(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
18 June 2020 (18.06.2020)

WIPO | PCT

(10) International Publication Number
WO 2020/122902 A1

(51) International Patent Classification:
G06N 20/00 (2019.01) *G06F 21/62* (2013.01)

(21) International Application Number:
PCT/US2018/065255

(22) International Filing Date:
12 December 2018 (12.12.2018)

(25) Filing Language: English

(26) Publication Language: English

(71) Applicant: **HEWLETT-PACKARD DEVELOPMENT COMPANY, L.P.** [US/US]; 10300 Energy Drive, Spring, Texas 77389 (US).

(72) Inventor: **MACEDO, Augusto Queiroz de**; Av. Ipiranga, 6681, Predios 5-6, Predio 95 -5, 6 e 7 andares, 90619-900 Porto Alegre (BR).

(74) Agent: **LEMMON, Marcus B.** et al.; HP Inc., 3390 E. Harmony Road, Mail Stop 35, Fort Collins, Colorado 80528 (US).

(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, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, 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, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States *(unless otherwise indicated, for every kind of regional protection available)*: ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, 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, 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).

Declarations under Rule 4.17:
— *as to the identity of the inventor (Rule 4.17(i))*
— *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*

Published:
— *with international search report (Art. 21(3))*

(54) Title: UPDATES OF MACHINE LEARNING MODELS BASED ON CONFIDENTIAL DATA

(57) Abstract: An example of a computer-readable medium to store machine-readable instructions. The instructions may cause a processor to update a machine learning model based on first confidential data set. The machine learning model may have been previously trained based on a training data set distinct from the confidential data set. The instructions may cause the processor to apply the machine learning model after the update to a second confidential data set.

WO 2020/122902 A1

UPDATES OF MACHINE LEARNING MODELS BASED ON CONFIDENTIAL DATA

BACKGROUND

[0001] Machine learning models may be trained on data to detect patterns and make predictions. The data used to train machine learning models may be kept confidential by a corporate entity and not shared with other entities.

BRIEF DESCRIPTION OF THE DRAWINGS

[0002] Various examples will be described below referring to the following figures:

[0003] Fig. 1 shows a computer system to update and apply a machine learning model in accordance with various examples;

[0004] Fig. 2 shows a method of updating and applying a machine learning model in accordance with various examples;

[0005] Fig. 3 shows a method of tuning a machine learning model in accordance with various examples; and

[0006] Fig. 4 shows a method of modifying a machine learning model based on telemetry data in accordance with various examples.

DETAILED DESCRIPTION

[0007] In preparing a machine learning model for use by customers, a vendor may use initial data to train the machine learning model. The initial data used to train the model may not be fully representative of a customer's circumstances. The customer may have confidential data that would be useful in developing the machine learning model, but the customer may be unwilling to share the confidential data with the vendor. If confidential data from both the customer and vendor were used to train the machine learning model, an improved machine learning model may result.

[0008] To address this problem, the vendor may perform an initial training of the machine learning model, which is to be provided to the customer. The customer may tune the machine learning model based upon the customer's confidential data. The tuned machine learning model may be applied to the customer's data to make predictions for the customer. The resulting tuned machine learning

model is thus based on both the confidential data of the customer and the confidential data of the vendor, without either party disclosing the data to the other. Instead, the resulting machine learning model or predictions made by the machine learning model may be shared. This tuned machine learning model may better predict the operating circumstances of the customer.

[0009] Fig. 1 shows a computer system 100 to update and apply a machine learning model in accordance with various examples. Computer system 100 includes a processor 110 and storage 120. Storage 120 stores update instructions 123 and apply instructions 126. Update instructions 123 and apply instructions 126 may be machine-readable instructions for execution by the processor 110. The update instructions 123 and apply instructions 126 may include instructions to implement the methods described herein, such as the methods described in connection with the other figures. The update instructions 123 may include instructions to update a machine learning model based on confidential data. The machine learning model may have been trained on a different computer system using a different data set. The confidential data may be used to tune the machine learning model. The apply instructions 126 may include instructions to apply the updated machine learning model to a data set.

