



(51) International Patent Classification:

H04L 41/14 (2022.01) H04L 41/122 (2022.01)
H04L 41/16 (2022.01) H04L 41/40 (2022.01)

(21) International Application Number:

PCT/EP2024/060374

(22) International Filing Date:

17 April 2024 (17.04.2024)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

20240100238 03 April 2024 (03.04.2024) GR

(71) Applicant: **LENOVO (SINGAPORE) PTE. LTD.**
[SG/SG]; 151, Lorong Chuan #02-01, New Tech Park, Singapore 556741 (SG).

(72) Inventors: **KARAMPATIS, Dimitrios**; 9 Lawn Close, Ruislip Middlesex HA4 6ED (GB). **PATEROMICHELAKIS, Emmanouil**; Breyeller Str. 37, 41751 Viersen (DE). **SAMDANIS, Konstantinos**; Rablst. 18 VI, 81669 Munich (DE).

(74) Agent: **MARKS & CLERK LLP**; 15 Fetter Lane, London Greater, London EC4A 1BW (GB).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CV, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IQ, IR, IS, IT, JM, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, MG, MK, MN, MU, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, WS, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, CV, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SC, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, ME, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

(54) Title: SUPPORTING FEDERATED LEARNING IN A COMMUNICATION SYSTEM




Figure 1

(57) Abstract: Various aspects of the present disclosure relate to a network entity for wireless communication supporting network data analytics, the network entity comprising instructions executable by at least one processor to cause the network entity to: obtain a request message to train a model for supporting a network analytics operation identified by an analytics identifier; select a set of one or more model training functions to perform a federated learning training process for the model; select a weighting factor for each function of the selected set of one or more model training functions; and control the federated learning training process for the model and the one or more selected functions, based on the selected weighting factor for each said function.

Published:

- *with international search report (Art. 21(3))*
- *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))*
- *upon request of the applicant, before the expiration of the time limit referred to in Article 21(2)(a)*

SUPPORTING FEDERATED LEARNING IN A COMMUNICATION SYSTEM

TECHNICAL FIELD

[0001] The present disclosure relates to wireless communications, and more specifically to network analytics in wireless communication networks.

BACKGROUND

[0002] A wireless communications system may include one or multiple network communication devices, such as base stations, which may support wireless communications for one or multiple user communication devices, which may be otherwise known as user equipment (UE), or other suitable terminology. The wireless communications system may support wireless communications with one or multiple user communication devices by utilizing resources of the wireless communication system (e.g., time resources (e.g., symbols, slots, subframes, frames, or the like) or frequency resources (e.g., subcarriers, carriers, or the like). Additionally, the wireless communications system may support wireless communications across various radio access technologies including third generation (3G) radio access technology, fourth generation (4G) radio access technology, fifth generation (5G) radio access technology, among other suitable radio access technologies beyond 5G (e.g., sixth generation (6G)).

SUMMARY

[0003] An article “a” before an element is unrestricted and understood to refer to “at least one” of those elements or “one or more” of those elements. The terms “a,” “at least one,” “one or more,” and “at least one of one or more” may be interchangeable. As used herein, including in the claims, “or” as used in a list of items (e.g., a list of items prefaced by a phrase such as “at least one of” or “one or more of” or “one or both of”) indicates an inclusive list such that, for example, a list of at least one of A, B, or C means A or B or C or AB or AC or BC or ABC (i.e., A and B and C). Also, as used herein, the phrase “based on” shall not be construed as a reference to a closed set of conditions. For example, an example step that is described as “based on condition A” may be based on both a condition A and a condition B without departing from the scope of the present disclosure. In other words, as

used herein, the phrase “based on” shall be construed in the same manner as the phrase “based at least in part on. Further, as used herein, including in the claims, a “set” may include one or more elements.

[0004] Some implementations of the method and apparatuses described herein may include a network entity for wireless communication supporting network data analytics, the network entity comprising instructions executable by at least one processor to cause the network entity to: obtain a request message to train a model for supporting (e.g. performing) a network analytics operation identified by an analytics identifier; select a set of one or more model training functions to perform a federated learning training process for the model; select a weighting factor for each function of the selected set of one or more model training functions; and control the federated learning training process for the model and the one or more selected functions, based on the selected weighting factor for each said function.

[0005] The instructions may cause the network entity to, responsive to the federated learning training process being completed, output a response message, the response message comprising at least: an identifier of each said model training function; said contribution weight for each said model training function; and an identifier identifying the federated learning training process for the model. The identifier of each model training function may for example comprise an address of each said model training function. The identifier of the model training process may be a process ID of the model training process, for example a single process ID identifying the federated learning training process.

[0006] The response message may comprise one or more of: information identifying the network entity as a function controlling (and/or coordinating) the federated learning process; and information identifying said model training functions as participants in the federated learning process.

[0007] The instructions may cause the network entity to determine completion of the federated learning training process based on a stability of the model.

[0008] The response message may be a response to the request message.

[0009] The instructions may cause the network entity to, responsive to the federated learning training process being completed, output to at least one second network entity of the communication network a training completion confirmation message, the training completion confirmation message comprising at least: an identifier of each said model training function; said contribution weight for each said model training function; and an identifier identifying the federated learning training process for the model. Similarly to that described above, the identifier of the model training process may be a process ID of the model training process, for example a single process ID identifying the federated learning training process.

[0010] The at least one second network entity may comprise one or more of: a network repository function; and an analytics data repository function.

[0011] The instructions may cause the network entity to receive the request message from a consumer entity within the communications network.

[0012] The consumer entity may be a consumer analytics logical function.

[0013] The instructions may cause the network entity to select a first contribution weight, for a given function of said one or more model training functions, based on at least one of: a quantity of data available, for the federated learning training process, from the given function; and a quantified importance weight of data available from the given function.

