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

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
22 December 2016 (22.12.2016)

WIPO | PCT

(10) International Publication Number
WO 2016/205132 A1

(51) International Patent Classification:
G06Q 10/06 (2012.01) *G06Q 50/04* (2012.01)

(21) International Application Number:
PCT/US2016/037247

(22) International Filing Date:
13 June 2016 (13.06.2016)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
14/744,352 19 June 2015 (19.06.2015) US
14/744,369 19 June 2015 (19.06.2015) US
14/963,207 8 December 2015 (08.12.2015) US

(71) Applicant: **UPTAKE TECHNOLOGIES, INC.** [US/US]; Suite 775, 600 West Chicago Ave., Chicago, Illinois 60654 (US).

(72) Inventors: **NICHOLAS, Brad**; Suite 775, 600 West Chicago Ave., Chicago, Illinois 60654 (US). **KOLB, Jason**; Suite 775, 600 West Chicago Ave., Chicago, Illinois 60654 (US).

(74) Agent: **SHEA, Rory**; Lee Sullivan Shea & Smith LLP, 224 N. Desplaines St., Suite 250, Chicago, Illinois 60654 (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, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, 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).

[Continued on next page]

(54) Title: LOCAL ANALYTICS AT AN ASSET

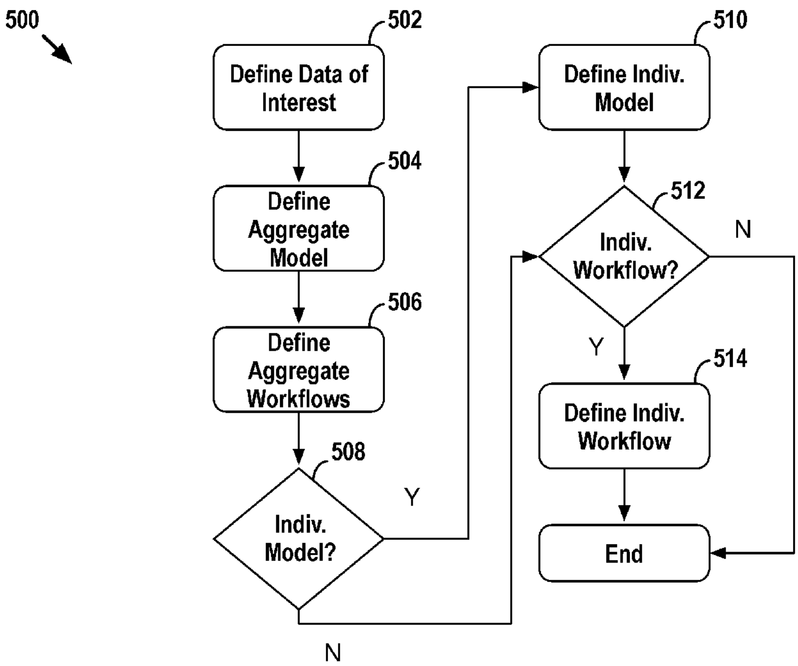

FIGURE 5

(57) Abstract: Disclosed herein are systems, devices, and methods related to assets and predictive models and corresponding workflows that are related to the operation of assets. In particular, examples involve defining and deploying aggregate, predictive models and corresponding workflows, defining and deploying individualized, predictive models and/or corresponding workflows, and dynamically adjusting the execution of model-workflow pairs. Additionally, examples involve assets configured to receive and locally execute predictive models, locally individualize predictive models, and/or locally execute workflows or portions thereof.

WO 2016/205132 A1

Published:

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

LOCAL ANALYTICS AT AN ASSET

CROSS-REFERENCE TO RELATED APPLICATIONS

This application claims priority to (i) U.S. Non-Provisional Patent Application No. 14/744,352, filed on June 19, 2015 and entitled Aggregate Predictive Model & Workflow for Local Execution, (ii) U.S. Non-Provisional Patent Application No. 14/744,369, filed on June 19, 2015 and entitled Individualized Predictive Model & Workflow for an Asset, and (iii) U.S. Non-Provisional Patent Application No. 14/963,207, filed on December 8, 2015 and entitled Local Analytics at an Asset, each of which is herein incorporated by reference in its entirety. This application also incorporates by reference U.S. Non-Provisional Patent Application No. 14/732,258, filed on June 5, 2015 and entitled Asset Health Score, in its entirety.

