual machine instance **122** (e.g., the ML training container **130**) to generate model data. For example, the ML training container **130** generates model data and stores the model data in a file system of the ML training container **130**. The model data includes characteristics of the machine learning model being trained, such as a number of layers in the machine learning model, hyperparameters of the machine learning model, coefficients of the machine learning model, weights of the machine learning model, and/or the like. In particular, the generated model data includes values for the characteristics that define a machine learning model being trained. In some embodiments, executing the executable instructions causes a modification to the ML training container **130** such that the model data is written to the top container layer of the ML training container **130** and/or the container image(s) that forms a portion of the ML training container **130** is modified to include the model data.

The virtual machine instance **122** (or the model training system **120** itself) pulls the generated model data from the ML training container **130** and stores the generated model data in the training model data store **175** in an entry associated with the virtual machine instance **122** and/or the machine learning model being trained. In some embodiments, the virtual machine instance **122** generates a single file that includes model data and stores the single file in the training model data store **175**. In some embodiments, the virtual machine instance **122** generates multiple files during the course of training a machine learning model, where each file includes model data. In some embodiments, each model data file includes the same or different model data information (e.g., one file identifies the structure of an algorithm, another file includes a list of coefficients, etc.). The virtual machine instance **122** can package the multiple files into a single file once training is complete and store the single file in the training model data store **175**. Alternatively, the virtual machine instance **122** stores the multiple files in the training model data store **175**. The virtual machine instance **122** stores the file(s) in the training model data store **175** while the training process is ongoing and/or after the training process is complete.

In some embodiments, the virtual machine instance **122** regularly stores model data file(s) in the training model data store **175** as the training process is ongoing. Thus, model data file(s) can be stored in the training model data store **175** at different times during the training process. Each set of model data files corresponding to a particular time or each set of model data files present in the training model data store **175** as of a particular time could be checkpoints that represent different versions of a partially-trained machine learning model during different stages of the training pro-

cess. Accordingly, before training is complete, a user, via the user device **102** can submit a deployment and/or execution request in a manner as described below to deploy and/or execute a version of a partially trained machine learning model (e.g., a machine learning model trained as of a certain stage in the training process). A version of a partially-trained machine learning model can be based on some or all of the model data files stored in the training model data store **175**.

In some embodiments, a virtual machine instance **122** executes code **136** stored in a plurality of ML training containers **130**. For example, the algorithm included in the container image can be in a format that allows for the parallelization of the training process. Thus, the model training system **120** can create multiple copies of the container image provided in a training request and cause the virtual machine instance **122** to load each container image copy in a separate ML training container **130**. The virtual machine instance **122** can then execute, in parallel, the code **136** stored in the ML training containers **130**. The virtual machine instance **122** can further provide configuration information to each ML training container **130** (e.g., information indicating that N ML training containers **130** are collectively training a machine learning model and that a particular ML training container **130** receiving the configuration information is ML training container **130** number X of N), which can be included in the resulting model data. By parallelizing the training process, the model training system **120** can significantly reduce the training time in some embodiments.

In some embodiments, a plurality of virtual machine instances **122** execute code **136** stored in a plurality of ML training containers **130**. For example, the resources used to train a particular machine learning model can exceed the limitations of a single virtual machine instance **122**. However, the algorithm included in the container image can be in a format that allows for the parallelization of the training process. Thus, the model training system **120** can create multiple copies of the container image provided in a training request, initialize multiple virtual machine instances **122**, and cause each virtual machine instance **122** to load a container image copy in one or more separate ML training containers **130**. The virtual machine instances **122** can then each execute the code **136** stored in the ML training containers **130** in parallel. The model training system **120** can further provide configuration information to each ML training container **130** via the virtual machine instances **122** (e.g., information indicating that N ML training containers **130** are collectively training a machine learning model and that a particular ML training container **130** receiving the configuration information is ML training container **130** number X of N, information indicating that M virtual machine instances **122** are collectively training a machine learning model and that a particular ML training container **130** receiving the configuration information is initialized in virtual machine instance **122** number Y of M, etc.), which can be included in the resulting model data. As described above, by parallelizing the training process, the model training system **120** can significantly reduce the training time in some embodiments.

In some embodiments, the model training system **120** includes a plurality of physical computing devices and two or more of the physical computing devices hosts one or more virtual machine instances **122** that execute the code **136**. Thus, the parallelization can occur over different physical computing devices in addition to over different virtual machine instances **122** and/or ML training containers **130**.

In some embodiments, the model training system **120** includes a ML model evaluator **128**. The ML model evalu-

ator **128** can monitor virtual machine instances **122** as machine learning models are being trained, obtaining the generated model data and processing the obtained model data to generate model metrics. For example, the model metrics can include quality metrics, such as an error rate of the machine learning model being trained, a statistical distribution of the machine learning model being trained, a latency of the machine learning model being trained, a confidence level of the machine learning model being trained (e.g., a level of confidence that the accuracy of the machine learning model being trained is known, etc. The ML model evaluator **128** can obtain the model data for a machine learning model being trained and evaluation data from the training data store **160**. The evaluation data is separate from the data used to train a machine learning model and includes both input data and expected outputs (e.g., known results), and thus the ML model evaluator **128** can define a machine learning model using the model data and execute the machine learning model by providing the input data as inputs to the machine learning model. The ML model evaluator **128** can then compare the outputs of the machine learning model to the expected outputs, and determine one or more quality metrics of the machine learning model being trained based on the comparison (e.g., the error rate can be a difference or distance between the machine learning model outputs and the expected outputs).

The ML model evaluator **128** periodically generates model metrics during the training process and stores the model metrics in the training metrics data store **165** in some embodiments. While the machine learning model is being trained, a user, via the user device **102**, can access and retrieve the model metrics from the training metrics data store **165**. The user can then use the model metrics to determine whether to adjust the training process and/or to stop the training process. For example, the model metrics can indicate that the machine learning model is performing poorly (e.g., has an error rate above a threshold value, has a statistical distribution that is not an expected or desired distribution (e.g., not a binomial distribution, a Poisson distribution, a geometric distribution, a normal distribution, Gaussian distribution, etc.), has an execution latency above a threshold value, has a confidence level below a threshold value) and/or is performing progressively worse (e.g., the quality metric continues to worsen over time). In response, in some embodiments, the user, via the user device **102**, can transmit a request to the model training system **120** to modify the machine learning model being trained (e.g., transmit a modification request). The request can include a new or modified container image, a new or modified algorithm, new or modified hyperparameter(s), and/or new or modified information describing the computing machine on which to train a machine learning model. The model training system **120** can modify the machine learning model accordingly. For example, the model training system **120** can cause the virtual machine instance **122** to optionally delete an existing ML training container **130**, create and initialize a new ML training container **130** using some or all of the information included in the request, and execute the code **136** stored in the new ML training container **130** to restart the machine learning model training process. As another example, the model training system **120** can cause the virtual machine instance **122** to modify the execution of code stored in an existing ML training container **130** according to the data provided in the modification request. In some embodiments, the user, via the user device **102**, can transmit a request to the model training system **120** to stop the machine learning model training process. The model training system

120 can then instruct the virtual machine instance **122** to delete the ML training container **130** and/or to delete any model data stored in the training model data store **175**.

As described below, in some embodiments, the model data stored in the training model data store **175** is used by the model hosting system **140** to deploy machine learning models. Alternatively or in addition, a user device **102** or another computing device (not shown) can retrieve the model data from the training model data store **175** to implement a learning algorithm in an external device. As an illustrative example, a robotic device can include sensors to capture input data. A user device **102** can retrieve the model data from the training model data store **175** and store the model data in the robotic device. The model data defines a machine learning model. Thus, the robotic device can provide the captured input data as an input to the machine learning model, resulting in an output. The robotic device can then perform an action (e.g., move forward, raise an arm, generate a sound, etc.) based on the resulting output.

While the virtual machine instances **122** are shown in FIG. 1 as a single grouping of virtual machine instances **122**, some embodiments of the present application separate virtual machine instances **122** that are actively assigned to execute tasks from those virtual machine instances **122** that are not actively assigned to execute tasks. For example, those virtual machine instances **122** actively assigned to execute tasks are grouped into an “active pool,” while those virtual machine instances **122** not actively assigned to execute tasks are placed within a “warming pool.” In some embodiments, those virtual machine instances **122** within the warming pool can be pre-initialized with an operating system, language runtimes, and/or other software required to enable rapid execution of tasks (e.g., rapid initialization of machine learning model training in ML training container(s) **130**) in response to training requests.

