

**SYSTEMS AND METHODS FOR PROVIDING A PRIVATE MULTI-MODAL
ARTIFICIAL INTELLIGENCE PLATFORM**

PRIORITY CLAIM

[0001] The present application claims priority to U.S. Provisional Application No. 63/020,930
5 (Docket No. 213-0104P), filed on May 6, 2020, the contents of which are incorporated herein
by reference. The present application is a continuation-in-part of U.S. Application No.
16/828,085 (Docket No. 213-0100), filed March 24, 2020, which claims priority to U.S.
Provisional App. No. 62/948,105, filed on December 13, 2019, the contents of which are
incorporated herein by reference. The present application is a continuation-in-part of U.S.
10 Application No. 16/828,216 (Docket No. 213-0101), filed March 24, 2020, which claims
priority to U.S. Provisional App. No. 62/948,105, filed on December 13, 2019, the contents of
which are incorporated herein by reference. The present application is a continuation-in-part
of U.S. Application No. 17/176,530 (213-0102-CON), filed February 16, 2021, which is a
continuation of U.S. Application No. 16/828,354 (213-0102), filed March 24, 2020, now U.S.
15 Patent No. 10,924,460, issued February 16, 2021, which claims priority to U.S. Provisional
App. No. 62/948,105, filed on December 13, 2019, now the contents of which are incorporated
herein by reference. The present application is a continuation-in-part of U.S. Application No.
16/828,420 (Docket No. 213-0103), filed on March 24, 2020, which claims priority to U.S.
Provisional App. No. 62/948,105, filed on December 13, 2019, the contents of which are
20 incorporated herein by reference.

TECHNICAL FIELD

[0001] The present disclosure generally relates to training neural networks and introduces
new techniques for training and deploying neural networks or other trained models in ways
which protect the training data from various sources from being discoverable.

BACKGROUND

[0002] There are existing approaches to training neural networks and that use a federated training approach or a centralized training approach. Each of the existing approaches to training neural networks is based on data received from different clients. The process of sharing of data in this context can result in the data leaking or becoming discoverable.

[0003] For example, deep learning or machine learning in the healthcare context can require large datasets to produce sufficient accuracy in the trained models to be meaningful for diagnostics. Such data can include data on X-rays or MRIs across multiple patients or patient categories. One approach to training such models is to pool raw data and then analyst can access the central repository of the pooled data to conduct machine learning over the large dataset. However, there are ethical issues and privacy issues as well as other issues with this approach such as single-point failure and custodial requirements.

BRIEF DESCRIPTION OF THE DRAWINGS

[0004] In order to describe the manner in which the above-recited and other advantages and features of the disclosure can be obtained, a more particular description of the principles briefly described above will be rendered by reference to specific embodiments thereof which are illustrated in the appended drawings. Understanding that these drawings depict only exemplary embodiments of the disclosure and are not therefore to be considered to be limiting of its scope, the principles herein are described and explained with additional specificity and detail through the use of the accompanying drawings in which:

[0005] FIG. 1 illustrates a federated learning model training approach;

[0006] FIG. 2 illustrates a split learning centralized model training approach;

[0007] FIG. 3 illustrates a split learning peer-to-peer approach;

[0008] FIG. 4 illustrates a federated split leaning approach;

- [0009] FIG. 5 illustrates an embodiment related to blind learning;
- [0010] FIG. 7 illustrates how blind correlation works across multiple clients;
- [0011] FIG. 8 illustrates a method embodiment;
- [0012] FIG. 9 illustrates a method embodiment;
- 5 [0013] FIG. 10 illustrates a method embodiment; and
- [0014] FIG. 11 illustrates a system embodiment.

INTRODUCTION

- [0015] Certain aspects and embodiments of this disclosure are provided below. Some of these aspects and embodiments may be applied independently and some of them may be
- 10 applied in combination as would be apparent to those of skill in the art. In the following description, for the purposes of explanation, specific details are set forth in order to provide a thorough understanding of embodiments of the application. However, it will be apparent that various embodiments may be practiced without these specific details. The figures and description are not intended to be restrictive.
- 15 [0016] The ensuing description provides example embodiments only, and is not intended to limit the scope, applicability, or configuration of the disclosure. Rather, the following description of the exemplary embodiments will provide those skilled in the art with an enabling description for implementing an exemplary embodiment. It should be understood that various changes may be made in the function and arrangement of elements without departing from the
- 20 spirit and scope of the application as set forth in the appended claims.

BRIEF DESCRIPTION

- [0017] What is needed in the art is a method and systems that combines known approaches in order to train neural network models that keeps the data private that the models have been
- 25 trained. The previous approaches can enable the training data to leak out or be discovered as

part of the training process. The improved approach disclosed herein can also address other issues such as custodial issues and remove the single-point of failure problem. This disclosure will first discuss the known approaches and then introduce the new approaches. In one aspect, a particular platform is used to enable a federated development or training of neural network models. The use of the disclosed platform for training models in this manner is disclosed as another embodiment herein. In yet another embodiment, data is encrypted as it is passed between a server and one or more client devices. Various types of federated learning (Shown in FIG. 1), split learning (shown in FIG. 2), and split-learning peer-to-peer (Shown in FIG. 3) are disclosed herein. This disclosure will provide several new improvements over traditional approaches.

[0018] Typical federated learning involves passing a whole model from a server to a client device for training using the client data. The process can include using a number of different clients, each with their respective data, for training purposes. The approach typically is performed in a linear and iterative fashion in which the whole model is sent to the first client with data, then after training at the first client, the whole model is received back to the server for “averaging”. Then whole updated model is sent to second client with data for additional processing. Then that updated model is sent back to the server for additional “averaging”, and so on. In a split learning approach, the model is split and part is sent to each client but there still is a linear and interactive training process that is inefficient. The split-learning peer-to-peer approach also is performed linearly as peer clients share data in the linear process. Improvements in maintaining the privacy of data and efficiency in the training process are needed.

[0019] This disclosure describes two primary improvements over federated learning and split learning. The first is a federated split learning approach (shown in FIGs. 4-5) in which client side processing occurs in parallel and independent of other clients. The second disclosed

approach (shown in FIGs. 6-10) relates to a multi-modal artificial intelligence (MMAI) training approach to handle different types of data from different clients.

[0020] As noted above, a federated split learning approach is disclosed as a variation on the typical federated learning approach above. A method in this regard includes splitting up, at a server, a neural network into a first portion and a second portion, and sending the second portion separately to a first client and a second client. The clients can have the data (MRIs, patient data, banking data for customers, etc.) and each receive a portion of the neutral network (a certain number of layers of the network up to a cut layer). The method includes performing the following operations until a threshold is met: (1) performing, at the first client and the second client, a forward step on the second portion simultaneously to generate data SA1 and SA2 (See FIGs. 1-4); (2) transmitting, from the first client and the second client, SA1 and SA2 to the server; (3) calculating, at the server, a loss value for the first client and the second client; (4) calculating, at the server, an average loss across the first client and the second client; (5) performing, at the server, backpropagation using the average loss and calculating gradients; and (6) sending, from the server, the gradients to the first client and the second client. This approach provides an improvement over the federated learning approach and the split learning approach by causing the processing on the client side (or the “data server” side) to operate in parallel and independent of each other. This approach also differs from the split learning peer-to-peer approach as well. The independent data servers send their activations up to the server side which aggregates, averages or otherwise processes the data depending on the network requirement to obtain the final trained model.