BACKGROUND

Today, machines (also referred to herein as “assets”) are ubiquitous in many industries. From locomotives that transfer cargo across countries to medical equipment that helps nurses and doctors to save lives, assets serve an important role in everyday life. Depending on the role that an asset serves, its complexity, and cost, may vary. For instance, some assets may include multiple subsystems that must operate in harmony for the asset to function properly (e.g., an engine, transmission, etc. of a locomotive).

Because of the key role that assets play in everyday life, it is desirable for assets to be repairable with limited downtime. Accordingly, some have developed mechanisms to monitor and detect abnormal conditions within an asset to facilitate repairing the asset, perhaps with minimal downtime.

OVERVIEW

The current approach for monitoring assets generally involves an on-asset computer that receives signals from various sensors and/or actuators distributed throughout an asset that monitor the operating conditions of the asset. As one representative example, if the asset is a locomotive, the sensors and/or actuators may monitor parameters such as temperatures, voltages, and speeds, among other examples. If sensor and/or actuator signals from one or more of these devices reach certain values, the on-asset computer may then generate an abnormal-condition indicator, such as a “fault code,” which is an indication that an abnormal condition has occurred within the asset.

In general, an abnormal condition may be a defect at an asset or component thereof, which may lead to a failure of the asset and/or component. As such, an abnormal condition may be associated with a given failure, or perhaps multiple failures, in that the abnormal condition is symptomatic of the given failure or failures. In practice, a user typically defines the sensors and respective sensor values associated with each abnormal-condition indicator. That is, the user defines an asset's "normal" operating conditions (e.g., those that do not trigger fault codes) and "abnormal" operating conditions (e.g., those that trigger fault codes).

After the on-asset computer generates an abnormal-condition indicator, the indicator and/or sensor signals may be passed to a remote location where a user may receive some indication of the abnormal condition and/or sensor signals and decide whether to take action. One action that the user might take is to assign a mechanic or the like to evaluate and potentially repair the asset. Once at the asset, the mechanic may connect a computing device to the asset and operate the computing device to cause the asset to utilize one or more local diagnostic tools to facilitate diagnosing the cause of the generated indicator.

While current asset-monitoring systems are generally effective at triggering abnormal-condition indicators, such systems are typically reactionary. That is, by the time the asset-monitoring system triggers an indicator, a failure within the asset may have already occurred (or is about to occur), which may lead to costly downtime, among other disadvantages. Additionally, due to the simplistic nature of on-asset abnormality-detection mechanisms in such asset-monitoring systems, current asset-monitoring approaches tend to involve a remote computing system performing monitoring computations for an asset and then transmitting instructions to the asset if a problem is detected. This may be disadvantageous due to network latency and/or infeasible when the asset moves outside of coverage of a communication network. Further still, due to the nature of local diagnostic tools stored on assets, current diagnosis procedures tend to be inefficient and cumbersome because a mechanic is required to cause the asset to utilize such tools.

The example systems, devices, and methods disclosed herein seek to help address one or more of these issues. In example implementations, a network configuration may include a communication network that facilitates communications between assets and a remote computing system. In some cases, the communication network may facilitate secure communications between assets and the remote computing system (e.g., via encryption or other security measures).

As noted above, each asset may include multiple sensors and/or actuators distributed throughout the asset that facilitate monitoring operating conditions of the asset. A number of assets may provide respective data indicative of each asset's operating conditions to the remote computing system, which may be configured to perform one or more operations based on the provided data. Typically, sensor and/or actuator data may be utilized for general asset-monitoring operations. However, as described herein, the remote computing system and/or assets may leverage such data to facilitate performing more complex operations.

In example implementations, the remote computing system may be configured to define and deploy to assets a predictive model and corresponding workflow (referred to herein as a “model-workflow pair”) that are related to the operation of the assets. The assets may be configured to receive the model-workflow pair and utilize a local analytics device to operate in accordance with the model-workflow pair.