In some embodiments, the model training system **120** includes a processing unit, a network interface, a computer-readable medium drive, and an input/output device interface, all of which can communicate with one another by way of a communication bus. The network interface can provide connectivity to one or more networks or computing systems. The processing unit can thus receive information and instructions from other computing systems or services (e.g., user devices **102**, the model hosting system **140**, etc.). The processing unit can also communicate to and from a memory of a virtual machine instance **122** and further provide output information for an optional display via the input/output device interface. The input/output device interface can also accept input from an optional input device. The memory can contain computer program instructions (grouped as modules in some embodiments) that the processing unit executes in order to implement one or more aspects of the present disclosure.

Example Model Hosting System

In some embodiments, the model hosting system **140** includes a single physical computing device or multiple physical computing devices that are interconnected using one or more computing networks (not shown), where the physical computing device(s) host one or more virtual machine instances **142**. The model hosting system **140** can handle the acquisition and configuration of compute capacity (e.g., containers, instances, etc.) based on demand for the execution of trained machine learning models. The model hosting system **140** can then execute machine learning models using the compute capacity, as is described in greater

detail below. The model hosting system **140** can automatically scale up and down based on the volume of execution requests received from user devices **102** via frontend **149** of the model hosting system **140**, thereby relieving the user from the burden of having to worry about over-utilization (e.g., acquiring too little computing resources and suffering performance issues) or under-utilization (e.g., acquiring more computing resources than necessary to run the machine learning models, and thus overpaying).

In some embodiments, the virtual machine instances **142** are utilized to execute tasks. For example, such tasks can include executing a machine learning model. As shown in FIG. 1, each virtual machine instance **142** includes an operating system (OS) **144**, a language runtime **146**, and one or more ML scoring containers **150**. The ML scoring containers **150** are similar to the ML training containers **130** in that the ML scoring containers **150** are logical units created within a virtual machine instance using the resources available on that instance, and can be utilized to isolate execution of a task from other processes (e.g., task executions) occurring in the instance. In some embodiments, the ML scoring containers **150** are formed from one or more container images and a top container layer. Each container image further includes one or more image layers, where each image layer represents an executable instruction. As described above, some or all of the executable instructions together represent an algorithm that defines a machine learning model. Changes made to the ML scoring containers **150** (e.g., creation of new files, modification of existing files, deletion of files, etc.) are stored in the top container layer. If a ML scoring container **150** is deleted, the top container layer is also deleted. However, the container image(s) that form a portion of the deleted ML scoring container **150** can remain unchanged. The ML scoring containers **150** can be implemented, for example, as Linux containers.

The ML scoring containers **150** each include individual copies of an OS **152**, runtime **154**, and code **156** in some embodiments. The OS **152** and/or the runtime **154** can be defined by one or more executable instructions that form at least a portion of a container image that is used to form the ML scoring container **150** (e.g., the executable instruction(s) in the container image that define the operating system and/or runtime to run in the container formed from the container image). The code **156** includes one or more executable instructions that form at least a portion of a container image that is used to form the ML scoring container **150**. For example, the code **156** includes the executable instructions in the container image that represent an algorithm that defines a machine learning model. The code **156** can also include model data that represent characteristics of the defined machine learning model, as described in greater detail below. The OS **152** and/or runtime **154** are configured to execute the code **156** in response to an instruction to begin execution of a machine learning model. Execution of the code **156** results in the generation of outputs (e.g., predicted results), as described in greater detail below.

In some embodiments, the OS **152** and the runtime **154** are the same as the OS **144** and runtime **146** utilized by the virtual machine instance **142**. In some embodiments, the OS **152** and/or the runtime **154** are different than the OS **144** and/or runtime **146** utilized by the virtual machine instance **142**.

In some embodiments, the model hosting system **140** uses one or more container images included in a deployment request (or a container image retrieved from the container data store **170** in response to a received deployment request)

to create and initialize a ML scoring container **150** in a virtual machine instance **142**. For example, the model hosting system **140** creates a ML scoring container **150** that includes the container image(s) and/or a top container layer.

As described above, a user device **102** can submit a deployment request and/or an execution request to the model hosting system **140** via the frontend **149** in some embodiments. A deployment request causes the model hosting system **140** to deploy a trained machine learning model into a virtual machine instance **142**. For example, the deployment request can include an identification of an endpoint (e.g., an endpoint name, such as an HTTP endpoint name) and an identification of one or more trained machine learning models (e.g., a location of one or more model data files stored in the training model data store **175**). Optionally, the deployment request also includes an identification of one or more container images stored in the container data store **170**.

Upon receiving the deployment request, the model hosting system **140** initializes one or more ML scoring containers **150** in one or more hosted virtual machine instance **142**. In embodiments in which the deployment request includes an identification of one or more container images, the model hosting system **140** forms the ML scoring container(s) **150** from the identified container image(s). For example, a container image identified in a deployment request can be the same container image used to form an ML training container **130** used to train the machine learning model corresponding to the deployment request. Thus, the code **156** of the ML scoring container(s) **150** includes one or more executable instructions in the container image(s) that represent an algorithm that defines a machine learning model. In embodiments in which the deployment request does not include an identification of a container image, the model hosting system **140** forms the ML scoring container (s) **150** from one or more container images stored in the container data store **170** that are appropriate for executing the identified trained machine learning model(s). For example, an appropriate container image can be a container image that includes executable instructions that represent an algorithm that defines the identified trained machine learning model(s).

The model hosting system **140** further forms the ML scoring container(s) **150** by retrieving model data corresponding to the identified trained machine learning model(s) in some embodiments. For example, the deployment request can identify a location of model data file(s) stored in the training model data store **175**. In embodiments in which a single model data file is identified in the deployment request, the model hosting system **140** retrieves the identified model data file from the training model data store **175** and inserts the model data file into a single ML scoring container **150**, which forms a portion of code **156**. In some embodiments, the model data file is archived or compressed (e.g., formed from a package of individual files). Thus, the model hosting system **140** unarchives or decompresses the model data file to obtain multiple individual files, and inserts the individual files into the ML scoring container **150**. In some embodiments, the model hosting system **140** stores the model data file in the same location as the location in which the model data file was stored in the ML training container **130** that generated the model data file. For example, the model data file initially was stored in the top container layer of the ML training container **130** at a certain offset, and the model hosting system **140** then stores the model data file in the top container layer of the ML scoring container **150** at the same offset.

In embodiments in which multiple model data files are identified in the deployment request, the model hosting system **140** retrieves the identified model data files from the training model data store **175**. The model hosting system **140** can insert the model data files into the same ML scoring container **150**, into different ML scoring containers **150** initialized in the same virtual machine instance **142**, or into different ML scoring containers **150** initialized in different virtual machine instances **142**. As an illustrative example, the deployment request can identify multiple model data files corresponding to different trained machine learning models because the trained machine learning models are related (e.g., the output of one trained machine learning model is used as an input to another trained machine learning model). Thus, the user may desire to deploy multiple machine learning models to eventually receive a single output that relies on the outputs of multiple machine learning models.

In some embodiments, the model hosting system **140** associates the initialized ML scoring container(s) **150** with the endpoint identified in the deployment request. For example, each of the initialized ML scoring container(s) **150** can be associated with a network address. The model hosting system **140** can map the network address(es) to the identified endpoint, and the model hosting system **140** or another system (e.g., a routing system, not shown) can store the mapping. Thus, a user device **102** can refer to trained machine learning model(s) stored in the ML scoring container(s) **150** using the endpoint. This allows for the network address of an ML scoring container **150** to change without causing the user operating the user device **102** to change the way in which the user refers to a trained machine learning model.

Once the ML scoring container(s) **150** are initialized, the ML scoring container(s) **150** are ready to execute trained machine learning model(s). In some embodiments, the user device **102** transmits an execution request to the model hosting system **140** via the frontend **149**, where the execution request identifies an endpoint and includes an input to a machine learning model (e.g., a set of input data). The model hosting system **140** or another system (e.g., a routing system, not shown) can obtain the execution request, identify the ML scoring container(s) **150** corresponding to the identified endpoint, and route the input to the identified ML scoring container(s) **150**.

In some embodiments, a virtual machine instance **142** executes the code **156** stored in an identified ML scoring container **150** in response to the model hosting system **140** receiving the execution request. In particular, execution of the code **156** causes the executable instructions in the code **156** corresponding to the algorithm to read the model data file stored in the ML scoring container **150**, use the input included in the execution request as an input parameter, and generate a corresponding output. As an illustrative example, the algorithm can include coefficients, weights, layers, cluster centroids, and/or the like. The executable instructions in the code **156** corresponding to the algorithm can read the model data file to determine values for the coefficients, weights, layers, cluster centroids, and/or the like. The executable instructions can include input parameters, and the input included in the execution request can be supplied by the virtual machine instance **142** as the input parameters. With the machine learning model characteristics and the input parameters provided, execution of the executable instructions by the virtual machine instance **142** can be completed, resulting in an output.