[0010] The processor 110 may comprise a microprocessor, a microcomputer, a microcontroller, a field programmable gate array (FPGA), or discrete logic to execute machine-readable instructions. The storage 120 may include a hard drive, solid state drive (SSD), flash memory, electrically erasable programmable read-only memory (EEPROM), or random access memory (RAM). The processor 110 may be coupled to the storage 120, such as via a bus.

[0011] Fig. 2 shows a method 200 of updating and applying a machine learning model in accordance with various examples. Method 200 includes updating a machine learning model based on a first confidential data set, the machine learning model being previously trained based on a training data set, the training data set being distinct from the first confidential data set, the previous training performed on a first computer system, and the update performed on a second computer system comprising the processor (210). Method 200 includes applying

the machine learning model after the update to a second confidential data set to make a prediction (220).

[0012] A machine learning model is a mathematical model trained on historical data sets with varying levels of human involvement to allow for predictions based on other data sets. In various examples, the machine learning model may include nodes for performing mathematical operations on inputs. A node may accept members of a data set as an input. The node may accept outputs of other nodes, or even its own output, as an input. The node may include a math function that is applied over the received inputs to summarize the inputs as a value, and it may add non-linearity to the model. The node may include a weight, which may indicate an amount of impact that node has on the model. The nodes of the machine learning model may be grouped into layers, with data passing from one layer to the next. For example, one layer may be an input layer, receiving members of the data set, fixed values, or output of other machine learning models as inputs. One layer may be an output layer, providing some prediction, an answer to a question, or another output from the machine learning model. There may be intermediate layers between the input layer and the output layer.

[0013] In various examples, a machine learning model may be tuned after training on an initial data set. The initial training may be performed using one data set, while the tuning may be performed using a separate data set. After the initial training, certain nodes of the machine learning model may be frozen. Freezing the nodes may prevent the alteration of the operations performed by the nodes and any weight of the nodes. Other nodes may be modifiable by the tuning. The weights and other attributes of such nodes may be modified as the machine learning model is tuned. Links between nodes, such as passing the outputs of nodes to be the inputs of other nodes, may be modified by the tuning. Some links between nodes may be frozen and not modifiable by the tuning. Where the nodes are organized into layers, the layers may be frozen. For example, a machine learning model may include an input layer, an output layer, and an intermediate layer between the input layer and the output layer. After the machine learning model is initially trained, the input layer and intermediate layer may be frozen, and the output layer may be modifiable by the tuning.

[0014] In various examples, the machine learning model may be initially trained by one entity and then tuned by a second entity. This may be done where the entities have confidential data sets to use in training the machine learning model. After the initial training, nodes and links between the nodes may be frozen. The tuning may be performed on the modifiable nodes and links. For example, a vendor entity may provide support services for computer systems. The vendor may initially train a machine learning model based on a data set it has gathered across various customers. One customer may have policies that prevent sharing data about the customer's computer systems with the vendor. The customer may consider the data to be confidential. The vendor may initially train a machine learning model and provide the machine learning model to the customer. The customer may tune the machine learning model based on the customer's confidential data. The tuning may enable the machine learning model to more accurately predict issues with the customer's computer system. For example, the vendor's data used to initially train the machine learning model may include a mix of servers, laptops, and desktop computers from the top three computer manufacturers. The customer may not include any servers and include a different mix of laptops and desktop computers. The customer may have computer systems from a small computer manufacturer that is not represented in the vendor's data set. Tuning the machine learning model may result in different predictions based on the differences between the computer systems that created the vendor's data set and the computer systems in actual use by the customer. After the machine learning model is tuned, the customer may execute the machine learning model to provide predictions regarding the customer's computer systems. The predictions may be provided to the vendor to assist the vendor in servicing the customer's computer systems. The initial training of the machine learning model may be performed on a computer system under control of the vendor. The tuning of the machine learning model may be performed on a different computer system under control of the vendor. The customer's computer system may be used to apply the machine learning model to data sets used to make predictions about the customer's computer systems.