Generally, a model-workflow pair may cause an asset to monitor certain operating conditions and when certain conditions exist, modify a behavior that may help facilitate preventing an occurrence of a particular event. Specifically, a predictive model may receive as inputs data from a particular set of asset sensors and/or actuators and output a likelihood that one or more particular events could occur at the asset within a particular period of time in the future. A workflow may involve one or more operations that are performed based on the likelihood of the one or more particular events that is output by the model.

In practice, the remote computing system may define an aggregate, predictive model and corresponding workflows, individualized, predictive models and corresponding workflows, or some combination thereof. An “aggregate” model/workflow may refer to a model/workflow that is generic for a group of assets, while an “individualized” model/workflow may refer to a model/workflow that is tailored for a single asset or subgroup of assets from the group of assts.

In example implementations, the remote computing system may start by defining an aggregate, predictive model based on historical data for multiple assets. Utilizing data for multiple assets may facilitate defining a more accurate predictive model than utilizing operating data for a single asset.

The historical data that forms the basis of the aggregate model may include at least operating data that indicates operating conditions of a given asset. Specifically, operating data may include abnormal-condition data identifying instances when failures occurred at assets and/or data indicating one or more physical properties measured at the assets at the time of those instances. The data may also include environment data indicating environments in which assets have been operated and scheduling data indicating dates and times when assets were utilized, among other examples of asset-related data used to define the aggregate model-workflow pair.

Based on the historical data, the remote computing system may define an aggregate model that predicts the occurrence of particular events. In a particular example implementation, an aggregate model may output a probability that a failure will occur at an asset within a particular period of time in the future. Such a model may be referred to herein as a “failure model.” Other aggregate models may predict the likelihood that an asset will complete a task within a particular period of time in the future, among other example predictive models.

After defining the aggregate model, the remote computing system may then define an aggregate workflow that corresponds to the defined aggregate model. Generally, a workflow may include one or more operations that an asset may perform based on a corresponding model. That is, the output of the corresponding model may cause the asset to perform workflow operations. For instance, an aggregate model-workflow pair may be defined such that when the aggregate model outputs a probability within a particular range an asset will execute a particular workflow operation, such as a local diagnostic tool.

After the aggregate model-workflow pair is defined, the remote computing system may transmit the pair to one or more assets. The one or more assets may then operate in accordance with the aggregate model-workflow pair.

In example implementations, the remote computing system may be configured to further define an individualized predictive model and/or corresponding workflow for one or multiple assets. The remote computing system may do so based on certain characteristics of each given asset, among other considerations. In example implementations, the remote computing system may start with an aggregate model-workflow pair as a baseline and individualize one or both of the aggregate model and workflow for the given asset based on the asset's characteristics.

In practice, the remote computing system may be configured to determine asset characteristics that are related to the aggregate model-workflow pair (e.g., characteristics of interest). Examples of such characteristics may include asset age, asset usage, asset class (e.g., brand and/or model), asset health, and environment in which an asset is operated, among other characteristics.

Then, the remote computing system may determine characteristics of the given asset that correspond to the characteristics of interest. Based at least on some of the given asset's characteristics, the remote computing system may be configured to individualize the aggregate model and/or corresponding workflow.

Defining an individualized model and/or workflow may involve the remote computing system making certain modifications to the aggregate model and/or workflow. For example, individualizing the aggregate model may involve changing model inputs, changing a model calculation, and/or changing a weight of a variable or output of a calculation, among other examples. Individualizing the aggregate workflow may involve changing one or more operations of the workflow and/or changing the model output value or range of values that triggers the workflow, among other examples.

After defining an individualized model and/or workflow for the given asset, the remote computing system may then transmit the individualized model and/or workflow to the given asset. In a scenario where only one of the model or workflow is individualized, the given asset may utilize the aggregate version of the model or workflow that is not individualized. The given asset may then operate in accordance with its individualized model-workflow pair.

In example implementations, a given asset may include a local analytics device that may be configured to cause the given asset to operate in accordance with a model-workflow pair provided by the remote computing system. The local analytics device may be configured to utilize operating data from the asset sensors and/or actuators (e.g., data that is typically utilized for other asset-related purposes) to run the predictive model. When the local analytics device receives certain operating data, it may execute the model and depending on the output of the model, may execute the corresponding workflow.