[0021] Another aspect of this disclosure relates to an improvement in developing an artificial intelligence model in which multiple different modes of data or types of data are available to be used for training. For example, different clients might have different types of data. One client might have images of X-rays or MRIs and another client may have text describing a

patient's health condition. In this regard, a method can include splitting a neural network into a first client-side network, a second client-side network and a server-side network, sending the first client-side network to a first client. The first client-side network is configured to process first data from the first client, the first data having a first type. The first client-side network
5 can include at least one first client-side layer. The method includes sending the second client-side network to a second client. The second client-side network is configured to process second data from the second client, the second data having a second type. The second client-side network can include at least one second client-side layer, wherein the first type and the second type have a common association.

10 [0022] The method can further include receiving, at the server-side network, first activations from a training of the first client-side network on first data from the first client, receiving, at the server-side network, second activations from a training of the second client-side network on second data from the second client, training at least one server-side layer of the server-side network based on the first activations and the second activations to generate gradients and
15 transmitting the gradients from the server-side network to the first client-side network and the second client-side network. In this manner, multiple different types of data, having a common relationship such as being related to single patient or a single type or category of patient, are used to train the model.

[0023] This summary is not intended to identify key or essential features of the claimed
20 subject matter, nor is it intended to be used in isolation to determine the scope of the claimed subject matter. The subject matter should be understood by reference to appropriate portions of the entire specification of this patent, any or all drawings, and each claim.

[0024] The foregoing, together with other features and embodiments, will become more apparent upon referring to the following specification, claims, and accompanying drawings.

25

DETAILED DESCRIPTION

[0025] Disclosed herein is a new system, a platform, compute environment, cloud environment, marketplace, or any other characterization of the system that will enable an improved approach to training neural networks. In one aspect, the approach is called a federated-split learning approach that combines features from known approaches but that provides a training process that maintains privacy for data used to train the model from various client devices. This disclosure first discusses in more detail the federated learning approach, follow by the split learning approach and a split learning peer-to-peer approach and then introduces the novel federated split-learning approach. Additionally, the multi-modal artificial intelligent (MMAI) learning approach for different types of data is introduced as well. The novel federated split-learning approach and the MMAI approach build on several models including those mentioned above. The application will review these first approaches in more detail and then introduce the two novel learning techniques.

Federated Learning

[0026] FIG. 1 illustrates the federated learning approach 100. This is an approach used by major companies now. A downside of this approach is that it proceeds “linearly” to one data provider at a time – rather than in parallel. The example neural network shown is a fully connected feed forward neural network that is being trained using a federated learning approach. The training process in this case includes a server 102 creating a model 104 and sharing the model 106A, 108A and 110A with respective clients 106, 108, 110 in a linear fashion. The clients train the respective model 106A, 108A, 110A separately when they receive the model on their turn and respectively send their trained model data back to the server 102 as shown. The server 102 averages the models and produces a new model 104 with updated weights (a.k.a a trained model). The server 102 sends the new model or weights to the

respective clients 106, 108, 110 in a linear fashion. The process is repeated a number of iterations or until a specific accuracy is achieved.

[0027] In each iteration, the server 102 averages all participating models to create a trained model B. Thus, the server has a fully-trained model 104 at any point of time. The term “global model” refers to the model that results from the training process. The global model is a trained object that will be used for an inference task. An inference task might be to evaluate a medical image to classify whether the patient has cancer or a broken bone or some other medical condition.

[0028] An example of this approach being used, devices such as an electronic watch, or a mobile device, a device charging at night for example, and connected to a Wi-Fi network, could have its processor used to train neural network models. Thus, client 1 (106) could be an Apple watch, client 2 (108) could be another person’s iPhone, and so forth. An example of a model is the Siri speech processing service offered by Apple. Every device is training the same model and the only difference is that the respective client is training on the data local to them. The model or data is transmitted back to the server 102 and the server averages the model together. The downside is that respective clients, such as client 1 (106), could be tricked into sharing something about the data being used to train the model. This would be a leakage of privacy data and raise the issued outlined above. The challenge of the federated learning approach is that there is no model privacy as the entire model is passed from client to client. There are high computational costs as each client processes the entire model, and a heavy communication overhead as the entire model is transmitted numerous times. A reconstruction attack can make training data venerable as well.

Split Learning

[0029] FIG. 2 illustrates a split learning centralized approach. A model (neural network) 204 is split into two parts: one part (206A, 208A, 210A) resides on the respective client side 206,

208, 210 and includes the input layer to the model and optionally other layers up to a cut layer, and the other part (B) resides on the server side 202 and often includes the output layer. Split layer (S) refers to the layer (the cut layer) where A and B are split. In FIG. 2, SA represents a split layer or data sent from A to B and SB represents a split layer sent from B to A.

5 [0030] In one example, the neural network between B 204 and client 1 (206) is the B portion 204 plus the A1 portion (206A) with the communication of data SB1 (206C) and SA1 (206B) to complete the entire neural network. The training process is as follows in this model. The server 202 creates A and B and sends a respective model A (206A, 208A, 210A) to the respective client 206, 208, 210. For every client, the operations include repeating the following

10 in a linear or iterative fashion across the group of clients until some conditions occurs. The respective client 206, 208, 210 on their turn downloads the most recent model A from the server 202 (Note that this step is different between the approach shown in FIG. 2 and FIG. 3.). The clients 206, 208, 210 in their respective turn do a forward step on the model A and sends the output of A (i.e., activations at S only or SA1 (206B), SA2(208B), SAN 210B)) to the server

15 202 in addition to the required labels. The server 202 does a forward step on B using the SAs received from the respective client 206, 208, 210. The server 202 calculates the loss function and the server 202 does backpropagation and calculates gradients at the S layer. The server 202 sends the gradients of S only (i.e., SB1 (206C), SB2 (208C), SBN (210C)) to the respective client 206, 208, 210. This process is performed linearly across the different clients such that

20 the operations occur first for client 206, followed by client 208, and then client 210. The client 206, 208, 210 does backpropagation using the SB gradients received from the server 202 and the client 206, 208, 210 shares their updated A (SA1 (206B), SA2 (208B), SAN (210B)) with the server 202.

[0031] The horizontal axis in FIG. 2 is time such that the processing occurs in like a round-robin fashion from client to client.

25

[0032] In one example, network A1 206A on client 1 can include a convolution layer and an activation layer. Having processed data, the client 1 (206) sends the result of that layer forward (SA1 (206B)) to the next layer in the network, which is at the server 202, which calculates the backpropagation and so forth as outlined above. The B network repeatedly (in round robin fashion) processes the different data from the different clients 206, 208, 210. It will ultimately arrive at an averaged reflection of the network. It never trains the network on all the data from all the clients 206, 208, 210 at the same time. It can process data faster and have a benefit of B being averaged across the data as it is built. The final algorithm has not seen all the data. The model B cannot be tricked into revealing its data as it has never been trained on all of the data.

Split Learning in a Peer-to-Peer Environment

[0033] FIG. 3 illustrates a split learning peer-to-peer approach. A model (neural network) is split into two parts: one part (A) resides on the client side and includes the input layer, and the other part (B) resides on the server side and often includes the output layer. In FIG. 3, the client side part (A) is shown respectively as A1 (306A) at client 306, A2 (308A) at client 308, AN (310A) at client 310. A split layer (S) refers to the layer where A and B are split. In FIG. 3, SA represents a split layer sent from A to B and SB represents a split layer sent from B to A.