In some embodiments, the virtual machine instance **142** stores the output in the model prediction data store **180**. Alternatively or in addition, the virtual machine instance **142** transmits the output to the user device **102** that submitted the execution result via the frontend **149**.

In some embodiments, the execution request corresponds to a group of related trained machine learning models. Thus, the ML scoring container **150** can transmit the output to a second ML scoring container **150** initialized in the same virtual machine instance **142** or in a different virtual machine instance **142**. The virtual machine instance **142** that initialized the second ML scoring container **150** can then execute second code **156** stored in the second ML scoring container **150**, providing the received output as an input parameter to the executable instructions in the second code **156**. The second ML scoring container **150** further includes a model data file stored therein, which is read by the executable instructions in the second code **156** to determine values for the characteristics defining the machine learning model. Execution of the second code **156** results in a second output. The virtual machine instance **142** that initialized the second ML scoring container **150** can then transmit the second output to the model prediction data store **180** and/or the user device **102** via the frontend **149** (e.g., if no more trained machine learning models are needed to generate an output) or transmit the second output to a third ML scoring container **150** initialized in the same or different virtual machine instance **142** (e.g., if outputs from one or more additional trained machine learning models are needed), and the above-referenced process can be repeated with respect to the third ML scoring container **150**.

While the virtual machine instances **142** are shown in FIG. 1 as a single grouping of virtual machine instances **142**, some embodiments of the present application separate virtual machine instances **142** that are actively assigned to execute tasks from those virtual machine instances **142** that are not actively assigned to execute tasks. For example, those virtual machine instances **142** actively assigned to execute tasks are grouped into an "active pool," while those virtual machine instances **142** not actively assigned to execute tasks are placed within a "warming pool." In some embodiments, those virtual machine instances **142** within the warming pool can be pre-initialized with an operating system, language runtimes, and/or other software required to enable rapid execution of tasks (e.g., rapid initialization of ML scoring container(s) **150**, rapid execution of code **156** in ML scoring container(s), etc.) in response to deployment and/or execution requests.

In some embodiments, the model hosting system **140** includes a processing unit, a network interface, a computer-readable medium drive, and an input/output device interface, all of which can communicate with one another by way of a communication bus. The network interface can provide connectivity to one or more networks or computing systems. The processing unit can thus receive information and instructions from other computing systems or services (e.g., user devices **102**, the model training system **120**, etc.). The processing unit can also communicate to and from a memory of a virtual machine instance **142** and further provide output information for an optional display via the input/output device interface. The input/output device interface can also accept input from an optional input device. The memory can contain computer program instructions (grouped as modules in some embodiments) that the processing unit executes in order to implement one or more aspects of the present disclosure.

Additional Embodiments of the Example Training and Hosting Environment

In some embodiments, the operating environment **100** supports many different types of machine learning models, such as multi arm bandit models, reinforcement learning models, ensemble machine learning models, deep learning models, and/or the like.

The model training system **120** and the model hosting system **140** depicted in FIG. **1** are not meant to be limiting. For example, the model training system **120** and/or the model hosting system **140** could also operate within a computing environment having a fewer or greater number of devices than are illustrated in FIG. **1**. Thus, the depiction of the model training system **120** and/or the model hosting system **140** in FIG. **1** may be taken as illustrative and not limiting to the present disclosure. For example, the model training system **120** and/or the model hosting system **140** or various constituents thereof could implement various Web services components, hosted or “cloud” computing environments, and/or peer-to-peer network configurations to implement at least a portion of the processes described herein. In some embodiments, the model training system **120** and/or the model hosting system **140** are implemented directly in hardware or software executed by hardware devices and may, for instance, include one or more physical or virtual servers implemented on physical computer hardware configured to execute computer-executable instructions for performing the various features that are described herein. The one or more servers can be geographically dispersed or geographically co-located, for instance, in one or more points of presence (POPs) or regional data centers.

The frontend **129** processes all training requests received from user devices **102** and provisions virtual machine instances **122**. In some embodiments, the frontend **129** serves as a front door to all the other services provided by the model training system **120**. The frontend **129** processes the requests and makes sure that the requests are properly authorized. For example, the frontend **129** may determine whether the user associated with the training request is authorized to initiate the training process.

Similarly, frontend **149** processes all deployment and execution requests received from user devices **102** and provisions virtual machine instances **142**. In some embodiments, the frontend **149** serves as a front door to all the other services provided by the model hosting system **140**. The frontend **149** processes the requests and makes sure that the requests are properly authorized. For example, the frontend **149** may determine whether the user associated with a deployment request or an execution request is authorized to access the indicated model data and/or to execute the indicated machine learning model.

The training data store **160** stores training data and/or evaluation data. The training data can be data used to train machine learning models and evaluation data can be data used to evaluate the performance of machine learning models. In some embodiments, the training data and the evaluation data have common data. In some embodiments, the training data and the evaluation data do not have common data. In some embodiments, the training data includes input data and expected outputs. While the training data store **160** is depicted as being located external to the model training system **120** and the model hosting system **140**, this is not meant to be limiting. For example, in some embodiments not shown, the training data store **160** is located internal to at least one of the model training system **120** or the model hosting system **140**.

In some embodiments, the training metrics data store **165** stores model metrics. While the training metrics data store **165** is depicted as being located external to the model training system **120** and the model hosting system **140**, this is not meant to be limiting. For example, in some embodiments not shown, the training metrics data store **165** is located internal to at least one of the model training system **120** or the model hosting system **140**.

The container data store **170** stores container images, such as container ML training images used to form ML training containers **130** and/or ML scoring containers **150**, that can be retrieved by various virtual machine instances **122** and/or **142**. While the container data store **170** is depicted as being located external to the model training system **120** and the model hosting system **140**, this is not meant to be limiting. For example, in some embodiments not shown, the container data store **170** is located internal to at least one of the model training system **120** and the model hosting system **140**.

The training model data store **175** stores model data files. In some embodiments, some of the model data files are comprised of a single file, while other model data files are packages of multiple individual files. While the training model data store **175** is depicted as being located external to the model training system **120** and the model hosting system **140**, this is not meant to be limiting. For example, in some embodiments not shown, the training model data store **175** is located internal to at least one of the model training system **120** or the model hosting system **140**.

The model prediction data store **180** stores outputs (e.g., execution results) generated by the ML scoring containers **150** in some embodiments. While the model prediction data store **180** is depicted as being located external to the model training system **120** and the model hosting system **140**, this is not meant to be limiting. For example, in some embodiments not shown, the model prediction data store **180** is located internal to at least one of the model training system **120** and the model hosting system **140**.

While the model training system **120**, the model hosting system **140**, the training data store **160**, the training metrics data store **165**, the container data store **170**, the training model data store **175**, and the model prediction data store **180** are illustrated as separate components, this is not meant to be limiting. In some embodiments, any one or all of these components can be combined to perform the functionality described herein. For example, any one or all of these components can be implemented by a single computing device, or by multiple distinct computing devices, such as computer servers, logically or physically grouped together to collectively operate as a server system. Any one or all of these components can communicate via a shared internal network, and the collective system (e.g., also referred to herein as a machine learning service) can communicate with one or more of the user devices **102** via the network **110**.

Various example user devices **102** are shown in FIG. **1**, including a desktop computer, laptop, and a mobile phone, each provided by way of illustration. In general, the user devices **102** can be any computing device such as a desktop, laptop or tablet computer, personal computer, wearable computer, server, personal digital assistant (PDA), hybrid PDA/mobile phone, mobile phone, electronic book reader, set-top box, voice command device, camera, digital media player, and the like. In some embodiments, the model training system **120** and/or the model hosting system **140** provides the user devices **102** with one or more user interfaces, command-line interfaces (CLI), application programming interfaces (API), and/or other programmatic interfaces for submitting training requests, deployment requests, and/

or execution requests. In some embodiments, the user devices **102** can execute a stand-alone application that interacts with the model training system **120** and/or the model hosting system **140** for submitting training requests, deployment requests, and/or execution requests.