Executing the corresponding workflow may help facilitate preventing an undesirable event from occurring at the given asset. In this way, the given asset may locally determine that an occurrence of a particular event is likely and may then execute a particular workflow to help prevent the occurrence of the event. This may be particularly useful if communication between the given asset and remote computing system is hindered. For example, in some situations, a failure might occur before a command to take preventative actions reaches the given asset from the remote computing system. In such situations, the local analytics device may be advantageous in that it may generate the command locally, thereby avoiding any network latency or any issues arising from the given asset being "off-line." As such, the local analytics device executing a model-workflow pair may facilitate causing the asset to adapt to its conditions.

In some example implementations, before or when first executing a model-workflow pair, the local analytics device may itself individualize the model-workflow pair that it received from the remote computing system. Generally, the local analytics device may individualize the model-workflow pair by evaluating some or all predictions, assumptions, and/or generalizations related to the given asset that were made when the model-workflow pair was defined. Based on the evaluation, the local analytics device may modify the model-workflow pair so that the underlying predictions, assumptions, and/or generalizations of the model-workflow pair more accurately reflect the actual state of the given asset. The local analytics device may then execute the individualized model-workflow pair instead of the model-workflow pair it originally received from the remote computing system, which may result in more accurate monitoring of the asset.

While a given asset is operating in accordance with a model-workflow pair, the given asset may also continue to provide operating data to the remote computing system. Based at least on this data, the remote computing system may modify the aggregate model-workflow pair and/or one or more individualized model-workflow pairs. The remote computing system may make modifications for a number of reasons.

In one example, the remote computing system may modify a model and/or workflow if a new event occurred at an asset that the model did not previously account for. For instance, in a failure model, the new event may be a new failure that had yet to occur at any of the assets whose data was used to define the aggregate model.

In another example, the remote computing system may modify a model and/or workflow if an event occurred at an asset under operating conditions that typically do not cause the event to occur. For instance, returning again to a failure model, the failure model or corresponding workflow may be modified if a failure occurred under operating conditions that had yet to cause the failure to occur in the past.

In yet another example, the remote computing system may modify a model and/or workflow if an executed workflow failed to prevent an occurrence of an event. Specifically, the remote computing system may modify the model and/or workflow if the output of the model caused an asset to execute a workflow aimed to prevent the occurrence of an event but the event occurred at the asset nonetheless. Other examples of reasons for modifying a model and/or workflow are also possible.

The remote computing system may then distribute any modifications to the asset whose data caused the modification and/or to other assets in communication with the remote computing system. In this way, the remote computing system may dynamically modify models and/or workflows and distribute these modifications to a whole fleet of assets based on operating conditions of an individual asset.

In some example implementations, an asset and/or the remote computing system may be configured to dynamically adjust executing a predictive model and/or workflow. In particular, the asset and/or remote computing system may be configured to detect certain events that trigger a change in responsibilities with respect to whether the asset and/or the remote computing system are executing a predictive model and/or workflow.

For instance, in some cases, after the asset receives a model-workflow pair from the remote computing system, the asset may store the model-workflow pair in data storage but then may rely on the remote computing system to centrally execute part or all of the model-workflow pair. On the other hand, in other cases, the remote computing system may rely on the asset to locally execute part or all of the model-workflow pair. In yet other cases, the remote computing system and the asset may share in the responsibilities of executing the model-workflow pair.

In any event, at some point in time, certain events may occur that trigger the asset and/or remote computing system to adjust the execution of the predictive model and/or workflow. For instance, the asset and/or remote computing system may detect certain characteristics of a communication network that couples the asset to the remote computing system. Based on the characteristics of the communication network, the asset may adjust whether it is locally executing a predictive model and/or workflow and the remote computing system may accordingly modify whether it is centrally executing the model and/or workflow. In this way, the asset and/or remote computing system may adapt to conditions of the asset.