[0014] The instructions may cause the network entity to perform at least one of: selecting said contribution weight such that a given model training function having a relatively higher quantity of available data has a respective contribution weight higher than the contribution weight of a given model training function having a relatively lower quantity of data; and selecting said contribution weight such that a given model training function having a relatively higher quantified importance level has a respective contribution weight higher than the contribution weight of a given model training function having a relatively lower quantified importance level.

[0015] The instructions may cause the network entity to select said contribution weight such that the given model training function has a respective contribution weight higher than

the contribution weight of a second given model training function, responsive to the given model training function having at least one of: a relatively higher quantity of data than the second given model training function; and a relatively higher quantified importance level than the second given model training function.

[0016] The controlling the federated learning training process may comprise: receiving intermediate values for the model from each of the one or more model training functions; computing at least one of one or more gradients and one or more losses based on the intermediate values; and outputting one or more gradient/loss messages comprising said at least one of the one or more gradients and the one or more losses to the one or more model training functions.

[0017] The controlling the federated learning training process may comprise updating the first contribution weight for each said model training function based on at least one of: said intermediate values; said one or more gradients; said one or more losses; and a number of times said intermediate values have been received from a given model training function within a preset time period.

[0018] The request message may comprise at least one analytics identifier identifying said set of networks analytics operations.

[0019] Said one or more network analytics operations may comprise at least one statistical analysis operation to be performed on the network.

[0020] Said one or more network analytics operations may comprise at least one prediction operation to be performed in respect of behavior of one or more entities within said communication network.

[0021] Said prediction operation may comprise at least one of: a prediction of a parameter associated with at least one entity within the communication network, said parameter optionally being a load of the at least one entity, said at least one entity optionally comprising at least one of a network entity and a UE of the communication network; and a prediction of a movement characteristic of a UE in the communication network, said movement characteristic optionally comprising a location of the UE.

[0022] Each of said one or more model training functions may be a network function for providing analytics data associated with one or more UEs within the communication network.

[0023] The model may be a machine learning model.

[0024] The federated learning model may be a vertical federated learning model.

[0025] The network entity may comprise at least one of: a model training function; and an application function.

[0026] The network entity may be configured to store one or more data labels associated with the model.

[0027] In some implementations of the method and apparatuses described herein, a method is performed by a network entity of a wireless communication network, the method comprising: obtaining a request message to train a model for supporting a network analytics operation identified by an analytics identifier; selecting a set of one or more model training functions to perform a federated learning training process for the model; selecting a weighting factor for each function of the selected set of one or more model training functions; and controlling the federated learning training process for the model and the one or more selected functions, based on the selected weighting factor for each said function.

BRIEF DESCRIPTION OF THE DRAWINGS

[0028] Figure 1 illustrates an example of a wireless communications system in accordance with aspects of the present disclosure.

[0029] Figures 2 and 3 depict a network architecture.

[0030] Figures 4A and 4B illustrate federated learning techniques.

[0031] Figures 5A and 5B depict communication flows for federated learning.

[0032] Figure 6 depicts architectures for federated learning, according to an example.

[0033] Figure 7 depicts a process flow for federated learning, according to an example.

[0034] Figure 8 illustrates an example of a network equipment (NE) 800 in accordance with aspects of the present disclosure.

[0035] Figure 9 illustrate a flowchart of a method performed by a NE in accordance with aspects of the present disclosure.

DETAILED DESCRIPTION

[0036] Some wireless communication systems, including one or more network entities may support network analytics operations (e.g., process, action), for example for performing statistical analyses of activity in these wireless communication systems. For example, analysis may be performed on parameters such as signal quality, traffic volumes, and the like. In some other examples, the network analytics operation may include the one or more network entities performing a prediction (e.g., estimation, determination, forecast) of a characteristic of the wireless communication systems (e.g., networks). For example, the one or more network entities may perform a prediction of a location (e.g., position) of a user equipment (UE) at a given time, or a prediction of a load (e.g., data traffic volume) for the one or more network entities at a given time in the future.

[0037] The present disclosure provides for improvements to the performance, for example the accuracy and/or processing, of such network analytics. Various aspects of the present disclosure relate to using one or more learning techniques, for example, a vertical federated learning technique, to perform network analytics. Federated learning techniques, described in more detail below, are techniques in which multiple entities collaborate to train a model, whilst their data is not centralized. It should be understood that other learning techniques (e.g., machine learning model, artificial intelligence models, or the like) may be used for support and enabling the one or more network entities to support network analytics. Some aspects of the present disclosure relate to control of the federated learning techniques (e.g., process) by one or more network entities, in order to effectively provide improved network analytics. This control, which is described in more detail below, may include transmitting messages to participants in the process to coordinate their training of a model. Such use of learning techniques can provide improved statistical analysis, and improved prediction quality, thereby improving the accuracy of network analytics. Improved network performance can thus be achieved.

[0038] Aspects of the present disclosure are described in the context of a wireless communications system.

[0039] Figure 1 illustrates an example of a wireless communications system 100 in accordance with aspects of the present disclosure. The wireless communications system 100 may include one or more NE 102, one or more UE 104, and a core network (CN) 106. The wireless communications system 100 may support various radio access technologies. In some implementations, the wireless communications system 100 may be a 4G network, such as an LTE network or an LTE-Advanced (LTE-A) network. In some other implementations, the wireless communications system 100 may be a NR network, such as a 5G network, a 5G-Advanced (5G-A) network, or a 5G ultrawideband (5G-UWB) network. In other implementations, the wireless communications system 100 may be a combination of a 4G network and a 5G network, or other suitable radio access technology including Institute of Electrical and Electronics Engineers (IEEE) 802.11 (Wi-Fi), IEEE 802.16 (WiMAX), IEEE 802.20. The wireless communications system 100 may support radio access technologies beyond 5G, for example, 6G. Additionally, the wireless communications system 100 may support technologies, such as time division multiple access (TDMA), frequency division multiple access (FDMA), or code division multiple access (CDMA), etc.