In some embodiments, the network **110** includes any wired network, wireless network, or combination thereof. For example, the network **110** may be a personal area network, local area network, wide area network, over-the-air broadcast network (e.g., for radio or television), cable network, satellite network, cellular telephone network, or combination thereof. As a further example, the network **110** may be a publicly accessible network of linked networks, possibly operated by various distinct parties, such as the Internet. In some embodiments, the network **110** may be a private or semi-private network, such as a corporate or university intranet. The network **110** may include one or more wireless networks, such as a Global System for Mobile Communications (GSM) network, a Code Division Multiple Access (CDMA) network, a Long Term Evolution (LTE) network, or any other type of wireless network. The network **110** can use protocols and components for communicating via the Internet or any of the other aforementioned types of networks. For example, the protocols used by the network **110** may include HTTP, HTTP Secure (HTTPS), Message Queue Telemetry Transport (MQTT), Constrained Application Protocol (CoAP), and the like. Protocols and components for communicating via the Internet or any of the other aforementioned types of communication networks are well known to those skilled in the art and, thus, are not described in more detail herein.

Example Block Diagram for Training a Machine Learning Model

FIG. 2 is a block diagram of the operating environment **100** of FIG. 1 illustrating the operations performed by the components of the operating environment **100** to train a machine learning model, according to some embodiments. As illustrated in FIG. 2, the user device **102** transmits a location of a container image and a location of training data to the frontend **129** at (1). The frontend **129** then causes a virtual machine instance **122** to be initialized and forwards the container image location and the training data location to the initialized virtual machine instance **122** at (2). In some embodiments, the container image location and the training data location are transmitted as part of a training request.

In some embodiments, the virtual machine instance **122** retrieves training data from the training data store **160** using the received location at (3). Before, during, or after retrieving the training data, the virtual machine instance **122** retrieves the container image from the container data store **170** using the received location at (4).

The virtual machine instance **122** initializes an ML training container within the virtual machine instance **122** using the received container image in some embodiments. The virtual machine instance **122** then executes code stored in the ML training container using the retrieved training data at (5) to train a machine learning model. For example, the code can include executable instructions originating in the container image that represent an algorithm that defines a machine learning model that is yet to be trained. The virtual machine instance **122** executes the code according to hyperparameter values that are provided by the user device **102**.

Executing the executable instructions causes the ML training container to generate model data that includes characteristics of the machine learning model being trained.

The virtual machine instance **122** stores the model data in the training model data store **175** at (6) in some embodiments. In some embodiments, the virtual machine instance **122** generates multiple model data files that are packaged into a single file stored in the training model data store **175**.

During the machine learning model training process, the ML model evaluator **128** can retrieve the model data from the training model data store **175** at (7). The ML model evaluator **128** further retrieves evaluation data from the training data store **160** at (8). For example, the evaluation data can be data that is separate from the data used to train machine learning models. The evaluation data can include input data and known results that occurred or were formed as a result of the input data. In some embodiments, the ML model evaluator **128** executes a machine learning model defined by the retrieved model data using input data included in the evaluation data at (9). The ML model evaluator **128** then compares outputs of the machine learning model defined by the retrieved model data with known results included in the evaluation data to determine a quality metric of the machine learning model at (10). For example, the quality metric can be determined based on an aggregated difference (e.g., average difference, median difference, etc.) between the machine learning model outputs and the known results. The ML model evaluator **128** can then store the quality metric in the training metrics data store **165** at (11).

In some embodiments, the ML model evaluator **128** also stores additional information in the training metrics data store **165**. For example, the ML model evaluator **128** can store the input data (or tags that represent the input data), the machine learning model outputs, and the known results. Thus, a user, via the user device **102**, can not only identify the quality metric(s), but can also identify which inputs resulted in small or no differences between machine learning model outputs and known results, which inputs resulted in large differences between machine learning model outputs and known results, etc.

Example Block Diagram for Modifying Machine Learning Model Training

FIG. 3 is a block diagram of the operating environment **100** of FIG. 1 illustrating the operations performed by the components of the operating environment **100** to modifying machine learning model training, according to some embodiments. As illustrated in FIG. 3, the user device **102** retrieves a quality metric stored in the training metrics data store **165** at (1). In some embodiments, a user, via the user device **102**, retrieves the quality metric to determine the accuracy of a machine learning model still being trained.

In some embodiments, the user device **102** transmits a location of a modified container image to the frontend **129** at (2). The frontend **129** then forwards the location of modified container image to the virtual machine instance **122** at (3). The user device **102** can transmit the modified container image as part of a modification request to modify the machine learning model being trained. In response, the virtual machine instance **122** stops execution of the code stored in the original ML training container formed from the original container image at (4). The virtual machine instance **122** then retrieves the modified container image from the container data store **170** at (5) using the received location. The virtual machine instance **122** can then form a modified ML training container from the modified container image, and execute code stored in the modified ML training container using previously retrieved training data at (6) to re-train a machine learning model.

Execution of the code causes the modified ML training container to generate updated model data, which the virtual machine instance 122 then stores in the training model data store 175 at (7). In some embodiments, not shown, the virtual machine instance 122 causes the training model data store 175 to delete any model data stored as a result of training performed using the original ML training container.

In some embodiments, not shown, while the user desires to modify a machine learning model being trained, the user, via the user device 102, does not provide a location of a modified container image because the user does not want to initialize a new ML training container. Rather, the user desires to modify the existing ML training container at runtime so that the machine learning model can be modified without re-starting the training process. Thus, the user device 102 instead provides code that the virtual machine instance 122 adds to the existing ML training container (or uses to replace other code already existing in the ML training container). For example, the original container image used to form the existing ML training container can include executable instructions that are constructed such that the executable instructions retrieve and execute additional code when executed. Such additional code can be provided by the user device 102 in conjunction with the container image (e.g., when the ML training container is initialized) and/or after the virtual machine instance 122 has already begun to execute code stored within the ML training container. In this embodiment, the container image, together with the additional code, form a complete ML training container.

Example Block Diagram for Parallelized Machine Learning Model Training

FIG. 4 is a block diagram of the operating environment 100 of FIG. 1 illustrating the operations performed by the components of the operating environment 100 to parallelize the machine learning model training process, according to some embodiments. As illustrated in FIG. 4, user device 102 transmits a container image location and a training data location to the frontend 129 at (1). In response, the frontend 129 initializes a first virtual machine instance 122A and a second virtual machine instance 122B such that the first virtual machine instance 122A can perform a partial training of a machine learning model using a first portion of the training data and the second virtual machine instance 122B can perform a partial training of the machine learning model using a second portion of the training data. The frontend 129 then transmits the container image location and the location of a first portion of the training data to the virtual machine instance 122A at (2A). Before, during, or after transmitting the container image location and the location of the first portion of the training data to the virtual machine instance 122A, the frontend 129 transmits the container image location and the location of a second portion of the training data to the virtual machine instance 122B at (2B). In some embodiments, the container image location and the training data location are transmitted as part of training requests.

In some embodiments, the virtual machine instance 122A retrieves the first portion of the training data from the training data store 160 using the received location at (3A). Before, during, or after the virtual machine instance 122A retrieves the first portion of the training data, the virtual machine instance 122B retrieves the second portion of the training data from the training data store 160 using the

received location at (3B). In some embodiments, not shown, the virtual machine instances 122A-122B retrieve the same training data.

The virtual machine instance 122A then forms an ML training container using a container image retrieved from the indicated location in some embodiments, and executes code stored in the ML training container using the retrieved first portion of the training data at (4A). Before, during, or after the virtual machine instance 122A executes the code, the virtual machine instance 122B forms an ML training container using a container image retrieved from the indicated location and executes code stored in the ML training container using the retrieved second portion of the training data at (4B). Thus, the virtual machine instances 122A-122B each include a copy of the same ML training container.

Executing the code causes the virtual machine instances 122A-122B (e.g., the ML training containers included therein) to generate model data. Thus, the virtual machine instance 122A transmits model data to the training model data store 175 at (5A) and the virtual machine instance 122B transmits model data to the training model data store 175 at (5B). In some embodiments, not shown, the model data generated by each virtual machine instance 122A-122B is packaged into a single model data file (e.g., by the training model data store 175).

In some embodiments, the virtual machine instances 122A-122B communicate with each other during the machine learning model training. For example, the virtual machine instances 122A-122B can share coefficients, weights, training strategies, and/or the like during the training process.

Example Block Diagram for Deploying and Executing a Machine Learning Model

FIG. 5A is a block diagram of the operating environment 100 of FIG. 1 illustrating the operations performed by the components of the operating environment 100 to deploy a trained machine learning model, according to some embodiments. As illustrated in FIG. 5A, user device 102 transmits a machine learning model deployment request to the frontend 149 at (1). The frontend 149 can initialize a virtual machine instance 142 at (2) and transmit the deployment request to the virtual machine instance 142 at (3). The deployment request includes a location of one or more model data files stored in the training model data store 175. In some embodiments, the deployment request includes an endpoint name. In some embodiments, the deployment request does not include an endpoint name.