[0034] In one example, the neural network between B and client 1 306 is the B portion plus the A1 portion 306A with the communication of data SB1 306C and SA1 306B to complete the entire neural network. The training process is as follows in this model. The server 302 creates A and B and sends A to the clients 306, 308, 310. For every client, the process includes repeating the following until some conditions occurs. First, the process includes downloading the most recent A from a previous client.

[0035] Note that this step is different between the approach shown in other figures. The process then includes performing a forward step on A and sending the output of A (i.e., activations at S only) to the server 302 in addition to the required labels. The server 302 performs a forward step on B using the SA received from the respective client 306, 308, 310.

5 The server 302 calculates a loss function and performs a backpropagation and calculates gradients at S. The server 302 sends the gradients of S only (i.e., SB) to the respective clients 306, 308, 310. The client does backpropagation using the SB gradients received from the server 302. The client shares their updated A with the server 302.

[0036] The peer-to-peer approach generally involves the respective client updating its A

10 model by directly downloading it from a last trained client, or more broadly, by a previously trained client. In this regard, the process of training clients can occur in a round-robin fashion where the clients are trained sequentially. For example, if client 1 306 gets trained first, then in a peer-to-peer model, rather than client 2 308 updating its client-side model A2 from the server 302 or another trusted server, client 2 308 updates its client model A2 by downloaded

15 the client side model A1 from client 1 306. The previously trained model can be the last trained client model or it could be a model from some other previously trained client based on some criteria. For example, client 1 306 and client 2 308 may have their respective models trained. Client 3 310 needs a client-side model update and might implement an algorithm or process to determine which client-side model to download between client 1 306 and client 2 308. Note

20 that the disclosure below implements a multi-model artificial intelligence training process that could apply here. If client 1 306 processes images and its model A1 focuses on image processing, and client 2 308 processes text and its model A2 focuses on text processing, and client 3 310 processes images, then the algorithm or process could cause, in a peer-to-peer environment, the downloading of the client side model A1 to the client 3 310 as its update.

[0037] In one scenario, there is not enough information from split learning to achieve proper training of the neural network. It is assumed in this model that a good training approach could be that A and B are aggregated at the server 302 in plain text by simply stacking them (A and B).

5 Federated Split Learning

[0038] FIG. 4 illustrates the improvement to training neural networks disclosed herein. This improvement can be characterized as a federated split learning approach and addresses some of the deficiencies of the approaches disclosed above. FIG. 4 introduces a parallel processing approach. The parallel and independent processing causes the model training to occur at a
10 faster pace than the other models described above.

[0039] The federated split learning approach does not perform the round robin processing described above. The server 402 splits the network at the “split layer” which is a user parameter inserted into the network definition codes. The “top portion” of the network is kept at the server 402 the “bottom portion” is sent to the respective data providers or clients 406,
15 408, 410 (the terms clients and data providers are used interchangeably here). The training starts at the very lowest network layer which is the layer closest to the data. Each layer reads either the data (from the first layer) or the output of the previous layer (all other layers).

[0040] The layers can calculate their output (these are termed “activations” because they come from an activation function) based on any valid network architecture command
20 (convolutions, dropouts, batch normalization, flatten layers, etc.) and activation function (relu, tanh, etc.). When the last layer on the data side 406, 408, 410 has calculated its appropriate activations (i.e., output) those outputs are sent to the first layer on “the other side of the split” – the first layer on the server side 402.

[0041] The following approach involves splitting the model up as before. A model is split
25 into two parts: (A) on the client side and includes the input layer, and (B) on the server side

and often includes the output layer. (S) is the split layer. The clients or data providers 406, 408, 410 run independently and send back the answer if they have it. The code on the server 402 processes the data and sends back its output equally to all the clients as SB (406C, 408C, 410C).

5 [0042] An example training process is as follows. The server 402 creates A and B and sends the portion A (406A, 408A, 410A) to the clients 406, 408, 410. The following steps are repeated until a condition is met (e.g., accuracy). All the clients 406, 408, 410 do the forward step on A simultaneously. Up to this point, all the calculations on the clients 406, 408, 410, are being done on independent servers and there is no dependency from one data server to the
10 other. This approach highlights a one of the innovations disclosed herein. All these calculations by the clients/data providers 406, 408, 410 can all operate in parallel, at the same time. This is in contrast to the linear or “round robin” fashion discussed above.

[0043] The clients 406, 408, 410 each run their portion A (406A, 408A, 410A) of the neural network and generate a respective output of A (i.e., SA (406B, 408B, 410B) and send the output
15 to the server 402. The server 402 receives 3 different ‘versions’ of the activations (one from each of SA1, SA2, SA3). At this point, the server 402 processes those activations “appropriately”, which can mean that the server 402 does different operations depending on the case. For example, the server 402 calculates the loss value for each client 406, 408, 410 and the server 402 calculates the average loss across all clients. The server 402 performs
20 backpropagation using the average loss and calculates gradients at S. The server 402 sends gradients at S (i.e., SB (406C, 408C, 410C)) to all the clients 406, 408, 410.

[0044] In other words, training on the server side 402 proceeds much like is described above. Once the first layer on the server side 402 is “complete” (either through averaging or aggregating what is received from the data providers 406, 408, 410) forward propagation
25 occurs until the “top” of the network is reached. An additional innovation described in this

disclosure is in the management of the activations coming from the data providers 406, 408, 410 and how they get averaged, aggregated or some other treatment. Once the system arrives at the top of the model, the server 402 calculates the gradients necessary for back propagation, and sends them back down and across the split networks as shown in FIG. 4.

5 [0045] As noted above, the processing and the management of the activations by the server 402 can vary depending on different factors. For example, assume a case where all three data providers 406, 408, 410 are supplying the same data (X-rays). In that case, the data will be combined horizontally which can conceptually mean that the data is “stacked” one file on top of the other. In this case, the activations that come up will most likely be averaged. The
10 “average of each activation” will then be sent forward into the “top half” of the network.

[0046] In a different case, the data can be “vertically” stacked, so Client 1 406 has the first 40 columns of data (say a blood test), Client 2 408 has the next 60 columns of data (say an Electronic Health record that includes data such as age, weight, etc.) and Client 3 410 has the last 100 columns of data (say insurance information – previous claims, etc.). In this instance,
15 the three clients can be considered as establishing a combined “record” of 200 columns (aggregated vertically across the page). In this case, the activations will be “combined vertically” and sent forward into the server network. This and other approaches to combining data can be implemented. Note that the multi-model artificial intelligence model described more fully below builds upon the concept just described with respect to combining vertically
20 the activations. More details will be provided below on this concept.

[0047] As noted above, the clients 406, 408, 410 run in parallel in this embodiment. This reduces the time it takes to train the model – as all the processing is done in parallel. Further, this data is delivered over a particular platform. The applications incorporated above provide examples of the particular platform that can be used to deliver the data as disclosed herein.
25 This will be discussed more below.

[0048] A global model in federated-split learning can be aggregated as follows. After the training is done, the system uses on the following approach to aggregate a global model, which will be used for the inference task. In a first approach, the server selects one of the models, A_i , to be aggregated with its model, B, to form the global model. The selection of A_i could be achieved using one of the following ways. For example, random selection could be used where the server selects a model (A_i) of any client 406, 408, 410 randomly. This random selection might be influenced by other factors, such as the currently available clients online, the types of data each client processes (text data, image data, temporal data) or based on the transmission speed or network delay between the two entities. The server then stacks both parts A_i and B to generate the global model.