[0040] The one or more NE 102 may be dispersed throughout a geographic region to form the wireless communications system 100. One or more of the NE 102 described herein may be or include or may be referred to as a network node, a base station, a network element, a network function, a network entity, a radio access network (RAN), a NodeB, an eNodeB (eNB), a next-generation NodeB (gNB), or other suitable terminology. An NE 102 and a UE 104 may communicate via a communication link, which may be a wireless or wired connection. For example, an NE 102 and a UE 104 may perform wireless communication (e.g., receive signaling, transmit signaling) over a Uu interface.

[0041] An NE 102 may provide a geographic coverage area for which the NE 102 may support services for one or more UEs 104 within the geographic coverage area. For example, an NE 102 and a UE 104 may support wireless communication of signals related to services (e.g., voice, video, packet data, messaging, broadcast, etc.) according to one or

multiple radio access technologies. In some implementations, an NE 102 may be moveable, for example, a satellite associated with a non-terrestrial network (NTN). In some implementations, different geographic coverage areas 112 associated with the same or different radio access technologies may overlap, but the different geographic coverage areas may be associated with different NE 102.

[0042] The one or more UE 104 may be dispersed throughout a geographic region of the wireless communications system 100. A UE 104 may include or may be referred to as a remote unit, a mobile device, a wireless device, a remote device, a subscriber device, a transmitter device, a receiver device, or some other suitable terminology. In some implementations, the UE 104 may be referred to as a unit, a station, a terminal, or a client, among other examples. Additionally, or alternatively, the UE 104 may be referred to as an Internet-of-Things (IoT) device, an Internet-of-Everything (IoE) device, or machine-type communication (MTC) device, among other examples.

[0043] A UE 104 may be able to support wireless communication directly with other UEs 104 over a communication link. For example, a UE 104 may support wireless communication directly with another UE 104 over a device-to-device (D2D) communication link. In some implementations, such as vehicle-to-vehicle (V2V) deployments, vehicle-to-everything (V2X) deployments, or cellular-V2X deployments, the communication link 114 may be referred to as a sidelink. For example, a UE 104 may support wireless communication directly with another UE 104 over a PC5 interface.

[0044] An NE 102 may support communications with the CN 106, or with another NE 102, or both. For example, an NE 102 may interface with other NE 102 or the CN 106 through one or more backhaul links (e.g., S1, N2, N2, or network interface). In some implementations, the NE 102 may communicate with each other directly. In some other implementations, the NE 102 may communicate with each other or indirectly (e.g., via the CN 106). In some implementations, one or more NE 102 may include subcomponents, such as an access network entity, which may be an example of an access node controller (ANC). An ANC may communicate with the one or more UEs 104 through one or more other access network transmission entities, which may be referred to as a radio heads, smart radio heads, or transmission-reception points (TRPs).

[0045] The CN 106 may support user authentication, access authorization, tracking, connectivity, and other access, routing, or mobility functions. The CN 106 may be an evolved packet core (EPC), or a 5G core (5GC), which may include a control plane entity that manages access and mobility (e.g., a mobility management entity (MME), an access and mobility management functions (AMF)) and a user plane entity that routes packets or interconnects to external networks (e.g., a serving gateway (S-GW), a Packet Data Network (PDN) gateway (P-GW), or a user plane function (UPF)). In some implementations, the control plane entity may manage non-access stratum (NAS) functions, such as mobility, authentication, and bearer management (e.g., data bearers, signal bearers, etc.) for the one or more UEs 104 served by the one or more NE 102 associated with the CN 106.

[0046] The CN 106 may communicate with a packet data network over one or more backhaul links (e.g., via an S1, N2, N2, or another network interface). The packet data network may include an application server. In some implementations, one or more UEs 104 may communicate with the application server. A UE 104 may establish a session (e.g., a protocol data unit (PDU) session, or the like) with the CN 106 via an NE 102. The CN 106 may route traffic (e.g., control information, data, and the like) between the UE 104 and the application server using the established session (e.g., the established PDU session). The PDU session may be an example of a logical connection between the UE 104 and the CN 106 (e.g., one or more network functions of the CN 106).

[0047] In the wireless communications system 100, the NEs 102 and the UEs 104 may use resources of the wireless communications system 100 (e.g., time resources (e.g., symbols, slots, subframes, frames, or the like) or frequency resources (e.g., subcarriers, carriers)) to perform various operations (e.g., wireless communications). In some implementations, the NEs 102 and the UEs 104 may support different resource structures. For example, the NEs 102 and the UEs 104 may support different frame structures. In some implementations, such as in 4G, the NEs 102 and the UEs 104 may support a single frame structure. In some other implementations, such as in 5G and among other suitable radio access technologies, the NEs 102 and the UEs 104 may support various frame structures (i.e., multiple frame structures). The NEs 102 and the UEs 104 may support various frame structures based on one or more numerologies.

[0048] One or more numerologies may be supported in the wireless communications system 100, and a numerology may include a subcarrier spacing and a cyclic prefix. A first numerology (e.g., $\mu=0$) may be associated with a first subcarrier spacing (e.g., 15 kHz) and a normal cyclic prefix. In some implementations, the first numerology (e.g., $\mu=0$) associated with the first subcarrier spacing (e.g., 15 kHz) may utilize one slot per subframe. A second numerology (e.g., $\mu=1$) may be associated with a second subcarrier spacing (e.g., 30 kHz) and a normal cyclic prefix. A third numerology (e.g., $\mu=2$) may be associated with a third subcarrier spacing (e.g., 60 kHz) and a normal cyclic prefix or an extended cyclic prefix. A fourth numerology (e.g., $\mu=3$) may be associated with a fourth subcarrier spacing (e.g., 120 kHz) and a normal cyclic prefix. A fifth numerology (e.g., $\mu=4$) may be associated with a fifth subcarrier spacing (e.g., 240 kHz) and a normal cyclic prefix.