[0015] Confidential data includes restrictions about sharing the data. Data may be confidential for various reasons. Data may be confidential as it was generated internally to a company and the company considers the data to provide it some kind of benefit. For example, data gathered regarding a company's computer systems may be confidential as it indicates the model or brand of computer system the company tends to purchase. The data may be confidential as it indicates the applications used in developing a company's products or services. Data may be confidential due to a privacy aspect. For example, video data collected for facial recognition of employees may indicate the hours worked by the employees. The video data may be affected by laws or regulations regarding privacy of the employees. Data may be confidential due to a contract or a court order.

[0016] The data set used to initially train the machine learning model may be distinct from the data set used to tune the model. The data sets may be distinct, even if there is some overlap of the data, as the data sets may still be different. For example, tuning of the data set may include both confidential data as well as non-confidential data. A customer may share some data with a vendor, or the vendor may share some data with the customer. Both the customer and the vendor may include that data in the data sets used to train and to tune the machine learning model.

[0017] In various examples, the machine learning model may be transferred between computer systems after the machine learning model has been initially trained and before the machine learning model is tuned. This transfer may be via a network connection, such as via Bluetooth or over the Internet. The network connection may be via a secure connection. The connection may be made secure in various ways, such as use of a password or personal identification number (PIN), use of a physical object such as a magnetic card, sending credentials to a specified cell phone number, use of biometrics, encrypting of the data stream, or limiting connections to a set of specified network addresses. The transfer may be via a removable computer-readable medium. A removable computer-readable medium allows for removal of the storage medium without powering down the computer system. For example, a universal serial bus (USB)

memory stick is a removable computer-readable medium. Removable computer-readable medium may include media such as an optical disk, magnetic tape, or a secure digital (SD) card.

[0018] In various examples, a new machine learning model may be initially trained and provided to customers for tuning. The customer may compare predictions by the new machine learning model against predictions by the previous tuned machine learning model and determine whether or not to update to the new machine learning model. The customer may tune the new machine learning model before making such a comparison. For example, a new machine learning model regarding maintenance of computer systems may be trained to include information for newer computer systems sold by manufacturers. Depending on whether the customer uses newer or older computer systems, the new machine learning model may make more accurate predictions than the older machine learning model.

[0019] Fig. 3 shows a method 300 of tuning a machine learning model in accordance with various examples. Method 300 includes tuning a machine learning model into a tuned machine learning model via a first processor based on a first data set, the machine-learning model being previously trained by a second processor based on a second data set, the first data set being distinct from the second data set (310). Method 300 includes applying the tuned machine learning model to a third data set to make a prediction (320). Method 300 includes causing a corrective action to be performed based on the prediction (330).

[0020] A corrective action may be an action to be taken due to an issue identified by the machine learning model. For example, the machine learning model may predict hard drive failures based on data collected from hard drives. The corrective action may be replacement of a hard drive that is showing signs of imminent failure.

[0021] In various examples, a corrective action may include maintenance of a computer system. The maintenance may include replacement or repair of a component of the computer system. The maintenance may include running

diagnostics on the computer system or modifying the settings of the operating system or applications installed on the computer system.

[0022] In various examples the data set used to initially train the machine learning model may be non-overlapping with the data set used to tune the machine learning model. For example, the data sets may be collected by two entities from different sets of devices. The two entities may not share data from the data sets with the other. This may, for example, occur when both entities consider their data sets to be confidential.

[0023] Fig. 4 shows a method 400 of modifying a machine learning model based on telemetry data in accordance with various examples. Method 400 includes modifying a machine learning model based on telemetry data collected regarding a first set of computers, the machine learning model being previously created based on data collected regarding a second set of computers, the second set of computers being non-overlapping with the first set of computers (410). Method 400 includes using the machine learning model after the modification to make a prediction regarding the first set of computers (420). Method 400 includes causing a corrective action to be taken with regard to a computer in the first set of computers based on the prediction (430).