[0049] In another example, a weighted client selection could be used. For this selection criteria, the server 402 assigns each client a weight (i.e., a numerical value) that reflects their importance based on their data, computational powers, and other valuable assets they possess and contribute during the training process. For example, a particular model set (say data for a certain language, data associated with a type of image, data associated with a patient set, or data from a particular country or region) could get weighted heavily in the model development. Thus, if a country is selected, then the client devices from that country can be weighted more heavily than clients from other countries. Japanese-based client devices can be used for 80% of the model data, for example. Australia could be 10% and Canada could be the other 10%. In another example, data from a certain clinic associated with an outbreak of the flu or COVID could be weighted more heavily. In yet another example, the type of data might be weighted more heavily as well. Image data may be used for 70% of a model, while textual data for 20% and temporal data for 10%.

[0050] Yet another model could be an accuracy-based selection. In this case, the server 402 can test the accuracy generated from each client model A_i and then select the model that

generates the “best” accuracy. The “best” can be identified by stakeholders, through a machine learning approach, or otherwise. These are all models of the first approach.

[0051] A second approach can be where the global model is aggregated by averaging all clients’ models $A_i \in \{1, N\}$. Each client first encrypts their model using homomorphic encryption and then sends the encrypted A_i data to the server 402. The server 402 adds all the encrypted models, decrypts the addition results, and then calculates their average. The averaged A is then stacked with B to generate a global model. One approach could be a default approach, and optional approaches could be provided as well. The decryption processes and averaging process could also be spread between different servers, for example, with one process occurring on the client side and another process being performed by the server 402 to achieve the global model.

[0052] The approaches may vary through the development of the model. For example, the model may begin to be trained using a default approach and then the training could be adjusted such that a weighted approach is used to complete the model training.

[0053] A method example is shown in FIG. 5 and can include splitting up, at a server, a neural network into a first portion and a second portion (502), sending the second portion separately to a first client and a second client (504) and performing the following operations until a threshold is met:

- (1) performing, at the first client and the second client, a forward step on the second portion simultaneously to generate data $SA1$ and $SA2$;
- (2) transmitting, from the first client and the second client, $SA1$ and $SA2$ to the server;
- (3) calculating, at the server, a loss value for the first client and the second client;
- (4) calculating, at the server, an average loss across the first client and the second client;

(5) performing, at the server, backpropagation using the average loss and calculating gradients; and

(6) sending, from the server, the gradients to the first client and the second client (506).

5 [0054] A computing device or devices performing the above operations can also be covered as well as a computer-readable storage device storing instructions which, when executed, cause the processor to perform these operations. The operations can be performed in any order and the method can include one or more of the operations.

[0055] In another aspect of this disclosure, the platforms described in the patent applications
10 incorporated above can provide the basis for communicating data back and forth in any of the federated models. For example, each of the clients and/or the server as well may be required to be logged onto a platform or one of the versions of the platform referenced in the applications incorporated herein. Therefore, delivering this functionality over a platform or an exchange configured as disclosed in these applications is also covered as an aspect of this disclosure.

15 [0056] In another aspect, a customer could chose SA, SB lines (vectors and numbers) which represent weights that need to be propagated. If a client wanted their data to be locked down without the server knowing anything about the data, that data can be homomorphically encrypted. The encryption process (which can include any encryption process) could be used in any approach disclosed above.

20 [0057] The incorporated patent applications above provide example platforms that client devices and/or servers can log into or may be required to be logged into in order to perform the federated-split learning approach disclosed herein.

[0058] It is noted that in one aspect, the steps disclosed herein can be practiced by a “system.” The system can include the server and one or more clients together, or might just be
25 functionality performed by the server. The system could also be a client or a group of clients,

such as clients in a particular geographic area or clients groups in some manner that are performing the client-based functions disclosed herein. In one aspect, the “server” can also be a computing device (physical or virtual) on the server side as well as a computing device (physical or virtual) on the client side. In one example, a server can be on the client side and
5 can receive back-propagation output of the respective client side models A_i and can synchronize a client-side global model in a round of training.

[0059] Thus, each of the server side system and the client side system can perform any one or more of the operations disclosed herein. Claims can be included which outline the steps that occur from the standpoint of any device disclosed herein. For example, the steps of
10 transmission, calculation, and receiving of data can be claimed from the standpoint of a server device, a client device, or group of client devices depending on which embodiment is being covered. All such communication from the standpoint of an individual component or device can be included as within the scope of a particular embodiment focusing on that device.

[0060] In another aspect, the system can include a platform as disclosed in the patent
15 applications incorporated by reference also performing steps in coordination with the concept disclosed above. Therefore, the platform as used to provide the federated-split learning process described herein is also an embodiment of this disclosure and steps can be recited in connection with the use of that platform for training models in a manner that maintains privacy of the data as described herein.

20 [0061] Typically the training of a neural network is performed on similar data types. For example, a neural network trained to identify cancer by receiving a patient image or a kidney is trained on images of kidneys that are and are not cancerous. Next is discussed a new approach to training which uses different types of training data together to train a neural network, using the federated split learning approaches disclosed herein.

25

Multi-Model Artificial Intelligence Approach

[0062] As mentioned above, the MMAI innovation builds on the “vertical aggregation” idea described in an example of federated split learning. The example related to all three clients 406, 408, 410 providing the same type of data – either images (for stacking) or tabular data to be combined vertically. When the inventors were considering the vertical aggregation concept, they realized that this could be done with different types of data. For example, Client 1 could provide images, Client 2 could provide a blood test, and Client 3 could provide doctors textual notes. The significant difference is all of those data types require different network architectures. In this case, the developers of the system can’t define one network and then let the server “split” it. Thus, part of the solution is to let the users define the network “before the split” for each data provider, and then define the network and aggregation technique on the server. This approach is illustrated in FIGs. 6-10.

[0063] FIG. 6 illustrates the multi-modal artificial intelligence (MMAI) platform or a machine learning (ML) platform 600. The MMAI approach reduces the computational requirements and communication overhead of other approaches. Additionally, the training speed is much faster and the process maintains a much higher privacy in the data, including the fact that the model stays private as well.

[0064] The MMAI platform 600 applies AI/ML techniques to multiple data types in one large AI model. Typically, different data types require different AI network architectures to yield accurate results. Images, for example, typically require special filters (convolutions), whereas text or speech require different “time series-like” treatment, and tabular data frequently works best with ML or feed forward architectures. The issue is that images are best understood by looking at all of the pixels together and “convoluting” them in various ways, whereas speech is best understood in the context of what came before and/or after a certain sound (i.e. in a

manner similar to time-series data), etc. Because of these differences in processing, "state of the art" systems today typically process one data type (i.e. images, text, speech, tabular, etc.).

[0065] Most AI researchers recognize that breakthroughs in "next generation" accuracy can be achieved by adding more unique data to their models. This is essentially the equivalent to providing more data to the model to give it more context with which to discover interesting differences in cases. An example of this concept is a model that diagnoses Atrial Fibrillation (A-fib) by examining ECG (electro-cardiogram) data. The model can reach a certain level of accuracy based on the ECG data alone, but when the researchers add age, sex, height and weight to the ECG data, the model becomes far more accurate. The increase in accuracy is due to the four additional data types being able to help the model better understand what would otherwise look to the model like "equivalent" ECGs. Adding the four items or characterizations of the data can make the data more granular.

[0066] The MMAI platform 600 shown in FIG. 6 introduces a new generation cryptography toolset to improve the training and protection of private data. The MMAI platform 600 provides the model with more data than is typically used to train AI/ML models and expands on the data. The approach adds a significant amount of data by combining different data types - i.e. images and tabular data, for instance.