[0049] A time interval of a resource (e.g., a communication resource) may be organized according to frames (also referred to as radio frames). Each frame may have a duration, for example, a 10 millisecond (ms) duration. In some implementations, each frame may include multiple subframes. For example, each frame may include 10 subframes, and each subframe may have a duration, for example, a 1 ms duration. In some implementations, each frame may have the same duration. In some implementations, each subframe of a frame may have the same duration.

[0050] Additionally or alternatively, a time interval of a resource (e.g., a communication resource) may be organized according to slots. For example, a subframe may include a number (e.g., quantity) of slots. The number of slots in each subframe may also depend on the one or more numerologies supported in the wireless communications system 100. For instance, the first, second, third, fourth, and fifth numerologies (i.e., $\mu=0$, $\mu=1$, $\mu=2$, $\mu=3$, $\mu=4$) associated with respective subcarrier spacings of 15 kHz, 30 kHz, 60 kHz, 120 kHz, and 240 kHz may utilize a single slot per subframe, two slots per subframe, four slots per subframe, eight slots per subframe, and 16 slots per subframe, respectively. Each slot may include a number (e.g., quantity) of symbols (e.g., OFDM symbols). In some implementations, the number (e.g., quantity) of slots for a subframe may depend on a numerology. For a normal cyclic prefix, a slot may include 14 symbols. For an extended cyclic prefix (e.g., applicable for 60 kHz subcarrier spacing), a slot may include 12

symbols. The relationship between the number of symbols per slot, the number of slots per subframe, and the number of slots per frame for a normal cyclic prefix and an extended cyclic prefix may depend on a numerology. It should be understood that reference to a first numerology (e.g., $\mu=0$) associated with a first subcarrier spacing (e.g., 15 kHz) may be used interchangeably between subframes and slots.

[0051] In the wireless communications system 100, an electromagnetic (EM) spectrum may be split, based on frequency or wavelength, into various classes, frequency bands, frequency channels, etc. By way of example, the wireless communications system 100 may support one or multiple operating frequency bands, such as frequency range designations FR1 (410 MHz – 7.125 GHz), FR2 (24.25 GHz – 52.6 GHz), FR3 (7.125 GHz – 24.25 GHz), FR4 (52.6 GHz – 114.25 GHz), FR4a or FR4-1 (52.6 GHz – 71 GHz), and FR5 (114.25 GHz – 300 GHz). In some implementations, the NEs 102 and the UEs 104 may perform wireless communications over one or more of the operating frequency bands. In some implementations, FR1 may be used by the NEs 102 and the UEs 104, among other equipment or devices for cellular communications traffic (e.g., control information, data). In some implementations, FR2 may be used by the NEs 102 and the UEs 104, among other equipment or devices for short-range, high data rate capabilities.

[0052] FR1 may be associated with one or multiple numerologies (e.g., at least three numerologies). For example, FR1 may be associated with a first numerology (e.g., $\mu=0$), which includes 15 kHz subcarrier spacing; a second numerology (e.g., $\mu=1$), which includes 30 kHz subcarrier spacing; and a third numerology (e.g., $\mu=2$), which includes 60 kHz subcarrier spacing. FR2 may be associated with one or multiple numerologies (e.g., at least 2 numerologies). For example, FR2 may be associated with a third numerology (e.g., $\mu=2$), which includes 60 kHz subcarrier spacing; and a fourth numerology (e.g., $\mu=3$), which includes 120 kHz subcarrier spacing.

[0053] In the wireless communication system 100, a NE 102 may be configured (as described in more detail below) to control a federated learning process to perform network analytics.

[0054] Figure 2 illustrates a network architecture that may allow data collection relating to the use of network data analytics functions (NWDAFs). In some cases, an analytics logical function (AnLF) requests a trained machine learning (ML) model from the model training logical function (MTLF) by including in the request an Analytic ID corresponding to the Analytic ID requested by an analytics consumer. Information on the available Analytic ID(s) supported by an AnLF is provided in 3GPP TS 23.288. The MTLF trains an ML model corresponding to the Analytic ID requested by collecting data from one or more data sources (NF, OAM or UEs).

[0055] The functionality of a NWDAF, such as that of Figure 1, may be divided between an AnLF and a MTLF. Figure 3 illustrates an example of such a division of functionality.

[0056] In Figure 3, MTLFs each train a ML model for performing a respective network analytic. Each MTLF may be associated with an analytic ID. Each such analytic ID identifies a particular network analytic. A given ML model may be configured by a ML model designer.

[0057] Each MTLF may train its respective model based on data collected from data producer network functions (NFs). Data producer NFs may be any network function that produces data. For example, one data producer NF may be an access and mobility function (AMF), producing location information in respect of UEs. Alternatively or additionally, MTLFs may train their respective models based on historical data received from a data collection coordination function (DCCF) or from an analytics data repository function (ADRF).

[0058] The AnLF then receives the trained ML models (and/or data indicative of the trained ML models) for a specific analytic ID from the MTLF corresponding to that analytic ID. The AnLF uses the model, potentially in combination with new data from the data producer NF(s) and/or DCCF, to derive analytics that have been requested by a consumer NF.

[0059] The AnLF may subscribe to the MTLF to receive trained ML models for an Analytic ID using a specific API (prior art). The AnLF subscribes based on an analytics

request from a consumer NF. An NWDAF may support both Analytics inference and ML Model Training functions.

[0060] In some cases it may not be possible to exchange data directly between data producer NFs and the NWDAF (i.e. MTLF or AnLF). This may for example be because of privacy concerns: an operator of an edge network or a private slice may not want to share data to external networks. As another example, this may be because of a high signalling load: signalling load can be increased considerably when data is provided to a centralized training function such as that provided by the MTLFs. It can be more efficient for a network at the edge to receive trained models rather than collect data to derive analytics.