[0024] In various examples, the machine learning model may be used to support or maintain a set of computer systems. Telemetry data may be collected regarding those computers. The telemetry data may include model and manufacturer information, versions of the operating system and applications installed on the computer system, and measurements of components of the computer system, such as processor utilization, processor temperature, fan speed, storage data, and battery data. The battery data may include information about a charge and temperature of the battery. Based on the telemetry data, the machine learning model may make a prediction regarding the battery of a specific computer system. Based on the prediction, a corrective action may include replacement of the battery. The storage data may include information about a temperature, capacity, unused capacity, and age of the storage. Based on the telemetry data, the machine learning model may make a prediction regarding the

storage of a specific computer system. Based on the prediction, a corrective action may include replacement of the storage.

[0025] In various examples, the computer system tuning and applying the machine learning model may be a secured computer system, due to its use of confidential data. The secured computer system may be connected to a network, but limit its communications to a known set of network addresses. The set of network addresses may be stored in boot-level memory coupled to the processor. For example, they may be stored on an EEPROM on a motherboard of the computer system. When a new machine learning model is available, the secured computer system may retrieve the new machine learning model from one of the allowed network address locations.

[0026] Allowing an initial training of the machine learning model, followed by tuning on a separate data set, allows an entity to maintain the confidentiality of its data. The entity may also receive the benefit of training using a large data set and also receive the benefit of tuning based on data specific to the entity. For example, an entity may not have enough confidential data to perform adequate training of a machine learning model, but it may be sufficient for tuning a trained machine learning model. Different entities may thus preserve their sets of confidential data, but still benefit from a machine learning model trained and updated using both sets of data. Thus, benefits comparable to sharing of the confidential data may be achieved while still preserving full confidentiality.

[0027] The above discussion is meant to be illustrative of the principles and various examples of the present disclosure. Numerous variations and modifications will become apparent to those skilled in the art once the above disclosure is fully appreciated. It is intended that the following claims be interpreted to embrace all such variations and modifications.

CLAIMS

What is claimed is:

1. A non-transitory computer-readable medium to store machine-readable instructions that, when executed by a processor, cause the processor to:

update a machine learning model based on a first confidential data set, the machine learning model being previously trained based on a training data set, the training data set being distinct from the first confidential data set, the previous training performed on a first computer system, and the update performed on a second computer system comprising the processor; and

apply the machine learning model after the update to a second confidential data set to make a prediction.

2. The computer-readable medium of claim 1, wherein execution of the machine-readable instructions by the processor causes the processor to:

create a second updated model based on the machine learning model after the update and based on the second confidential data set; and

replace the machine learning model with the second updated model.

3. The computer-readable medium of claim 1, wherein the machine learning model includes an input layer, an intermediate layer, and an output layer, and wherein the machine learning model after the update includes the input layer, the intermediate layer, and an updated output layer, the updated output layer based on the output layer and the first confidential data set.

4. The computer-readable medium of claim 1, wherein execution of the machine-readable instructions by the processor causes the processor to cause repair of a device based on the prediction, the device being a source of the second confidential data set.

5. The computer-readable medium of claim 1, wherein execution of the machine-readable instructions by the processor causes the processor to:

receive a new machine learning model via a removable computer-readable medium, the new machine learning model being previously trained based on a second training data set;

replace the machine learning model with the new machine learning model; and

update the new machine learning model based on a third confidential data set.

6. A method comprising:

tuning a machine learning model into a tuned machine learning model via a first processor based on a first data set, the machine-learning model being previously trained by a second processor based on a second data set, the first data set being distinct from the second data set;

applying the tuned machine learning model to a third data set to make a prediction; and

causing a corrective action to be performed based on the prediction.

7. The method of claim 6, wherein the first data set is non-overlapping with the second data set.

8. The method of claim 6, wherein the machine learning model comprises a first set of nodes with frozen weights and a second set of nodes with modifiable weights, and wherein the tuning comprises modifying the modifiable weights of the second set of nodes.

9. The method of claim 6, wherein the second data set includes data regarding a first computer system and a second computer system, and the corrective action regards the first computer system.