[0067] FIG. 6 illustrates a first outside source of data 602, which is shown as Wells Fargo bank. The Wells Fargo data 602a is encrypted 602b and the package of encrypted data 602c is transmitted to a private AI infrastructure 603. A second outside source of data 604 is shown as Citibank. The Citibank data 604a is encrypted 604b and the package of encrypted data 604c is transmitted to the private AI infrastructure 603. A third outside source of data 606 is shown as from Bank of America. The Bank of America data 606a is encrypted 606b and the package of encrypted data 606c is transmitted to the private AI infrastructure 603. The AI infrastructure 603 includes a first module 608 that will privately explore, select and preprocess all of the data

610 from the disparate sources 602, 604, 606. In this example, all of the sources are identified as banks but they will have different structures for their data, and the respective data can be disparate as well. Of course, it is not a requirement that all of the outside sources 602, 604, 606 of data be of the same type, i.e., banks. The use of banks is just an example. The outside
5 sources 602, 604, 606 could be, for example, a hospital, a clinic, a university, and so forth. The basic concept is that the data types can be different from the various different outside sources 602, 604, 606.

[0068] The private AI infrastructure 603 can include a component that privately explores, selects and preprocesses the relevant features from all of the data 602c, 604c, 606c it receives
10 for training. Feature 612 represents the subset of the data 610 which can result from the processing of the component in the private AI infrastructure 603. In operations 614, 616, the AI infrastructure 603 privately trains new deep and statistical models on the selected data 612 and in operation 618 will predict on any private and sensitive data, which can include images, video, text and/or other data types. The AI infrastructure 603 can then sell or grant access to
15 the new models which is presented in operation 620.

[0069] FIG. 7 illustrates another variation on the split learning technique 700. This approach provides low compute requirements and low communication overhead to improve the training of models by using a blind correlation process for training based on disparate types of data. Building on the A-fib model example above, another source of even more data for the model
20 would be to include a chest X-ray for each case the model considers. Unfortunately, the typical processing of the X-ray image is not consistent with the typical processing of the tabular ECG data. With a few minor engineering additions, the above-disclosed split-federated learning tool can be used to address this incompatibility problem. Namely, new instructions can be provided to the tool to allow different data types to process in the existing pipeline.

[0070] In this case rather than an “automatic” split of the network architecture this variation on the idea allows the network architect (i.e. the data scientist developing the algorithm) to specify the specific network components desired for each data type. Each data type will need network architecture layers relevant to its data type (i.e. convolutional layers for images, Recurrent layers/Long Short Term Memory layers for speech, feed forward layers for tabular data, etc.). These disparate layers, each specific to the data type in question, will be specified such that they run on the “data server” side (almost like independent networks in and of themselves). The last layer of each “independent network” (per data type) will send its activations “across the split” to the “server side”. The algorithm server side will have one consistent “network” that processes the incoming activations (from the data server side) appropriately. In some respects this approach is similar to an “ensemble of networks” (on the data server side) being aggregated into one final network on the algorithm server side (which ultimately produces the final “answer” from the “ensemble” of networks).

[0071] Split learning is a collaborative deep learning technique, where a deep learning network or neural network (NN) can be split into two portions, a client-side network A and a server-side network B, as discussed above. The NN includes weights, bias, and hyperparameters. In FIG. 7, the clients 702, 704, 706, where the data reside, commit only to the client-side portion of the network, and the server 710 commits only to the server-side portion of the network. The client-side and server-side portions collectively form the full network NN.

[0072] The training of the network is done by a sequence of distributed training processes. The forward propagation and the back-propagation can take place as follows. With the raw data, a client (say client 702) trains the client-side network 702A up to a certain layer of the network, which can be called the cut layer or the split layer, and sends the activations of the cut layer to the server 710. The server 710 trains the remaining layers of the NN with the activations that it received from the client 702. This completes a single forward propagation step. A similar

process occurs in parallel for the second client 704 and its client side network 704A and its data and generated activations which are transmitted to the server 710. A further similar process occurs in parallel for the third client 706 and its client side network 706A and its data and generated activations which are transmitted to the server 710.

5 [0073] Next, the server 710 carries out the back-propagation up to the cut layer and sends the gradients of the activations to the respective clients 702, 704, 706. With the gradients, each respective client 702, 704, 706 performs back-propagation on the remaining network 702A, 704A, 706A. This completes a single pass of the back-propagation between a client 702, 704, 706 and the server 710.

10 [0074] This process of forward propagation and back-propagation continues until the network gets trained with all the available clients 702, 704, 706 and reaches its convergence. In split learning, the architectural configurations are assumed to be conducted by a trusted party that has direct access to the main server 710. This authorized party selects the ML model (based on the application) and network splitting (finding the cut layer) at the beginning of the learning.

15 [0075] As noted above, a concept introduced in this disclosure relates to the clients 702, 704, 706 each providing a different type of data but also where the different types of data have a common association. Thus, the selection of the machine learning model can be based on the types of data that are being processed on the client side, and the process of finding the cut layer can also depend on what types of data or the disparity in the different types of data. For
20 example, for widely disparate data types across the clients 702, 704, 706, the cut layer may be chosen to have more or less layers on the client-side networks 702A, 704A, 706A. In another aspect, the number of layers before the cut layer or split layer may vary across clients. Client 702 may be processing images and require 8 layers before the cut layer, while client 704 may process text and only need 4 layers before the cut layer. In this regard, as long as the vectors,
25 activations or activation layer at the cut layer is consistent across the different clients 702, 704,

706 having different types of data, there is no requirement that the number of layers at the client-side networks 702A, 704A, 706A be the same.

[0076] The synchronization of the learning process with multiple clients 702, 704, 706 can be done either in centralized mode or peer-to-peer mode. In the centralized mode, before starting
5 training with the server 710, a client 702, 704, 706 updates its client-side model 702A, 704A, 706A by downloading the model parameters from a trusted third-party server 710, which retains the updated client-side model uploaded by the last trained client. On the other hand, in peer-to-peer mode, the client 702, 704, 706 updates its client-side model by directly downloading it from the last trained client. As noted above, previously-trained models may
10 have a data type similarity to a current client that needs to update its model. For example, the similarity may be based on the data be images, textual data, speech data, video data, temporal data, and so forth. Thus, there may be an intelligent selection of which previously-trained client model to use to download from a peer. The processing by the server 710 can also be split in some cases between some processing on the server side and other processing at a
15 federated server on the client side.

[0077] As introduced above, client one 702, client two 704 and client three 706 could have different data types. The server 710 will create two parts of the network and sends one part 702A, 704A, 706A to all the clients 702, 704, 706. The system repeats certain steps until an accuracy condition or other condition is met, such as all the clients sending data to the part of
20 the network that they have, and sends the output to the server 710. The server 710 calculates the loss value for each client and the average loss across all the clients. The server 710 can update its model using a weighted average of the gradients that it computes during back-propagation and sends the gradients back to all the clients 702, 704, 706. The clients 702, 704, 706 receives the gradients from the server 710 and each client 702, 704, 706 performs the back-
25 propagation on their client-side network 702A, 704A, 706A and computes the respective

gradients for each client-side-network 702A, 704A, 706A. The respective gradients from the client-side networks 702A, 704A, 706A can then be transmitted back to the server 710 which conducts an averaging of the client-side updates and sends the global result back to all the clients 702, 704, 706.