[0061] One way for addressing situations in which data is not to be directly exchanged (for example because of privacy concerns or signalling load efficiency) is via federated learning. Federated learning allows local model training functions to exchange model parameters and aggregate trained model and/or intermediate training results, instead of centrally training models by collecting raw data. Federated learning is thus a distributed machine learning framework that allows a model to be trained collectively from data that is distributed across different data owners. This provides an advantage that AI/ML models can be trained closer to the data source, rather than sending raw data to a centralised training node. Instead of sending raw data, parameters, weights or intermediate results of a ML model can be sent back to the centralized node to assist generic model training.

[0062] It will be understood that federated learning can be conceptually divided into horizontal and vertical federated learning. These are schematically depicted in Figure 4A and 4B, respectively.

[0063] Horizontal federated learning, or sample-based federated learning, can be applied in scenarios having data sets that share the same feature space (e.g. each data set includes the same features) but have different samples (for example being collected from different users). This is illustrated in Figure 4A.

[0064] Vertical federated learning, or feature-based federated learning, can be applied in scenarios having data sets that share the same sample space (e.g. each applies to the same

set of users) but differ in feature space (e.g. different features are collected in each data set). This is illustrated in Figure 4B.

[0065] In both horizontal and vertical federated learning models, parameters from each local model training function are sent to a model aggregator to calculate an aggregate model. The model aggregator provides updated model parameters to each MTLF that each MTLF uses to re-train its own model thus allowing every local model training function to have a trained model using data from multiple sources.

[0066] There are different approaches of how to train a model with vertical federated learning (VFL), including “split” and “non-split”. Examples of these approaches are now described in Vertical Federated Learning: Taxonomies, Threats, and Prospects, Qun Li et.al. In particular, Figure 2 thereof depicts a taxonomy of VFL algorithms including split and non-split. Figure 3 thereof depicts non-split VFL in more detail, Figure 4 thereof depicts split VFL.

[0067] In both split and non-split VFL, one party is the “label owner” or “active participant”. This party has the ability to classify input data. The other parties are “passive participants” or “workers” which participate in the VFL training process.

[0068] In non-split VFL, passive participants send intermediate results based on data collected locally using their own model to active participants and the active participants computes gradients/losses using as basis the labels and its own ML model. Some scenarios also include a coordinator that ensures exchanges of messages between VFL parties are encrypted (as depicted in Li Fig. 3).

[0069] In split VFL, in contrast, the model is split between several parties. One party own the top model (label owner/VLF server) and other parties own one or more bottom models (passive participants). This is illustrated in Li Fig. 4. The label owner may also be the active participant. In contrast to non-split VFL, the VFL server is aware of the labels and is able to compute gradient/losses that are shared to passive participants

[0070] Figures 5A and 5B depict communication flows for federated learning amongst different NWDAFs. These communication flows are described in 3GPP TS 23.288.

[0071] Figure 5A depicts the registration, discovery and selection of clients for federated learning. Each MTLF registers its FL capability (FL client or FL server) within the NWDAF profile in the NRF. The profile includes additionally information on the analytic IDs supported for FL and/or location area and/or Event ID(s) where data collection is supported. The NWDAF MTLF supporting FL server capability discovers candidate FL clients for training a model using federated learning by receiving a list of FL clients from the NRF that support the analytic ID and/or location area and/or Event IDs needed to train the model. The FL server sends a preparation request to all candidates to check if they are capable to join the FL process, i.e., in terms of the data availability and required time schedule, (steps 7-9) and each FL client responds if it can join the process or not.

[0072] 5B depicts the procedures for FL once the FL clients have been selected and joined the FL process.

[0073] The general procedure for FL is that the FL server sends a subscription request related to federated learning for model training to each selected FL client including local model parameters and each FL client trains the model using data from local NFs.

[0074] Examples of the present disclosure provide solutions to support vertical federated learning enabling the 5GS to assist in collaborative AI/ML operation involving the 5GC/NWDAF or AF. In particular, in order to support VFL, the present disclosure provides methods for VFL entities (either server, active participants or passive participants) to participate in the VFL training process. The following description also sets out actions that the server, active participant may take after a VFL training process is completed (e.g. model information storage).

[0075] Figure 6 depicts various architectures for performing VFL according to examples of the present disclosure. These architectures depict communication between a VFL server 610, an active participant 615 and passive participants 625. In this example, a VFL server 710, which may also be termed a VFL coordinator or controller, is a function that manages the VFL procedures, selects VFL participants 615, 625 for model training, and assigns functions to act as VFL active participant 615 or passive participant 625.

[0076] In this example, a VFL active participant 615 is a VFL function that owns part of a ML model for a given analytic ID and knows the labels for the ML model. The active participant 615 is the main function for training an ML model for a given analytic ID.

[0077] In this example, a VFL passive participant 625 is a VFL function that owns part of an ML model for a given analytic ID, does not know the labels of the ML model, but is able to collect local data for one or more features.

[0078] A given VFL function in the 3GPP network may support a combination of the functions above, e.g., support both VFL server and active participant functions.

[0079] In the present example, the VFL server 610 or active participant 615 identifies how each VFL passive participant 625 contributes in the model training process and assigns weights which are subsequently used for inference. The VFL server 610 or active participant 615 stores the ML model information in a function so that a VFL server can subsequently retrieve the details if the ML model needs to be retrained.

[0080] In this example, a given VFL function (i.e. a VFL server 610, active participant 615, and/or passive participant 625) may be part of the NWDAF AnLF, NWDAF MTLF, or a function in an AF. When interaction between VFL functions from NWDAF and VFL function from AF takes place, then an exchange of messages may be performed via a network exposure function (NEF) 628.

[0081] Figure 6 sets out four architecture scenarios including a VFL server 610, active participants 615, passive participants 625, with involvement of NWDAF AnLF, MTLF and AF. Specifically, the top row of Figure 6 depicts examples in which VFL is performed within a core network, and the bottom row depicts examples in which VFL is initiated outside the core network. Within those rows, the left hand side of Figure 6 depicts examples in which a VFL server 610 is included, and the right hand side depicts examples in which no VFL server 610 is used.