10. The method of claim 9, wherein the corrective action includes replacement of a component of the first computer system.

11. A non-transitory computer-readable medium to store machine-readable instructions that, when executed by a processor, cause the processor to:

modify a machine learning model based on telemetry data collected regarding a first set of computers, the machine learning model being previously created based on data collected regarding a second set of computers, the second set of computers being non-overlapping with the first set of computers;

use the machine learning model after the modification to make a prediction regarding the first set of computers; and

cause a corrective action to be taken with regard to a computer in the first set of computers based on the prediction.

12. The computer-readable medium of claim 11, wherein the telemetry data includes data regarding a storage of the computer.

13. The computer-readable medium of claim 12, wherein the corrective action includes replacing the storage of the computer.

14. The computer-readable medium of claim 11, wherein the processor is coupled to a network, and wherein communication of the processor with the network is limited to communication with a set of network addresses, the set of network addresses being stored in a boot-level memory coupled to the processor.

15. The computer-readable medium of claim 14, wherein execution of the machine-readable instructions by the processor causes the processor to:

receive a new machine learning model via a network address in the set of network addresses;

replace the machine learning model with the new machine learning model; and

modify the new machine learning model based on the telemetry data.

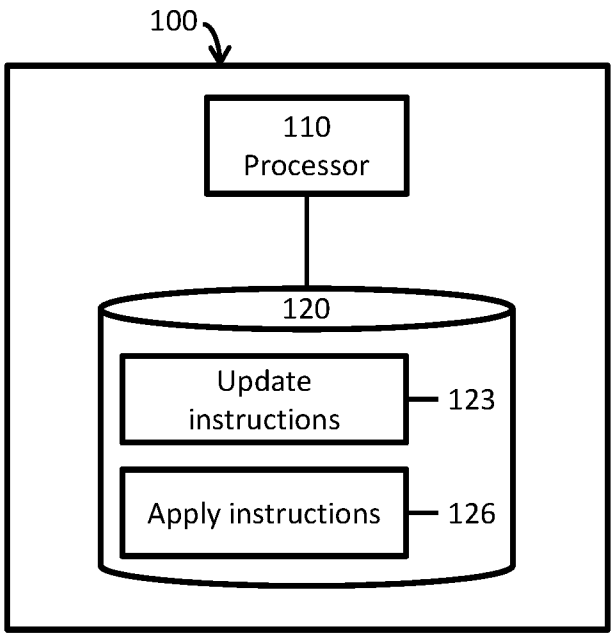
100
110
Processor
120
Update
instructions
123
Apply instructions
126

Fig. 1

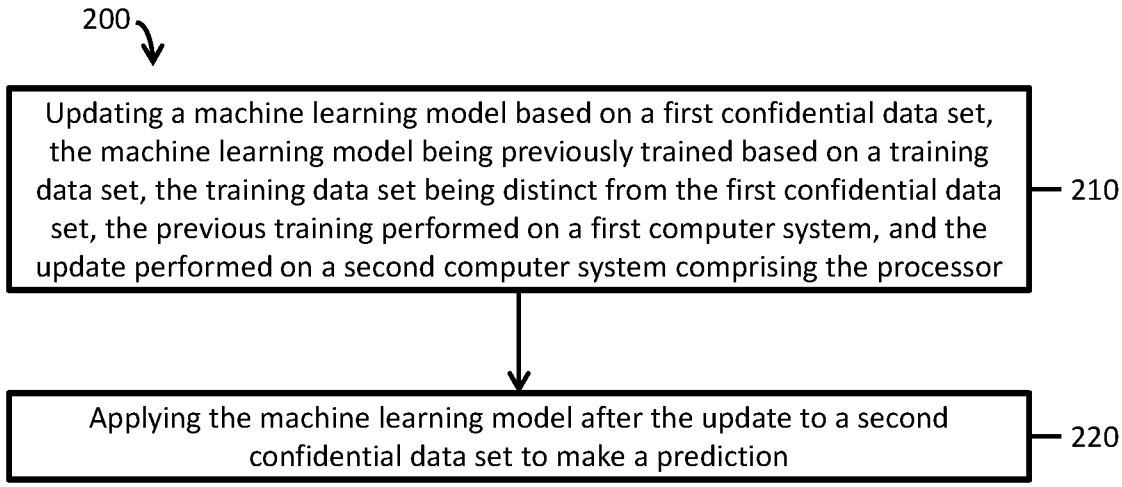
200
Updating a machine learning model based on a first confidential data set, the machine learning model being previously trained based on a training data set, the training data set being distinct from the first confidential data set, the previous training performed on a first computer system, and the update performed on a second computer system comprising the processor
210
Applying the machine learning model after the update to a second confidential data set to make a prediction
220