5 [0078] It is noted that the server 710 functionality can be also broken into several servers that each perform the different operations (such as updating its model by one server and averaging the local client updates by another server, each located in different areas). In the case of FIG. 7, the clients 702, 704, 706 all process disparate types of data which normally would or could not be processed to develop an AI model.

10 [0079] For example purposes, the A-fib model from above can be used to illustrate the process. Client one 702 could have ECG data, client two 704 could have X-ray data, and client three 706 could have genetic data. Client one 702, for example, could be a hospital, client two 704 could be a medical diagnostics imaging company and client three 706 could be a bank or financial institution, in a manner depicted in FIG. 6. One of the clients could also have time-
15 based data such as progressive information about the patient relative to weekly visits to the hospital for checkups.

[0080] The approach shown in FIG. 7 illustrates how the system can implement new user instructions that allow a user to bring different data types together with the "correct" processing before the split or cut layer or as shown in the blind decorrelation block 708. Each of those
20 parts of the model can be independent, and will operate independently. In one aspect, the processing performed by the blind correlation block 708 will result in an activation layer or activations that are transferred to the server 710. This approach is similar to the approach described above with the addition of the differences in data type amongst the clients 702, 704, 706.

[0081] The server 710 will combine those activation layers in one of a multitude of ways. The server 710 can average them (which is also described above), but it could also concatenate them into one long activation layer. In another aspect, the server 710 could apply any mathematical function to achieve the desired combination of the activation layers. The server
5 710 can then process the combined activation layers further using any appropriate network architecture. In one aspect, a server on the client side can receive gradients and average the gradients to generate a global model of the various clients 702, 704, 706 and send the global model to the server 710 for concatenation or for further processing.

[0082] The ideas shown in FIGs. 6 and 7 represent an expansion and application of the split-
10 federated learning tool set and provides a platform of off-the-shelf tools to bring disparate data types together into a superset AI model. The processing can be done all privately and the offering can also be included in a marketplace as described in the incorporated patent applications referenced above.

[0083] Not only can the system combine different data types, but the system can also combine
15 different AI/ML techniques. For example, client one 702 can be a CNN (convolutional neural network), client two 704 can be an ML routine (i.e. XGBoost), and client 3 706 can apply a different technique as well. In this regard, although the different AI/ML techniques are different, as long as the resulting data at the cut layer is consistent and properly configured, the forward propagation and back propagation can occur and the models can be trained.

[0084] In order to assist one of skill in the art to understand how the MMAI approach might
20 work, the following is an example of actual commands per data type coming from the three data providers 702, 704, 706. This code uses the python numbering convention so it starts with builder0 (tabular data from data provider 1 702). Builder1 in this example is for a CT Scan or image data. The commands would be similar for Xray, MRI, and/or any other picture. Builder2
25 (from data provider 704) is text data. Note the “lstm” command, which is short for “long/short

term memory”. The “server” builder commands define the network that aggregates the other three at the “top” on the other side of the split.

```
builder0 = tb.NetworkBuilder()
5 builder0.add_dense_layer(100, 120)
  builder0.add_relu()
  builder0.add_dense_layer(120, 160)
  builder0.add_relu()
  builder0.add_dropout(0.25)
10 builder0.add_dense_layer(160, 200)
  builder0.add_relu()
  builder0.add_split()

builder1 = tb.NetworkBuilder()
15 builder1.add_conv2d_layer(1, 32, 3, 1)
  builder1.add_batchnorm2d(32)
  builder1.add_relu()
  builder1.add_max_pool2d_layer(2, 2)
  builder1.add_conv2d_layer(32, 64, 3, 1)
20 builder1.add_batchnorm2d(64)
  builder1.add_relu()
  builder1.add_max_pool2d_layer(2, 2)
  builder1.add_flatten_layer()
  builder1.add_split()
25
builder2 = tb.NetworkBuilder()
  builder2.add_lstm_layer(39, 100, batch_first=True)
  builder2.add_dense_layer(100, 39)
  builder2.add_split()
30
server_builder = tb.NetworkBuilder()
```

```

server_builder.add_dense_layer(60000, 8000),
server_builder.add_relu()
server_builder.add_dense_layer(8000, 1000),
server_builder.add_relu()
5 server_builder.add_dense_layer(1000, 128),
server_builder.add_relu()
server_builder.add_dense_layer(128, 1)

```

[0085] FIG. 8 illustrate an example method 800 for providing a MMAI concept from the standpoint of the clients. The method includes receiving a first set of data from a first data source, the first set of data having a first data type (802), training a first client-side network on the first set of data and generating first activations (804), receiving a second set of data from a second data source, the second set of data having a second data type (806) and training a second client-side network on the second set of data and generating second activations (808).

15 [0086] The method can further include transmitting the first activations and the second activations to a server-side network, wherein the server-side network is trained based on the first activations and the second activations to generate gradients (810), and receiving the gradients at the first client-side network and the second client-side network (812). The first data type and the second data type can be different data types, such as one being image-based and the other being textual or temporally based as in speech.

[0087] FIG. 9 illustrates an example method 900 from the standpoint of both a server 710 and one or more clients 702, 704, 706. The method can include splitting a neural network into a first client-side network, a second client-side network and a server-side network (902), sending the first client-side network to a first client, wherein the first client-side network is configured to process first data from the first client, the first data having a first type and wherein the first client-side network can include at least one first client-side layer (904), and sending the second client-side network to a second client, wherein the second client-side network is configured to

process second data from the second client, the second data having a second type and wherein the second client-side network can include at least one second client-side layer, wherein the first type and the second type have a common association (906).

[0088] The method can further include training the first client-side network on first data from the first client and generating first activations (908), transmitting the first activations from the first client-side network to the server-side network (910), training the second client-side network on second data from the second client and generating second activations (912), transmitting the second activations from the second client-side network to the server-side network (914), training at least one server-side layer of the server-side network based on the first activations and the second activations to generate gradients (916) and transmitting the gradients from the server-side network to the first client-side network and the second client-side network (918).

[0089] The common association between the disparate types of data can include at least one of a device, a person, a consumer, a patient, a business, a concept, a medical condition, a group of people, a process, a product and/or a service. Any concept, device or person can be the common association or theme of the various disparate types of data that come from different clients and that are processed by different and independent client-side networks up to a cut or split layer. The server-side network can include a global machine learning model. The neural network can include weights, bias and hyperparameters. Hyperparameters typically relate to a parameter whose value is used to control the learning process, such as a topology parameter or a size of a neural network. For example, a learning rate, a mini-batch size, a number of layers on client side, or any parameter related to controlling the process that might impact or relate to different data types can represent a hyperparameter.

[0090] The at least one first client-side layer and the at least one second client-side layer each can include a same number of layers or a different number of layers. Because they operate

independently, the client-side networks can have a different number of layers as long as they process their data to generate vectors or activations that are in a proper format for passing on to the server-side network for further training. A cut layer can exist between the server-side network and the first client-side network and the second client-side network.

- 5 [0091] FIG. 10 illustrates an example method 1000 from the standpoint of the server 710. A method can include splitting a neural network into a first client-side network, a second client-side network and a server-side network (1002), sending the first client-side network to a first client, wherein the first client-side network is configured to process first data from the first client, the first data having a first type and wherein the first client-side network can include at
- 10 least one first client-side layer (1004) and sending the second client-side network to a second client, wherein the second client-side network is configured to process second data from the second client, the second data having a second type and wherein the second client-side network can include at least one second client-side layer, wherein the first type and the second type have a common association (1006).
- 15 [0092] The method can further include receiving, at the server-side network, first activations from a training of the first client-side network on first data from the first client (1008), receiving, at the server-side network, second activations from a training of the second client-side network on second data from the second client (1010), training at least one server-side layer of the server-side network based on the first activations and the second activations to
- 20 generate gradients (1012) and transmitting the gradients from the server-side network to the first client-side network and the second client-side network (1014).