[0082] Figure 7 depicts a process flow in accordance with aspects of the present disclosure. The process flow may implement aspects of the wireless communications system 100 as described with reference to Figure 1, and more particularly may be implemented within one or more of the architectures depicted in Figure 6. The process flow

may involve one or more of a consumer AnLF 705, a VFL server 710, a VFL active participant 715, a data producer 720, a VFL passive participant 725, a data producer 730, a VFL passive participant 735, a data producer 740, a NRF 745 and a ADRF 750. The process flow can thus be implemented within the architecture depicted in Figure 6, with same-named entities having the same functions as described above. In the following description, the operations performed by one or more of the entities 705-750 may be performed in a different order from that shown, or at different times. Some operations may also be omitted from the process flow, and other operations may be added.

[0083] In the example of Figure 7, the VFL server 710 controls (i.e. coordinates) a VFL operation performed with the VFL active participant 715 and passive participants 725, 735. This is based on data from data producers 720, 730, 740. In other examples, a smaller or larger number of participants and/or data producers may take part in the VFL process. The process flow of Figure 7 thus provides an effective way for conducting a VFL process in a wireless communication network. As explained above, this improves the accuracy and performance of network analytics, thereby improving system performance.

[0084] At 1, a consumer 705 (in this example, a NWDAF supporting AnLF) requests a trained ML model for supporting derivation of analytics (e.g. DN performance analytics). This request includes the Analytics ID and/or Model ID of the trained model needed. A list of some analytics, one or more of which may be supported by the NWDAF, is set out in 3GPP TS 23.288. The request is transmitted to the VFL server 710. In an alternative example, the request is transmitted to the VFL active participant 715 or a server supporting a combination of the functions of the VFL server 710 and VFL active participant 715.

[0085] At 2, the VFL server 710 (and/or VFL active participant 715) determines that model training using federated learning is to be performed. This may include one or more of identifying which VFL functions acting as passive participants (i.e. do not own the labels) can participate in the VFL process; identifying what data (feature or samples) and the size of the data available in each participant; performing feature/sample alignment (so that all VFL participants perform VFL using the same samples); and/or selecting the VFL passive participants 725, 735 that can participate in the VFL process for model training

[0086] At 3, based on the selected VFL passive participants 725, 735, the VFL server 710 and/or active participant 715 identifies initial contribution weights from each passive participant 725, 735. The contribution weights may be based on the importance of the features and/or the size of the data available by each active participant. A VFL process ID is also assigned for the training process so that all VFL participants 715, 725, 735 associate the exchange of VFL related messages (including the VFL process ID) with the single VFL process initiated by the VFL server 710 and/or active participant 715.

[0087] At 4, each participant 715, 725, 735 in the VFL process trains their model using the available data (features) and derives an intermediate value (4a, 4b, 4c)

[0088] At 5, the intermediate value from each passive participant 725, 735 is sent to the active participant 715 and/or directly to the VFL server 710. If the VFL server 710 is used the active participant 715 may provide its derived intermediate value to the VFL server 710.

[0089] At 6, the VFL server 710 or active participant 715 computes the gradient/loss based on the intermediate values received. The gradient/loss may be an aggregate value based on all intermediate values received or a gradient/loss value per VFL passive participant 725, 735.

[0090] At 7, the gradient/loss is sent to each passive participant 725, 735 or active participant 715 (if the VFL server 710 is used).

[0091] At 8, each VFL active/passive participant 715, 725, 735 updates their ML model using the gradient/loss and further input data (8a, 8b, 8c).

[0092] At 9, one or more iterations of 4 to 8 take place until the VFL server 710 or VFL active participant 715 determines that the ML model has been trained with the confidence level requested by the ML model consumer.

[0093] At 10, after the VFL iterations are complete, the VFL server/active participant 710, 715 updates the contribution weights from each passive participant 725, 730 e.g., based on the number of times feedback (i.e. intermediate values) has been provided or

feedback provided within a time limit, by each VFL passive participant 725, 735, gradient/loss estimation on per passive participant.

[0094] At 11, once the ML model is trained the VFL server/active participant 710, 715 indicates to the ML model consumer 705 that the trained ML model is available. This response may include one or more of the addresses and/or NF IDs of each VFL participant 715, 725, 735 and their respective contribution weights, an address of the VFL server 710 (if used), an address of the VFL active participant 715, and or a process ID of the VFL process.

[0095] At 12, the VFL server/active participant 710, 715 may also update the NRF 745 with an indication that the model for a given analytic ID was trained using VFL. This indication may include the address/NF IDs of one or more of each VFL participant 715, 725, 735, the VFL server 710 and/or VFL active participant 715, and/or the VFL process ID.

[0096] At 13, the VFL server/active participant 710, 715 may also indicate to the ADRF 750 that the model for a given analytic ID was trained using VFL. This indication may include the address/NF IDs of one or more of each VFL participant 715, 725, 735, the VFL server 710, and/or the VFL active participant 715. The address where the global model is stored may also be included in the request. The VFL process ID may also be stored.

[0097] The architectures of Figure 6 in combination with the process flow of Figure 7 thus provide an effective way for performing VFL within a wireless communication network.

[0098] Figure 8 illustrates an example of a NE 800 in accordance with aspects of the present disclosure. The NE 800 may include a processor 802, a memory 804, a controller 806, and a transceiver 808. The processor 802, the memory 804, the controller 806, or the transceiver 808, or various combinations thereof or various components thereof may be examples of means for performing various aspects of the present disclosure as described herein. These components may be coupled (e.g., operatively, communicatively, functionally, electronically, electrically) via one or more interfaces.

[0099] The processor 802, the memory 804, the controller 806, or the transceiver 808, or various combinations or components thereof may be implemented in hardware (e.g., circuitry). The hardware may include a processor, a digital signal processor (DSP), an application-specific integrated circuit (ASIC), or other programmable logic device, or any combination thereof configured as or otherwise supporting a means for performing the functions described in the present disclosure.