In some embodiments, the virtual machine instance 142 retrieves model data from the training model data store 175 at (4). For example, the virtual machine instance 142 retrieves the model data corresponding to the location identified in the deployment request. In some embodiments, not shown, the virtual machine instance 142 does not retrieve the model data. Rather, the model data can be embedded in the container image retrieved by the virtual machine instance 142. The virtual machine instance 142 also retrieves a container image from the container data store 170 at (5). The container image can correspond to a container image identified in the deployment request.

The virtual machine instance 142 can initialize an ML scoring container at (6) in some embodiments. For example, the virtual machine instance 142 can form the ML scoring container using the retrieved container image. The virtual machine instance 142 can further store the model data in the ML scoring container (e.g., in a location that is the same as

the location in which the model data is stored in an ML training container **130** when a machine learning model is trained) at (7).

In some embodiments, if the deployment request did not include an endpoint name, the virtual machine instance **142** can transmit an endpoint name to the frontend **149** at (8). The frontend **149** can then forward the endpoint name to the user device **102** at (9). Thus, the user device **102** can use the endpoint name to access the initialized ML scoring container in the future (e.g., to submit a machine learning model execution request).

FIG. 5B is a block diagram of the operating environment **100** of FIG. 1 illustrating the operations performed by the components of the operating environment **100** to execute a trained machine learning model, according to some embodiments. As illustrated in FIG. 5B, user device **102** transmits a machine learning model execution request to the frontend **149** at (1). The frontend **149** then forwards the execution request to the virtual machine instance **142** at (2). In some embodiments, the execution request includes an endpoint name, which the model hosting system **140** uses to route the execution request to the appropriate virtual machine instance **142**.

In some embodiments, the virtual machine instance **142** executes code stored in an ML scoring container initialized in the virtual machine instance **142** using input data included in the execution request to generate an output at (3). In some embodiments, the virtual machine instance **142** stores the output in the model prediction data store **180** at (4). Alternatively or in addition, the virtual machine instance **142** transmits the output to the frontend **149** at (5), and the frontend **149** transmits the output to the user device **102** at (6).

Example Block Diagram for Executing Related Machine Learning Models

FIG. 6 is a block diagram of the operating environment **100** of FIG. 1 illustrating the operations performed by the components of the operating environment **100** to execute related machine learning models, according to some embodiments. As illustrated in FIG. 6, user device **102** transmits a machine learning model execution request to the frontend **149** at (1). The frontend **149** then forwards the execution request to a first ML scoring container **150A** initialized in a virtual machine instance **142** at (2). In some embodiments, the execution request can include a request for an output from a second machine learning model executed by a second ML scoring container **150B** initialized in the virtual machine instance **142**. However, to generate an output, the ML scoring container **150B** needs data from the execution of a first machine learning model executed by the ML scoring container **150A**. Thus, the virtual machine instance **142** initially routes the execution request to the ML scoring container **150A**. In some embodiments, the ML scoring container **150A** servers as a master container, managing communications to and from other ML scoring containers (e.g., ML scoring container **150B**).

In some embodiments, virtual machine instance **142** causes the ML scoring container **150A** to execute first code to generate a first output at (3). For example, execution of the first code represents the execution of a first machine learning model using input data included in the execution request. The ML scoring container **150A** then transmits the first output to the ML scoring container **150B** at (4).

The virtual machine instance **142** then causes the second ML scoring container **150B** to execute second code using the

first output to generate a second output at (5). For example, execution of the second code represents the execution of a second machine learning model using the first output as an input to the second machine learning model. The second ML scoring container **150B** then transmits the second output to the first ML scoring container **150A** at (6).

In some embodiments, the virtual machine instance **142** pulls the second output from the first ML scoring container **150A** and stores the second output in the model prediction data store **180** at (7). Alternatively or in addition, the virtual machine instance **142** pulls the second output from the first ML scoring container **150A** and transmits the second output to the frontend **149** at (8). The frontend **149** then transmits the second output to the user device **102** at (9).

In some embodiments, not shown, the ML scoring containers **150A-150B** are initialized in different virtual machine instances **142**. Thus, the transmissions of the first output and the second output can occur between virtual machine instances **142**.

Example Machine Learning Model Accuracy Improvement Routine

FIG. 7 is a flow diagram depicting a machine learning model training routine **700** (e.g., a method) illustratively implemented by a model training system, according to some embodiments. As an example, the model training system **120** of FIG. 1 can be configured to execute the machine learning model training routine **700**. The machine learning model training routine **700** begins at block **702**.

At block **704**, in some embodiments, a container image location and a training data location are received. For example, the container image location and the training data location are received as part of a training request.

At block **706**, in some embodiments, a virtual machine instance is initialized. For example, the initialized virtual machine instance is the instance that will perform the machine learning model training.

At block **708**, in some embodiments, the container image and training data are retrieved. For example, the container image can be retrieved from the container data store **170** and the training data can be retrieved from the training data store **160**.

At block **710**, in some embodiments, an ML training container is initialized in the virtual machine instance. For example, the ML training container is formed using the received container image. The container image includes executable instructions that define an algorithm. Thus, the ML training container includes code that includes executable instructions that define an algorithm.

At block **712**, in some embodiments, code stored in the ML training container is executed using the retrieved training data. For example, the retrieved training data (e.g., input data in the training data) is supplied as inputs to the executable instructions that define the algorithm (e.g., using as values for input parameters of the executable instructions).

At block **714**, in some embodiments, model data generated as a result of execution of the code is stored. For example, the model data is stored in the training model data store **175**. Model data can be periodically generated during the machine learning model training process.

At block **716**, in some embodiments, a determination is made as to whether a new container image is received during the machine learning model training process. If a new container image is received, the machine learning model training routine **700** proceeds to block **718**. Otherwise, if no

new container image is received during the machine learning model training process, the machine learning model training routine 700 proceeds to block 720 and ends.

At block 718, in some embodiments, the original ML training container is replaced with a new ML training container. For example, the new ML training container is formed using the new container image. Once the original ML training container is replaced, the machine learning model training routine 700 proceeds back to block 712 such that code stored in the new ML training container is executed using the training data.

In some embodiments, not shown, a new container image is not received. However, one or more new hyperparameters (e.g., a change to the number of clusters, a change to the number of layers, etc.), new code, and/or the like is received. The model training system 120 can modify the original ML training container during runtime (instead of replacing the original ML training container with a new ML training container) to train the machine learning model using the new hyperparameter(s), using the new code, and/or the like.

Example Architecture of Model Training and Hosting Systems

FIG. 8 depicts some embodiments of an architecture of an illustrative model system 800, such as the model training system 120 and the model hosting system 140, that train and/or host machine learning models in accordance with the present application. The general architecture of the model system depicted in FIG. 8 includes an arrangement of computer hardware and software components that can be used to implement aspects of the present disclosure. As illustrated, the model system 800 includes a processing unit 804, a network interface 806, a computer-readable medium drive 807, an input/output device interface 820, all of which may communicate with one another by way of a communication bus.

In some embodiments, the network interface 806 provides connectivity to one or more networks or computing systems, such as the network 110 of FIG. 1. The processing unit 804 can thus receive information and instructions from other computing systems or services via a network. The processing unit 804 can also communicate to and from memory 810 and further provide output information. In some embodiments, the model system 800 includes more (or fewer) components than those shown in FIG. 8.

In some embodiments, the memory 810 includes computer program instructions that the processing unit 804 executes in order to implement one or more embodiments. The memory 810 generally includes RAM, ROM, or other persistent or non-transitory memory. The memory 810 can store an operating system 814 that provides computer program instructions for use by the processing unit 804 in the general administration and operation of the functionality implemented by the model training system 120 and/or the model hosting system 140. The memory 810 can further include computer program instructions and other information for implementing aspects of the present disclosure. For example, in some embodiments, the memory 810 includes a model training component 816 that corresponds to functionality provided by the model training system 120 illustrated in FIG. 1. In some embodiments, the memory 810 includes a model execution component 818 that corresponds to functionality provided by the model hosting system 140.

Example Architecture of an End User Device

FIG. 9 depicts some embodiments of an architecture of an illustrative end user device 102 that can receive data, prepare

data, transmit training requests to the model training system 120, and transmit deployment and/or execution requests to the model hosting system 140 in accordance with the present application. The general architecture of the end user device 102 depicted in FIG. 9 includes an arrangement of computer hardware and software components that can be used to implement and access aspects of the present disclosure. As illustrated, the end user device 102 includes a processing unit 904, a network interface 906, a computer readable medium drive 907, an input/output device interface 920, an optional display 930, and an input device 940, all of which may communicate with one another by way of a communication bus.