- [0093] Note that in each case, part of the process of the server 710 in terms of training could be performed by the server 710 and other parts such as an averaging of values over the various clients could be performed by a different server (not shown) that could be at a client site, a
- 25 separate location, or across different clients.

[0094] This approach enables the use of the federated split-learning tool set in a new way that when the system splits up the neural network, at the blind correlation 708, the system can make it harder to take the resulting trained model, break it and apply a training inference attack. Because the system can break the neural network in half (or in two portions), and the way it is described above, all that is exchanged from the neural network parts 702A, 704A, 706A is a string or array of numbers, also described as activation layer numbers. Since these are only numbers or an array of characters, what happens at a first neural network portion 702A could be different from what happens at a second neural network portion 704A. For example, the first neural network portion 702A could be 2 layers deep and the second neural network portion 704A could be 90 layers deep. As long as each output resolves to a string of numbers that is structured appropriately for transmission to the top part of the neural network 710, then the forward propagation and the back propagation can work and the training can be achieved. This understanding paves the way for a new concept disclosed herein that different types of data handled across the different portions 702A, 704A, 706A of the neural network can be received and processed properly to train the models(). If the system can create a different bottom half 702A, 704A, 706A for each of different clients, then the clients 702, 704, 706 don't have to produce or process the same type of data (between text and images, for example), but the properly formatted neural network portions 702A, 704A, 706A can process that disparate data, and produce the structured output that can be sent to the server 710.

[0095] In one example, client one 702 might provide a person's ECG, client two 704 can provide a chest X-ray of a heart can client three 706 can provide the genetic profile of the most four interesting proteins in the patient's blood. If the neural network portions 702A, 704A, 706A can process the different respective types of data down to the right vector structure for output, and provide the disparate types of data to the server 710, the server 710 can be

configured with the proper neural network to combine all of that information to train a model to be used to make a diagnosis which can utilize the different and disparate types of data.

[0096] In one aspect, while the neural network portions 702A, 704A, 706A each process a different type of data, there is some correlating factor associated with the data. In the above
5 example, all of the data may relate generally to the same person, although some data is ECG related and other data is associated with a genetic profile, yet they all are for the same person. Thus, one aspect of this disclosure is that the data does have a common association. In another aspect, the data may not be related to the same person but the common association could be related to an age, gender, race, project, concept, the weather, the stock market, or other factors.
10 All of the data might relate to women between the ages of 30-35, for example. Thus, the common association has some flexibility to how it would be applied.

[0097] In another example, the data could be images from a camera of a jet engine stream, another stream of data could be sensor data, and other data could be flight characteristics from an airplane, and the common association could be the airplane. In another aspect, the common
15 association could be a consumer with one type of data being purchasing habits, another type of data being web-surfing patterns, another type of data being emails that the user sends, another type of data being audio from Siri or other speech processing tools, and another type of data being what physical stores the consumer frequents or what is the user's current location. The output of the server could be an advertisement to provide to the user based on the analysis of
20 the disparate types of input. Thus, the common association can relate to any concept that could be used in which disparate types of data can relate to the concept.

[0098] FIG. 11 illustrates example computer device that can be used in connection with any of the systems disclosed herein. In this example, FIG. 11 illustrates a computing system 1100 including components in electrical communication with each other using a connection 1105,
25 such as a bus. System 1100 includes a processing unit (CPU or processor) 1110 and a system

connection 1105 that couples various system components including the system memory 1115, such as read only memory (ROM) 1120 and random access memory (RAM) 1125, to the processor 1110. The system 1100 can include a cache of high-speed memory connected directly with, in close proximity to, or integrated as part of the processor 1110. The system

5 1100 can copy data from the memory 1115 and/or the storage device 1130 to the cache 1112 for quick access by the processor 1110. In this way, the cache can provide a performance boost that avoids processor 1110 delays while waiting for data. These and other modules can control or be configured to control the processor 1110 to perform various actions. Other system memory 1115 may be available for use as well. The memory 1115 can include multiple

10 different types of memory with different performance characteristics. The processor 1110 can include any general purpose processor and a hardware or software service or module, such as service (module) 1 1132, service (module) 2 1134, and service (module) 3 1136 stored in storage device 1130, configured to control the processor 1110 as well as a special-purpose processor where software instructions are incorporated into the actual processor design. The

15 processor 1110 may be a completely self-contained computing system, containing multiple cores or processors, a bus, memory controller, cache, etc. A multi-core processor may be symmetric or asymmetric.

[0099] To enable user interaction with the device 1100, an input device 1145 can represent any number of input mechanisms, such as a microphone for speech, a touch-sensitive screen for

20 gesture or graphical input, keyboard, mouse, motion input, speech and so forth. An output device 1135 can also be one or more of a number of output mechanisms known to those of skill in the art. In some instances, multimodal systems can enable a user to provide multiple types of input to communicate with the device 1100. The communications interface 1140 can generally govern and manage the user input and system output. There is no restriction on

operating on any particular hardware arrangement and therefore the basic features here may easily be substituted for improved hardware or firmware arrangements as they are developed.

[0100] Storage device 1130 is a non-volatile memory and can be a hard disk or other types of computer readable media which can store data that are accessible by a computer, such as
5 magnetic cassettes, flash memory cards, solid state memory devices, digital versatile disks, cartridges, random access memories (RAMs) 1125, read only memory (ROM) 1120, and hybrids thereof.

[0101] The storage device 1130 can include services or modules 1132, 1134, 1136 for controlling the processor 1110. Other hardware or software modules are contemplated. The
10 storage device 1130 can be connected to the system connection 1105. In one aspect, a hardware module that performs a particular function can include the software component stored in a computer-readable medium in connection with the necessary hardware components, such as the processor 1110, connection 1105, output device 1135, and so forth, to carry out the function.

[0102] In some cases, such a computing device or apparatus may include a processor,
15 microprocessor, microcomputer, or other component of a device that is configured to carry out the steps of the methods disclosed above. In some examples, such computing device or apparatus may include one or more antennas for sending and receiving RF signals. In some examples, such computing device or apparatus may include an antenna and a modem for sending, receiving, modulating, and demodulating RF signals, as previously described.

[0103] The components of the computing device can be implemented in circuitry. For
20 example, the components can include and/or can be implemented using electronic circuits or other electronic hardware, which can include one or more programmable electronic circuits (e.g., microprocessors, graphics processing units (GPUs), digital signal processors (DSPs), central processing units (CPUs), and/or other suitable electronic circuits), and/or can include
25 and/or be implemented using computer software, firmware, or any combination thereof, to

perform the various operations described herein. The computing device may further include a display (as an example of the output device or in addition to the output device), a network interface configured to communicate and/or receive the data, any combination thereof, and/or other component(s). The network interface may be configured to communicate and/or receive

5 Internet Protocol (IP) based data or other type of data.

[0104] The methods discussed above are illustrated as a logical flow diagram, the operations of which represent a sequence of operations that can be implemented in hardware, computer instructions, or a combination thereof. In the context of computer instructions, the operations represent computer-executable instructions stored on one or more computer-readable storage

10 media that, when executed by one or more processors, perform the recited operations. Generally, computer-executable instructions include routines, programs, objects, components, data structures, and the like that perform particular functions or implement particular data types. The order in which the operations are described is not intended to be construed as a limitation, and any number of the described operations can be combined in any order and/or in

15 parallel to implement the processes.