Fig. 2

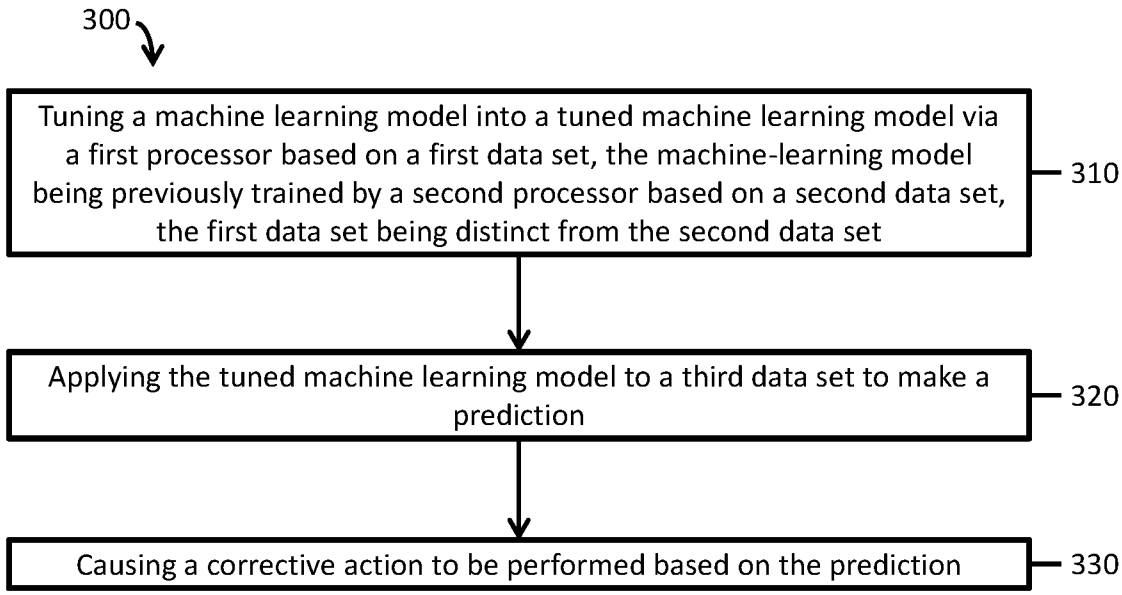
300
Tuning a machine learning model into a tuned machine learning model via a first processor based on a first data set, the machine-learning model being previously trained by a second processor based on a second data set, the first data set being distinct from the second data set
310
Applying the tuned machine learning model to a third data set to make a prediction
320
Causing a corrective action to be performed based on the prediction
330