In some embodiments, the network interface 906 provides connectivity to one or more networks or computing systems, such as the network 110 of FIG. 1. The processing unit 904 can thus receive information and instructions from other computing systems or services via a network. The processing unit 904 can also communicate to and from memory 910 and further provide output information for the optional display 930 via the input/output device interface 920. The input/output device interface 920 can also accept input from the optional input device 940, such as a keyboard, mouse, digital pen, touchscreen, etc. In some embodiments, the end user devices 102 include more (or fewer) components than those shown in FIG. 9.

In some embodiments, the memory 910 includes computer program instructions that the processing unit 904 executes in order to receive data, prepare data, and transmit the requests described herein. The memory 910 generally includes RAM, ROM, or other persistent or non-transitory memory. The memory 910 can store an operating system 914 that provides computer program instructions and interface software 912 for use by the processing unit 904 in the general administration and operation of the end user device 102. The memory 910 can further include computer program instructions and other information for implementing aspects of the present disclosure. For example, in some embodiments, the memory 910 includes a network application 916, such as browser application, media player, CLI, stand-alone application, etc., for accessing content and communicating with the model training system 120 and/or the model hosting system 140.

FIG. 10 illustrates an example provider network (or “service provider system”) environment according to some embodiments. A provider network 1000 may provide resource virtualization to customers via one or more virtualization services 1010 that allow customers to purchase, rent, or otherwise obtain instances 1012 of virtualized resources, including but not limited to computation and storage resources, implemented on devices within the provider network or networks in one or more data centers. Local Internet Protocol (IP) addresses 1016 may be associated with the resource instances 1012; the local IP addresses are the internal network addresses of the resource instances 1012 on the provider network 1000. In some embodiments, the provider network 1000 may also provide public IP addresses 1014 and/or public IP address ranges (e.g., Internet Protocol version 4 (IPv4) or Internet Protocol version 6 (IPv6) addresses) that customers may obtain from the provider 1000.

Conventionally, the provider network 1000, via the virtualization services 1010, may allow a customer of the service provider (e.g., a customer that operates one or more client networks 1050A-1050C including one or more customer device(s) 1052) to dynamically associate at least some public IP addresses 1014 assigned or allocated to the cus-

customer with particular resource instances **1012** assigned to the customer. The provider network **1000** may also allow the customer to remap a public IP address **1014**, previously mapped to one virtualized computing resource instance **1012** allocated to the customer, to another virtualized computing resource instance **1012** that is also allocated to the customer. Using the virtualized computing resource instances **1012** and public IP addresses **1014** provided by the service provider, a customer of the service provider such as the operator of customer network(s) **1050A-1050C** may, for example, implement customer-specific applications and present the customer's applications on an intermediate network **1040**, such as the Internet. Other network entities **1020** on the intermediate network **1040** may then generate traffic to a destination public IP address **1014** published by the customer network(s) **1050A-1050C**; the traffic is routed to the service provider data center, and at the data center is routed, via a network substrate, to the local IP address **1016** of the virtualized computing resource instance **1012** currently mapped to the destination public IP address **1014**. Similarly, response traffic from the virtualized computing resource instance **1012** may be routed via the network substrate back onto the intermediate network **1040** to the source entity **1020**.

Local IP addresses, as used herein, refer to the internal or "private" network addresses, for example, of resource instances in a provider network. Local IP addresses can be within address blocks reserved by Internet Engineering Task Force (IETF) Request for Comments (RFC) 1918 and/or of an address format specified by IETF RFC 4193, and may be mutable within the provider network. Network traffic originating outside the provider network is not directly routed to local IP addresses; instead, the traffic uses public IP addresses that are mapped to the local IP addresses of the resource instances. The provider network may include networking devices or appliances that provide network address translation (NAT) or similar functionality to perform the mapping from public IP addresses to local IP addresses and vice versa.

Public IP addresses are Internet mutable network addresses that are assigned to resource instances, either by the service provider or by the customer. Traffic routed to a public IP address is translated, for example via 1:1 NAT, and forwarded to the respective local IP address of a resource instance.

Some public IP addresses may be assigned by the provider network infrastructure to particular resource instances; these public IP addresses may be referred to as standard public IP addresses, or simply standard IP addresses. In some embodiments, the mapping of a standard IP address to a local IP address of a resource instance is the default launch configuration for all resource instance types.

At least some public IP addresses may be allocated to or obtained by customers of the provider network **1000**; a customer may then assign their allocated public IP addresses to particular resource instances allocated to the customer. These public IP addresses may be referred to as customer public IP addresses, or simply customer IP addresses. Instead of being assigned by the provider network **1000** to resource instances as in the case of standard IP addresses, customer IP addresses may be assigned to resource instances by the customers, for example via an API provided by the service provider. Unlike standard IP addresses, customer IP addresses are allocated to customer accounts and can be remapped to other resource instances by the respective customers as necessary or desired. A customer IP address is associated with a customer's account, not a particular

resource instance, and the customer controls that IP address until the customer chooses to release it. Unlike conventional static IP addresses, customer IP addresses allow the customer to mask resource instance or availability zone failures by remapping the customer's public IP addresses to any resource instance associated with the customer's account. The customer IP addresses, for example, enable a customer to engineer around problems with the customer's resource instances or software by remapping customer IP addresses to replacement resource instances.

FIG. **11** is a block diagram of an example provider network that provides a storage virtualization service and a hardware virtualization service to customers, according to some embodiments. Hardware virtualization service **1120** provides multiple computation resources **1124** (e.g., VMs) to customers. The computation resources **1124** may, for example, be rented or leased to customers of the provider network **1100** (e.g., to a customer that implements customer network **1150**). Each computation resource **1124** may be provided with one or more local IP addresses. Provider network **1100** may be configured to route packets from the local IP addresses of the computation resources **1124** to public Internet destinations, and from public Internet sources to the local IP addresses of computation resources **1124**.

Provider network **1100** may provide a customer network **1150**, for example coupled to intermediate network **1140** via local network **1156**, the ability to implement virtual computing systems **1192** via hardware virtualization service **1120** coupled to intermediate network **1140** and to provider network **1100**. In some embodiments, hardware virtualization service **1120** may provide one or more APIs **1102**, for example a web services interface, via which a customer network **1150** may access functionality provided by the hardware virtualization service **1120**, for example via a console **1194** (e.g., a web-based application, standalone application, mobile application, etc.). In some embodiments, at the provider network **1100**, each virtual computing system **1192** at customer network **1150** may correspond to a computation resource **1124** that is leased, rented, or otherwise provided to customer network **1150**.

From an instance of a virtual computing system **1192** and/or another customer device **1190** (e.g., via console **1194**), the customer may access the functionality of storage virtualization service **1110**, for example via one or more APIs **1102**, to access data from and store data to storage resources **1118A-1118N** of a virtual data store **1116** provided by the provider network **1100**. In some embodiments, a virtualized data store gateway (not shown) may be provided at the customer network **1150** that may locally cache at least some data, for example frequently accessed or critical data, and that may communicate with virtualized data store service **1110** via one or more communications channels to upload new or modified data from a local cache so that the primary store of data (virtualized data store **1116**) is maintained. In some embodiments, a user, via a virtual computing system **1192** and/or on another customer device **1190**, may mount and access virtual data store **1116** volumes, which appear to the user as local virtualized storage **1198**.

While not shown in FIG. **11**, the virtualization service(s) may also be accessed from resource instances within the provider network **1100** via API(s) **1102**. For example, a customer, appliance service provider, or other entity may access a virtualization service from within a respective virtual network on the provider network **1100** via an API **1102** to request allocation of one or more resource instances within the virtual network or within another virtual network.

As discussed, different approaches can be implemented in various environments in accordance with the described embodiments. For example, FIG. 12 illustrates an example of an environment 1200 for implementing aspects in accordance with various embodiments. For example, in some embodiments requests and responses are HyperText Transfer Protocol (HTTP) messages that are received/sent by a web server (e.g., web server 1206), and the users, via electronic devices, may interact with the provider network via a web portal provided via the web server 1206 and application server 1208. As will be appreciated, although a web-based environment is used for purposes of explanation, different environments may be used, as appropriate, to implement various embodiments. The system includes an electronic client device 1202, which may also be referred to as a client device and can be any appropriate device operable to send and receive requests, messages or information over an appropriate network 1204 and convey information back to a user of the device 1202. Examples of such client devices include personal computers (PCs), cell phones, handheld messaging devices, laptop computers, set-top boxes, personal data assistants, electronic book readers, wearable electronic devices (e.g., glasses, wristbands, monitors), and the like. The one or more networks 1204 can include any appropriate network, including an intranet, the Internet, a cellular network, a local area network, or any other such network or combination thereof. Components used for such a system can depend at least in part upon the type of network and/or environment selected. Protocols and components for communicating via such a network are well known and will not be discussed herein in detail. Communication over the network can be enabled via wired or wireless connections and combinations thereof. In this example, the network 1204 includes the Internet, as the environment includes a web server 1206 for receiving requests and serving content in response thereto, although for other networks an alternative device serving a similar purpose could be used, as would be apparent to one of ordinary skill in the art.