[0100] The processor 802 may include an intelligent hardware device (e.g., a general-purpose processor, a DSP, a CPU, an ASIC, an FPGA, or any combination thereof). In some implementations, the processor 802 may be configured to operate the memory 804. In some other implementations, the memory 804 may be integrated into the processor 802. The processor 802 may be configured to execute computer-readable instructions stored in the memory 804 to cause the NE 800 to perform various functions of the present disclosure.

[0101] The memory 804 may include volatile or non-volatile memory. The memory 804 may store computer-readable, computer-executable code including instructions when executed by the processor 802 cause the NE 800 to perform various functions described herein. The code may be stored in a non-transitory computer-readable medium such the memory 804 or another type of memory. Computer-readable media includes both non-transitory computer storage media and communication media including any medium that facilitates transfer of a computer program from one place to another. A non-transitory storage medium may be any available medium that may be accessed by a general-purpose or special-purpose computer.

[0102] In some implementations, the processor 802 and the memory 804 coupled with the processor 802 may be configured to cause the NE 800 to perform one or more of the functions described herein (e.g., executing, by the processor 802, instructions stored in the memory 804). For example, the processor 802 may support wireless communication at the NE 800 in accordance with examples as disclosed herein. The NE 800 may be configured to support a means for obtaining a request message to train a model for supporting a network analytics operation identified by an analytics identifier; selecting a set of one or more model training functions to perform a federated learning training process for the model; selecting a weighting factor for each function of the selected set of one or more

model training functions; and controlling the federated learning training process for the model and the one or more selected functions, based on the selected weighting factor for each said function.

[0103] The controller 806 may manage input and output signals for the NE 800. The controller 806 may also manage peripherals not integrated into the NE 800. In some implementations, the controller 806 may utilize an operating system such as iOS®, ANDROID®, WINDOWS®, or other operating systems. In some implementations, the controller 806 may be implemented as part of the processor 802.

[0104] In some implementations, the NE 800 may include at least one transceiver 808. In some other implementations, the NE 800 may have more than one transceiver 808. The transceiver 808 may represent a wireless transceiver. The transceiver 808 may include one or more receiver chains 810, one or more transmitter chains 812, or a combination thereof.

[0105] A receiver chain 810 may be configured to receive signals (e.g., control information, data, packets) over a wireless medium. For example, the receiver chain 810 may include one or more antennas for receive the signal over the air or wireless medium. The receiver chain 810 may include at least one amplifier (e.g., a low-noise amplifier (LNA)) configured to amplify the received signal. The receiver chain 810 may include at least one demodulator configured to demodulate the receive signal and obtain the transmitted data by reversing the modulation technique applied during transmission of the signal. The receiver chain 810 may include at least one decoder for decoding the processing the demodulated signal to receive the transmitted data.

[0106] A transmitter chain 812 may be configured to generate and transmit signals (e.g., control information, data, packets). The transmitter chain 812 may include at least one modulator for modulating data onto a carrier signal, preparing the signal for transmission over a wireless medium. The at least one modulator may be configured to support one or more techniques such as amplitude modulation (AM), frequency modulation (FM), or digital modulation schemes like phase-shift keying (PSK) or quadrature amplitude modulation (QAM). The transmitter chain 812 may also include at least one power amplifier configured to amplify the modulated signal to an appropriate power level suitable

for transmission over the wireless medium. The transmitter chain 812 may also include one or more antennas for transmitting the amplified signal into the air or wireless medium.

[0107] Figure 9 illustrates a flowchart of a method in accordance with aspects of the present disclosure. The operations of the method may be implemented by a NE as described herein. In some implementations, the NE may execute a set of instructions to control the function elements of the NE to perform the described functions. The operations of the method may be performed in accordance with examples as described herein. In some implementations, aspects of the operations may be performed by a NE as described with reference to Figure 8. One or more of the operations may be omitted, or performed in a different order from the example shown.

[0108] At 902, the method may include obtaining a request message to train a model for supporting a network analytics operation identified by an analytics identifier;

[0109] At 904, the method may include selecting a set of one or more model training functions to perform a federated learning training process for the model.

[0110] At 906, the method may include selecting a weighting factor for each function of the selected set of one or more model training functions

[0111] At 908, the method may include controlling the federated learning training process for the model and the one or more selected functions, based on the selected weighting factor for each said function.

[0112] It should be noted that the method described herein describes A possible implementation, and that the operations and the steps may be rearranged or otherwise modified and that other implementations are possible.

[0113] The description herein is provided to enable a person having ordinary skill in the art to make or use the disclosure. Various modifications to the disclosure will be apparent to a person having ordinary skill in the art, and the generic principles defined herein may be applied to other variations without departing from the scope of the disclosure. Thus, the disclosure is not limited to the examples and designs described herein but is to be accorded the broadest scope consistent with the principles and novel features disclosed herein.

CLAIMS

What is claimed is:

1. A network entity for wireless communication supporting network data analytics, the network entity comprising instructions executable by at least one processor to cause the network entity to:

obtain a request message to train a model for supporting a network analytics operation identified by an analytics identifier;

select a set of one or more model training functions to perform a federated learning training process for the model;

select a weighting factor for each function of the selected set of one or more model training functions; and

control the federated learning training process for the model and the one or more selected functions, based on the selected weighting factor for each said function.

2. The network entity of claim 1, wherein the instructions cause the network entity to, responsive to the federated learning training process being completed, output a response message, the response message comprising at least:

an identifier of each said model training function;

said contribution weight for each said model training function; and

an identifier identifying the federated learning training process for the model.

3. The network entity of claim 2, wherein the response message comprises at least one of:

information identifying the network entity as a controlling function of the federated learning process; and

information identifying said model training functions as participants in the federated learning process.

4. The network entity of claim 2 or claim 3, wherein the instructions cause the network entity to determine completion of the federated learning training process based on a stability of the model.