[0105] Additionally, the methods disclosed herein may be performed under the control of one or more computer systems configured with executable instructions and may be implemented as code (e.g., executable instructions, one or more computer programs, or one or more applications) executing collectively on one or more processors, by hardware, or combinations

20 thereof. As noted above, the code may be stored on a computer-readable or machine-readable storage medium, for example, in the form of a computer program including a plurality of instructions executable by one or more processors. The computer-readable or machine-readable storage medium may be non-transitory.

[0106] The term “computer-readable medium” includes, but is not limited to, portable or non-

25 portable storage devices, optical storage devices, and various other mediums capable of storing,

containing, or carrying instruction(s) and/or data. A computer-readable medium may include a non-transitory medium in which data can be stored and that does not include carrier waves and/or transitory electronic signals propagating wirelessly or over wired connections. Examples of a non-transitory medium may include, but are not limited to, a magnetic disk or tape, optical storage media such as compact disk (CD) or digital versatile disk (DVD), flash memory, memory or memory devices. A computer-readable medium may have stored thereon code and/or machine-executable instructions that may represent a procedure, a function, a subprogram, a program, a routine, a subroutine, a module, a software package, a class, or any combination of instructions, data structures, or program statements. A code segment may be coupled to another code segment or a hardware circuit by passing and/or receiving information, data, arguments, parameters, or memory contents. Information, arguments, parameters, data, etc. may be passed, forwarded, or transmitted via any suitable means including memory sharing, message passing, token passing, network transmission, or the like.

[0107] In some embodiments the computer-readable storage devices, mediums, and memories can include a cable or wireless signal containing a bit stream and the like. However, when mentioned, non-transitory computer-readable storage media expressly exclude media such as energy, carrier signals, electromagnetic waves, and signals per se.

[0108] Specific details are provided in the description above to provide a thorough understanding of the embodiments and examples provided herein. However, it will be understood by one of ordinary skill in the art that the embodiments may be practiced without these specific details. For clarity of explanation, in some instances the present technology may be presented as including individual functional blocks including devices, device components, steps or routines in a method embodied in software, or combinations of hardware and software. Additional components may be used other than those shown in the figures and/or described herein. For example, circuits, systems, networks, processes, and other components may be

shown as components in block diagram form in order not to obscure the embodiments in unnecessary detail. In other instances, well-known circuits, processes, algorithms, structures, and techniques may be shown without unnecessary detail in order to avoid obscuring the embodiments.

- 5 [0109] Individual embodiments may be described above as a process or method which is depicted as a flowchart, a flow diagram, a data flow diagram, a structure diagram, or a block diagram. Although a flowchart may describe the operations as a sequential process, many of the operations can be performed in parallel or concurrently. In addition, the order of the operations may be re-arranged. A process is terminated when its operations are completed, but
10 can have additional steps not included in a figure. A process may correspond to a method, a function, a procedure, a subroutine, a subprogram, etc. When a process corresponds to a function, its termination can correspond to a return of the function to the calling function or the main function.

- [0110] Processes and methods according to the above-described examples can be implemented
15 using computer-executable instructions that are stored or otherwise available from computer-readable media. Such instructions can include, for example, instructions and data which cause or otherwise configure a general purpose computer, special purpose computer, or a processing device to perform a certain function or group of functions. Portions of computer resources used can be accessible over a network. The computer executable instructions may be, for example,
20 binaries, intermediate format instructions such as assembly language, firmware, source code. Examples of computer-readable media that may be used to store instructions, information used, and/or information created during methods according to described examples include magnetic or optical disks, flash memory, USB devices provided with non-volatile memory, networked storage devices, and so on.

[0111] Devices implementing processes and methods according to these disclosures can include hardware, software, firmware, middleware, microcode, hardware description languages, or any combination thereof, and can take any of a variety of form factors. When implemented in software, firmware, middleware, or microcode, the program code or code segments to perform the necessary tasks (e.g., a computer-program product) may be stored in a computer-readable or machine-readable medium. A processor(s) may perform the necessary tasks. Typical examples of form factors include laptops, smart phones, mobile phones, tablet devices or other small form factor personal computers, personal digital assistants, rackmount devices, standalone devices, and so on. Functionality described herein also can be embodied in peripherals or add-in cards. Such functionality can also be implemented on a circuit board among different chips or different processes executing in a single device, by way of further example.

[0112] The instructions, media for conveying such instructions, computing resources for executing them, and other structures for supporting such computing resources are example means for providing the functions described in the disclosure.

[0113] In the foregoing description, aspects of the application are described with reference to specific embodiments thereof, but those skilled in the art will recognize that the application is not limited thereto. Thus, while illustrative embodiments of the application have been described in detail herein, it is to be understood that the inventive concepts may be otherwise variously embodied and employed, and that the appended claims are intended to be construed to include such variations, except as limited by the prior art. Various features and aspects of the above-described application may be used individually or jointly. Further, embodiments can be utilized in any number of environments and applications beyond those described herein without departing from the broader spirit and scope of the specification. The specification and drawings are, accordingly, to be regarded as illustrative rather than restrictive. For the purposes

of illustration, methods were described in a particular order. It should be appreciated that in alternate embodiments, the methods may be performed in a different order than that described.

[0114] One of ordinary skill will appreciate that the less than (“<”) and greater than (“>”) symbols or terminology used herein can be replaced with less than or equal to (“≤”) and greater
5 than or equal to (“≥”) symbols, respectively, without departing from the scope of this description.

[0115] Where components are described as being “configured to” perform certain operations, such configuration can be accomplished, for example, by designing electronic circuits or other hardware to perform the operation, by programming programmable electronic circuits (e.g.,
10 microprocessors, or other suitable electronic circuits) to perform the operation, or any combination thereof.

[0116] The phrase “coupled to” refers to any component that is physically connected to another component either directly or indirectly, and/or any component that is in communication with another component (e.g., connected to the other component over a wired or wireless
15 connection, and/or other suitable communication interface) either directly or indirectly.

[0117] Claim language or other language reciting “at least one of” a set and/or “one or more” of a set indicates that one member of the set or multiple members of the set (in any combination) satisfy the claim. For example, claim language reciting “at least one of A and B” or “at least one of A or B” means A, B, or A and B. In another example, claim language reciting “at least
20 one of A, B, and C” or “at least one of A, B, or C” means A, B, C, or A and B, or A and C, or B and C, or A and B and C. The language “at least one of” a set and/or “one or more” of a set does not limit the set to the items listed in the set. For example, claim language reciting “at least one of A and B” or “at least one of A or B” can mean A, B, or A and B, and can additionally include items not listed in the set of A and B.

[0118] Although a variety of examples and other information was used to explain aspects within the scope of the appended claims, no limitation of the claims should be implied based on particular features or arrangements in such examples, as one of ordinary skill would be able to use these examples to derive a wide variety of implementations. Further and although some
5 subject matter may have been described in language specific to examples of structural features and/or method steps, it is to be understood that the subject matter defined in the appended claims is not necessarily limited to these described features or acts. For example, such functionality can be distributed differently or performed in components other than those identified herein. Rather, the described features and steps are disclosed as examples of
10 components of systems and methods within the scope of the appended claims.

[0119] Claim language reciting "at least one of" a set indicates that one member of the set or multiple members of the set satisfy the claim. For example, claim language reciting "at least one of A and B" means A, B, or A and B.