The illustrative environment includes at least one application server 1208 and a data store 1210. It should be understood that there can be several application servers, layers, or other elements, processes or components, which may be chained or otherwise configured, which can interact to perform tasks such as obtaining data from an appropriate data store. As used herein the term "data store" refers to any device or combination of devices capable of storing, accessing and retrieving data, which may include any combination and number of data servers, databases, data storage devices and data storage media, in any standard, distributed or clustered environment. The application server 1208 can include any appropriate hardware and software for integrating with the data store 1210 as needed to execute aspects of one or more applications for the client device 1202 and handling a majority of the data access and business logic for an application. The application server 1208 provides access control services in cooperation with the data store 1210 and is able to generate content such as text, graphics, audio, video, etc., to be transferred to the client device 1202, which may be served to the user by the web server in the form of HyperText Markup Language (HTML), Extensible Markup Language (XML), JavaScript Object Notation (JSON), or another appropriate unstructured or structured language in this example. The handling of all requests and responses, as well as the delivery of content between the client device 1202 and the application server 1208, can be handled by the web server 1206. It should be understood that the web server 1206 and application server 1208 are not required and are

merely example components, as structured code discussed herein can be executed on any appropriate device or host machine as discussed elsewhere herein.

The data store 1210 can include several separate data tables, databases, or other data storage mechanisms and media for storing data relating to a particular aspect. For example, the data store illustrated includes mechanisms for storing production data 1212 and user information 1216, which can be used to serve content for the production side. The data store 1210 also is shown to include a mechanism for storing log or session data 1214. It should be understood that there can be many other aspects that may need to be stored in the data store, such as page image information and access rights information, which can be stored in any of the above listed mechanisms as appropriate or in additional mechanisms in the data store 1210. The data store 1210 is operable, through logic associated therewith, to receive instructions from the application server 1208 and obtain, update, or otherwise process data in response thereto. In one example, a user might submit a search request for a certain type of item. In this case, the data store 1210 might access the user information 1216 to verify the identity of the user and can access a production data 1212 to obtain information about items of that type. The information can then be returned to the user, such as in a listing of results on a web page that the user is able to view via a browser on the user device 1202. Information for a particular item of interest can be viewed in a dedicated page or window of the browser.

The web server 1206, application server 1208, and/or data store 1210 may be implemented by one or more electronic devices 1220, which can also be referred to as electronic server devices or server end stations, and may or may not be located in different geographic locations. Each of the one or more electronic devices 1220 may include an operating system that provides executable program instructions for the general administration and operation of that device and typically will include computer-readable medium storing instructions that, when executed by a processor of the device, allow the device to perform its intended functions. Suitable implementations for the operating system and general functionality of the devices are known or commercially available and are readily implemented by persons having ordinary skill in the art, particularly in light of the disclosure herein.

The environment in one embodiment is a distributed computing environment utilizing several computer systems and components that are interconnected via communication links, using one or more computer networks or direct connections. However, it will be appreciated by those of ordinary skill in the art that such a system could operate equally well in a system having fewer or a greater number of components than are illustrated in FIG. 12. Thus, the depiction of the environment 1200 in FIG. 12 should be taken as being illustrative in nature and not limiting to the scope of the disclosure.

Various embodiments discussed or suggested herein can be implemented in a wide variety of operating environments, which in some cases can include one or more user computers, computing devices, or processing devices which can be used to operate any of a number of applications. User or client devices can include any of a number of general purpose personal computers, such as desktop or laptop computers running a standard operating system, as well as cellular, wireless, and handheld devices running mobile software and capable of supporting a number of networking and messaging protocols. Such a system also can include a number of workstations running any of a variety of com-

mercially-available operating systems and other known applications for purposes such as development and database management. These devices also can include other electronic devices, such as dummy terminals, thin-clients, gaming systems, and/or other devices capable of communicating via a network.

Most embodiments utilize at least one network that would be familiar to those skilled in the art for supporting communications using any of a variety of commercially-available protocols, such as Transmission Control Protocol/Internet Protocol (TCP/IP), File Transfer Protocol (FTP), Universal Plug and Play (UPnP), Network File System (NFS), Common Internet File System (CIFS), Extensible Messaging and Presence Protocol (XMPP), AppleTalk, etc. The network(s) can include, for example, a local area network (LAN), a wide-area network (WAN), a virtual private network (VPN), the Internet, an intranet, an extranet, a public switched telephone network (PSTN), an infrared network, a wireless network, and any combination thereof.

In embodiments utilizing a web server, the web server can run any of a variety of server or mid-tier applications, including HTTP servers, File Transfer Protocol (FTP) servers, Common Gateway Interface (CGI) servers, data servers, Java servers, business application servers, etc. The server(s) also may be capable of executing programs or scripts in response requests from user devices, such as by executing one or more Web applications that may be implemented as one or more scripts or programs written in any programming language, such as Java®, C, C # or C++, or any scripting language, such as Perl, Python, PHP, or TCL, as well as combinations thereof. The server(s) may also include database servers, including without limitation those commercially available from Oracle®, Microsoft®, Sybase®, IBM®, etc. The database servers may be relational or non-relational (e.g., “NoSQL”), distributed or non-distributed, etc.

The environment can include a variety of data stores and other memory and storage media as discussed above. These can reside in a variety of locations, such as on a storage medium local to (and/or resident in) one or more of the computers or remote from any or all of the computers across the network. In a particular set of embodiments, the information may reside in a storage-area network (SAN) familiar to those skilled in the art. Similarly, any necessary files for performing the functions attributed to the computers, servers, or other network devices may be stored locally and/or remotely, as appropriate. Where a system includes computerized devices, each such device can include hardware elements that may be electrically coupled via a bus, the elements including, for example, at least one central processing unit (CPU), at least one input device (e.g., a mouse, keyboard, controller, touch screen, or keypad), and/or at least one output device (e.g., a display device, printer, or speaker). Such a system may also include one or more storage devices, such as disk drives, optical storage devices, and solid-state storage devices such as random-access memory (RAM) or read-only memory (ROM), as well as removable media devices, memory cards, flash cards, etc.

Such devices also can include a computer-readable storage media reader, a communications device (e.g., a modem, a network card (wireless or wired), an infrared communication device, etc.), and working memory as described above. The computer-readable storage media reader can be connected with, or configured to receive, a computer-readable storage medium, representing remote, local, fixed, and/or removable storage devices as well as storage media for temporarily and/or more permanently containing, stor-

ing, transmitting, and retrieving computer-readable information. The system and various devices also typically will include a number of software applications, modules, services, or other elements located within at least one working memory device, including an operating system and application programs, such as a client application or web browser. It should be appreciated that alternate embodiments may have numerous variations from that described above. For example, customized hardware might also be used and/or particular elements might be implemented in hardware, software (including portable software, such as applets), or both. Further, connection to other computing devices such as network input/output devices may be employed.

Storage media and computer readable media for containing code, or portions of code, can include any appropriate media known or used in the art, including storage media and communication media, such as but not limited to volatile and non-volatile, removable and non-removable media implemented in any method or technology for storage and/or transmission of information such as computer readable instructions, data structures, program modules, or other data, including RAM, ROM, Electrically Erasable Programmable Read-Only Memory (EEPROM), flash memory or other memory technology, Compact Disc-Read Only Memory (CD-ROM), Digital Versatile Disk (DVD) or other optical storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other medium which can be used to store the desired information and which can be accessed by a system device. Based on the disclosure and teachings provided herein, a person of ordinary skill in the art will appreciate other ways and/or methods to implement the various embodiments.

In the preceding description, various embodiments are described. For purposes of explanation, specific configurations and details are set forth in order to provide a thorough understanding of the embodiments. However, it will also be apparent to one skilled in the art that the embodiments may be practiced without the specific details. Furthermore, well-known features may be omitted or simplified in order not to obscure the embodiment being described.

Bracketed text and blocks with dashed borders (e.g., large dashes, small dashes, dot-dash, and dots) are used herein to illustrate optional operations that add additional features to some embodiments. However, such notation should not be taken to mean that these are the only options or optional operations, and/or that blocks with solid borders are not optional in certain embodiments.