5. The network entity of any of claims 2 to 4, wherein the response message is a response to the request message.

6. The network entity of any preceding claim, wherein the instructions cause the network entity to, responsive to the federated learning training process being completed, output to at least one second network entity of the communication network a training completion confirmation message, the training completion confirmation message comprising at least:

an identifier of each said model training function; and

said contribution weight for each said model training function.

an identifier identifying the federated learning training process for the model.

7. The network entity of claim 6, wherein the at least one second network entity comprises one or more of:

a network repository function; and

an analytics data repository function.

8. The network entity of any preceding claim, wherein the instructions cause the network entity to receive the request message from a consumer entity within the communications network.

9. The network entity of claim 8, wherein the consumer entity is a consumer analytics logical function.

10. The network entity of any preceding claim, wherein the instructions cause the network entity to select a first contribution weight, for a given function of said one or more model training functions, based on at least one of:

a quantity of data available, for the federated learning training process, from the given function; and

a quantified importance weight of data available from the given function.

11. The network entity of claim 10, wherein the instructions cause the network entity to perform at least one of:

selecting said contribution weight such that a given model training function having a relatively higher quantity of available data has a respective contribution weight higher than the contribution weight of a given model training function having a relatively lower quantity of data; and

selecting said contribution weight such that a given model training function having a relatively higher quantified importance level has a respective contribution weight higher than the contribution weight of a given model training function having a relatively lower quantified importance level.

12. The network entity of claim 10 or claim 11, wherein the instructions cause the network entity to select said contribution weight such that the given model training function has a respective contribution weight higher than the contribution weight of a second given model training function, responsive to the given model training function having at least one of:

a relatively higher quantity of data than the second given model training function;
and

a relatively higher quantified importance level than the second given model training function.

13. The network entity of any preceding claim, wherein the controlling the federated learning training process comprises:

receiving intermediate values for the model from each of the one or more model training functions;

computing at least one of one or more gradients and one or more losses based on the intermediate values; and

outputting one or more gradient/loss messages comprising said at least one of the one or more gradients and the one or more losses to the one or more model training functions.

14. The network entity of claim 13, wherein the controlling the federated learning training process comprises updating the first contribution weight for each said model training function based on at least one of:

said intermediate values;

said one or more gradients;

said one or more losses; and

a number of times said intermediate values have been received from a given model training function within a preset time period.

15. The network entity of any preceding claim, wherein the request message comprises at least one analytics identifier identifying said set of networks analytics operations.

16. The network entity of any preceding claim, wherein the federated learning model is a vertical federated learning model.

17. The network entity of any preceding claim, wherein the network entity comprises at least one of:

a model training function; and

an application function.

18. The network entity of any preceding claim, wherein the network entity is configured to store one or more data labels associated with the model.

19. A method, performed by a network entity of a wireless communication network, the method comprising:

obtaining a request message to train a model for supporting a network analytics operation identified by an analytics identifier;

selecting a set of one or more model training functions to perform a federated learning training process for the model;

selecting a weighting factor for each function of the selected set of one or more model training functions; and

controlling the federated learning training process for the model and the one or more selected functions, based on the selected weighting factor for each said function.




Figure 1

2/10




Figure 2




Figure 3

4/10




Figure 4A




Figure 4B

5/10




Figure 5A




Figure 5B




Figure 6

8/10




figure 7




Figure 8

10/10




Figure 9

INTERNATIONAL SEARCH REPORT

International application No

PCT/EP2024/060374

A. CLASSIFICATION OF SUBJECT MATTER

INV. H04L41/14 H04L41/16 H04L41/122 H04L41/40
ADD.

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

EPO-Internal**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	"3rd Generation Partnership Project; Technical Specification Group Services and System Aspects; Study of Enablers for Network Automation for 5G 5G System (5GS); Phase 3 (Release 18)", 3GPP STANDARD; TECHNICAL REPORT; 3GPP TR 23.700-81, 3RD GENERATION PARTNERSHIP PROJECT (3GPP), MOBILE COMPETENCE CENTRE ; 650, ROUTE DES LUCIOLES ; F-06921 SOPHIA-ANTIPOLIS CEDEX ; FRANCE , no. V0.4.0 6 September 2022 (2022-09-06), pages 1-257, XP052210688, Retrieved from the Internet: URL:https://ftp.3gpp.org/Specs/archive/23_ series/23.700-81/23700-81-040.zip 23700-81-040_MCCclean.docx [retrieved on 2022-09-06] - / - -	1 - 19



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance;; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance;; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

19 November 2024

Date of mailing of the international search report

26/11/2024

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040,
Fax: (+31-70) 340-3016

Authorized officer

Camba, Sonia

INTERNATIONAL SEARCH REPORT

International application No

PCT/EP2024/060374

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
	page 23 - page 28 page 35 - page 40 page 60 page 79 - page 91 page 103 - page 104 page 111 - page 115 page 121 page 150 - page 151 page 177 - page 179 page 205 - page 206 page 235; figures 6.1.2.2-1, 6.1.2.2-2, 6.30.2-2, 6.42.2.2-1 -----	
X	THOMAS BELLING ET AL: "Accuracy vs Metrics terminology", 3GPP DRAFT; S2-2402986; TYPE DISCUSSION; ENA_PH3, 3RD GENERATION PARTNERSHIP PROJECT (3GPP), MOBILE COMPETENCE CENTRE ; 650, ROUTE DES LUCIOLES ; F-06921 SOPHIA-ANTIPOLIS CEDEX ; FRANCE , vol. SA WG2, no. Athens, GR; 20240226 - 20240301 16 February 2024 (2024-02-16), XP052566946, Retrieved from the Internet: URL:https://www.3gpp.org/ftp/tsg_sa/WG2_Arch/TSGS2_161_Athens_2024-02/Docs/S2-2402986.zip S2-2402986 eNA_Ph3 metrics discussion.docx [retrieved on 2024-02-16] page 4 - page 6 page 8 page 24 - page 30 page 64 -----	1-3,5, 8-12,15, 17,19