Fig. 3

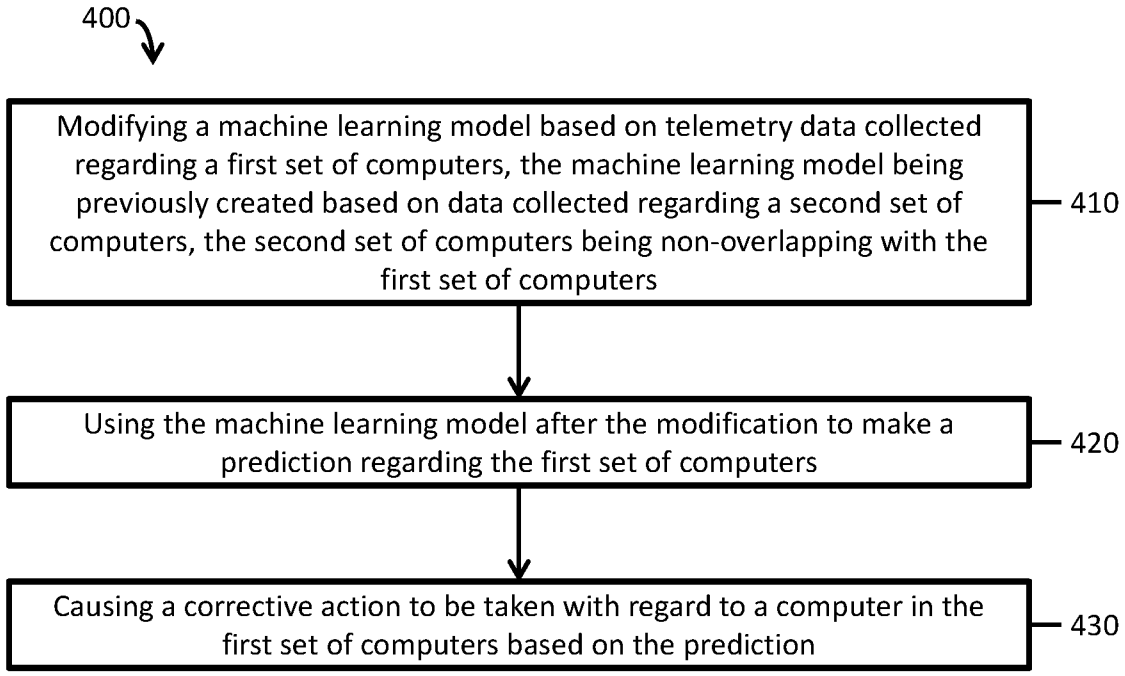
400
Modifying a machine learning model based on telemetry data collected regarding a first set of computers, the machine learning model being previously created based on data collected regarding a second set of computers, the second set of computers being non-overlapping with the first set of computers
410
Using the machine learning model after the modification to make a prediction regarding the first set of computers
420
Causing a corrective action to be taken with regard to a computer in the first set of computers based on the prediction
430

Fig. 4

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US 2018/065255

A. CLASSIFICATION OF SUBJECT MATTER

G06N 20/00 *(2019.01)*
G06F 21/62 *(2013.01)*

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)

G06N 20/00, 99/00, 3/00-3/12, G06F 21/00-21/62, 3/00-3/0484, H04L 29/00-29/08

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)

PatSearch (RUPTO Internal), USPTO, PAJ, Espacenet, Information Retrieval System of FIPS

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2015/0379429 A1 (AMAZON TECHNOLOGIES, INC.) 31.12.2015, abstract, paragraphs [0079] - [0081], [0084], [0092], [0098], [0110], [0113], [0115], [0119] - [0120], [0204], [0209], [0234], [0255], [0258], [0291] - [0293], [0323], [0341] - [0343]	1-3, 5-8
Y		4, 9-15
Y	US 2017/0006135 A1 (C3, INC.) 05.01.2017, abstract, paragraphs [0041], [0054], [0152], [0167], [0490], [0525] - [0528], [0538], [0543], [0609] - [0610]	4, 9-15
A	US 2017/0372226 A1 (MICROSOFT TECHNOLOGY LICENSING, LLC) 28.12.2017	1-15

☐ 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 document 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	Date of mailing of the international search report
12 August 2019 (12.08.2019)	22 August 2019 (22.08.2019)

Name and mailing address of the ISA/RU:	Authorized officer
Federal Institute of Industrial Property, Berezhkovskaya nab., 30-1, Moscow, G-59, GSP-3, Russia, 125993 Facsimile No: (8-495) 531-63-18, (8-499) 243-33-37	A. Tokarev Telephone No. 8(495) 531-64-81