Reference numerals with suffix letters may be used to indicate that there can be one or multiple instances of the referenced entity in various embodiments, and when there are multiple instances, each does not need to be identical but may instead share some general traits or act in common ways. Further, the particular suffixes used are not meant to imply that a particular amount of the entity exists unless specifically indicated to the contrary. Thus, two entities using the same or different suffix letters may or may not have the same number of instances in various embodiments.

References to “one embodiment,” “an embodiment,” “an example embodiment,” etc., indicate that the embodiment described may include a particular feature, structure, or characteristic, but every embodiment may not necessarily include the particular feature, structure, or characteristic. Moreover, such phrases are not necessarily referring to the same embodiment. Further, when a particular feature, structure, or characteristic is described in connection with an embodiment, it is submitted that it is within the knowledge of one skilled in the art to affect such feature, structure, or

31

characteristic in connection with other embodiments whether or not explicitly described.

Moreover, in the various embodiments described above, unless specifically noted otherwise, disjunctive language such as the phrase “at least one of A, B, or C” is intended to be understood to mean either A, B, or C, or any combination thereof (e.g., A, B, and/or C). As such, disjunctive language is not intended to, nor should it be understood to, imply that a given embodiment requires at least one of A, at least one of B, or at least one of C to each be present.

The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense. It will, however, be evident that various modifications and changes may be made thereunto without departing from the broader spirit and scope of the disclosure as set forth in the claims.

What is claimed is:

1. A computer-implemented method comprising:

receiving, at a service provider network, a first request to train a machine learning (ML) model; wherein the first request to train identifies a first ML training container image; wherein the first request to train identifies a set of training data; wherein the service provider network is implemented by one or more electronic devices;

in response to receiving the first request to train:

retrieving the first ML training container image from a container data store,

using the first ML training container image to initialize a first ML training container on a virtual machine instance, the first ML training container image comprising a first training algorithm code, and executing the first training algorithm code and using the set of training data to train the ML model in the first ML training container to yield a first trained ML model;

evaluating the first trained ML model to obtain a first set of output data;

determining a first quality metric based on comparing the first set of output data to a set of evaluation data;

receiving, at the service provider network, a second request to train the machine learning (ML) model; wherein the second request to train identifies a second ML training container image; wherein the second request to train identifies the set of training data;

in response to receiving the second request to train: retrieving the second ML training container image from a container data store,

using the second ML training container image to initialize a second ML training container on the virtual machine instance, the second ML training container image comprising a second training algorithm, executing the second training algorithm code and using the set of training data to train the ML model in the second ML training container to yield a second trained ML model, and

storing the second trained ML model in a training model data store;

evaluating the second trained ML model to obtain a second set of output data;

determining a second quality metric based on comparing the second set of output data to the set of evaluation data;

receiving, at the service provider network, a request to deploy the second trained ML model, wherein the request to deploy identifies a ML scoring container image, wherein the request to deploy identifies the second trained ML model;

32

in response to receiving the request to deploy:

retrieving the ML scoring container image from a container data store,

using the ML scoring container image to initialize an ML scoring container, the ML scoring container image comprising a scoring algorithm code,

retrieving the second trained ML model from the training model data store,

storing the second trained ML model in the ML scoring container, and

returning an endpoint name for the ML scoring container;

receiving, at the service provider network, a request to perform scoring, the request to perform scoring comprising the endpoint name, the request to perform scoring identifying input data;

and

in response to receiving the request to perform scoring:

executing the scoring algorithm code and using the second trained ML model on the input data in the ML scoring container to yield a result, and

returning the result.

2. The computer-implemented method of claim 1, wherein the set of training data is provided to the first ML training container as one or more files in a first local directory in the first ML training container or as one or more input streams accessible within the first ML training container, and wherein the method further comprises:

storing a set of one or more model artifacts at a storage location, wherein the storing comprises:

obtaining the set of one or more model artifacts from a second local directory in the first ML training container; and

sending the set of one or more model artifacts or an archived version of the set of one or more model artifacts to the storage location.

3. The computer-implemented method of claim 1, wherein a front end of the service provider network is to receive the request to perform scoring and return the result using HyperText Transfer Protocol (HTTP) messages.

4. A computer-implemented method comprising:

receiving, at a service provider network, a request to train a machine learning (ML) model; wherein the request to train identifies a set of training data; wherein the request to train identifies a ML training container image; wherein the service provider network is implemented by one or more electronic devices;

in response to receiving the request to train:

retrieving the ML training container image from a container data store,

using the ML training container image to initialize a ML training container,

executing training algorithm code and using the set of training data to train the ML model in the ML training container to yield a trained ML model, and storing the trained ML model in a training model data store;

receiving, at the service provider network, a request to deploy the trained ML model wherein the request to deploy identifies a ML scoring container image, wherein the request to deploy identifies the trained ML model;

in response to receiving the request to deploy:

retrieving the ML scoring container image from a container data store,

using the ML scoring container image to initialize an ML scoring container, the ML scoring container image comprises a scoring algorithm code, retrieving the trained ML model from the training model data store, storing the trained ML model in the ML scoring container, and returning an endpoint name for the ML scoring container;

receiving, at the service provider network, a request to perform scoring, the request to perform scoring comprising the endpoint name, the request to perform scoring identifying input data;

and

in response to receiving the request to perform scoring: executing scoring algorithm code and using the trained ML model on the input data in the ML scoring container to yield a result, and returning the result.

5. The computer-implemented method of claim 4, wherein the request to deploy identifies a location of the ML scoring container image.

6. The computer-implemented method of claim 4, wherein the ML scoring container further includes a runtime.

7. The computer-implemented method of claim 4, wherein training the ML model based on the set of training data is based on obtaining the set of training data from a storage service in the service provider network.

8. The computer-implemented method of claim 4, wherein training the ML model further comprises: providing the set of training data to the ML training container as one or more files in a local directory in the ML training container or as one or more input streams accessible within the ML training container.

9. The computer-implemented method of claim 4, wherein the request to train the ML model identifies one or more hyperparameters to be used for training the ML model, and wherein the training the ML model further comprises providing the one or more hyperparameters to the ML training container as one or more files in a local directory in the ML training container.

10. The computer-implemented method of claim 4, further comprising: performing scoring using a graphical processing unit.

11. The computer-implemented method of claim 4, wherein a front end of the service provider network is to receive the request to perform scoring and provide the result.

12. The computer-implemented method of claim 4, wherein the request to perform scoring and result are transmitted using HyperText Transfer Protocol (HTTP) messages.

13. The computer-implemented method of claim 4, wherein the request to perform scoring comprises an HTTP endpoint name.

14. The computer-implemented method of claim 4, wherein the result is stored in a model prediction data store.

15. A system comprising:
 a first one or more electronic devices to implement a storage service in a service provider network; and
 a second one or more electronic devices to implement a machine learning service in the service provider net-

work, the machine learning service including instructions that upon execution cause the machine learning service to:

receive a request to train a machine learning (ML) model; wherein the request to train identifies a ML training container image; wherein the request to train identifies a set of training data stored at the storage service;

responsive to receiving the request to train:
 retrieve the ML training container image from a container data store,
 use the ML training container image to initialize a ML training container,
 execute training algorithm code and using the set of training data to train the ML model in the ML training container to yield a trained ML model, and store the trained ML model in a training model data store;

receive a request to deploy the trained ML model, wherein the request to deploy identifies a ML scoring container image, wherein the request to deploy identifies the trained ML model;

response to receiving the request to deploy:
 retrieve the ML scoring container image from a container data store,
 initialize an ML scoring container using the ML scoring container image, the ML scoring container image comprises a scoring algorithm code,
 retrieving the trained ML model from the training model data store,
 storing the trained ML model in the ML scoring container, and
 returning an endpoint name for the ML scoring container;

receive a request to perform scoring, the request to perform scoring comprising the endpoint name, the request to perform scoring identifying input data;

in response to receiving the request to perform scoring:
 executing the scoring algorithm code and using the trained ML model on the input data in the ML scoring container to yield a result, and
 return the result.

16. The system of claim 15, wherein the request to deploy identifies a location of the ML scoring container image comprising the scoring algorithm code.

17. The system of claim 15, wherein the request to train the ML model comprises an identifier of the set of training data usable to obtain the set of training data from the storage service.

18. The system of claim 15, wherein the machine learning service is further to provide the set of training data to the ML training container as one or more files in a local directory in the ML training container or as one or more input streams accessible within the ML training container.

19. The system of claim 15, wherein the request to train the ML model is to identify one or more hyperparameters to be used for training the ML model, and wherein the machine learning service is to provide the one or more hyperparameters to the ML training container as one or more files in a local directory in the ML training container.

20. The system of claim 15, wherein the request to perform scoring and result are to be transmitted using HyperText Transfer Protocol (HTTP) messages.