In a particular example, the asset may detect an indication that a signal strength of a communication link between the asset and the remote computing system is relatively weak (e.g., the asset may determine that is about to go “off-line”), that a network latency is relatively high, and/or that a network bandwidth is relatively low. Accordingly, the asset may be programmed to take on responsibilities for executing the model-workflow pair that were previously being handled by the remote computing system. In turn, the remote computing system may cease centrally executing some or all of the model-workflow pair. In this way, the asset may locally execute the predictive model and then, based on executing the predictive model, execute the corresponding workflow to potentially help prevent an occurrence of a failure at the asset.

Moreover, in some implementations, the asset and/or the remote computing system may similarly adjust executing (or perhaps modify) a predictive model and/or workflow based on various other considerations. For example, based on the processing capacity of the asset, the asset may adjust locally executing a model-workflow pair and the remote computing system may accordingly adjust as well. In another example, based on the bandwidth of the communication network coupling the asset to the remote computing system, the asset may execute a modified workflow (e.g., transmitting data to the remote computing system according to a data-transmission scheme with a reduced transmission rate). Other examples are also possible.

As discussed above, examples provided herein are related to deployment and execution of predictive models. In one aspect, a computing system is provided. The computing system comprises at least one processor, a non-transitory computer-readable medium, and program instructions stored on the non-transitory computer-readable medium that are executable by the at least one processor to cause the computing system to: (a) receive respective operating data for a plurality of assets, (b) based on the received operating data, define a predictive model and a corresponding workflow that are related to the operation of the plurality of assets, and (c) transmit to at least one asset of the plurality of assets the predictive model and the corresponding workflow for local execution by the at least one asset.

In another aspect, a non-transitory computer-readable medium is provided having instructions stored thereon that are executable to cause a computing system to: (a) receive respective operating data for a plurality of assets, (b) based on the received operating data, define a predictive model and a corresponding workflow that are related to the operation of the plurality of assets, and (c) transmit to at least one asset of the plurality of assets the predictive model and the corresponding workflow for local execution by the at least one asset.

In yet another aspect, a computer-implemented method is provided. The method comprises: (a) receiving respective operating data for a plurality of assets, (b) based on the received operating data, defining a predictive model and a corresponding workflow that are related to the operation of the plurality of assets, and (c) transmitting to at least one asset of the plurality of assets the predictive model and the corresponding workflow for local execution by the at least one asset.

As discussed above, examples provided herein are related to deployment and execution of predictive models. In one aspect, a computing system is provided. The computing system comprises at least one processor, a non-transitory computer-readable medium, and program instructions stored on the non-transitory computer-readable medium that are executable by the at least one processor to cause the computing system to: (a) receive operating data for a plurality of assets, wherein the plurality of assets comprises a first asset, (b) based on the received operating data, define an aggregate predictive model and an aggregate corresponding workflow that are related to the operation of the plurality of assets, (c) determine one or more characteristics of the first asset, (d) based on the one or more characteristics of the first asset and the aggregate predictive model and the aggregate corresponding workflow, define at least one of an individualized predictive model or an individualized corresponding workflow that is related to the operation of the first asset, and (e) transmit to the first asset the defined at least one individualized predictive model or individualized corresponding workflow for local execution by the first asset.

In another aspect, a non-transitory computer-readable medium is provided having instructions stored thereon that are executable to cause a computing system to: (a) receive operating data for a plurality of assets, wherein the plurality of assets comprises a first asset, (b) based on the received operating data, define an aggregate predictive model and an aggregate corresponding workflow that are related to the operation of the plurality of assets, (c) determine one or more characteristics of the first asset, (d) based on the one or more characteristics of the first asset and the aggregate predictive model and the aggregate corresponding workflow, define at least one of an individualized predictive model or an individualized corresponding workflow that is related to the operation of the first asset, and (e) transmit to the first asset the defined at least one individualized predictive model or individualized corresponding workflow for local execution by the first asset.

In yet another aspect, a computer-implemented method is provided. The method comprises: (a) receiving operating data for a plurality of assets, wherein the plurality of assets comprises a first asset, (b) based on the received operating data, defining an aggregate predictive model and an aggregate corresponding workflow that are related to the operation of the plurality of assets, (c) determining one or more characteristics of the first asset, (d) based on the one or more characteristics of the first asset and the aggregate predictive model and the aggregate corresponding workflow, defining at least one of an individualized predictive model or an individualized corresponding workflow that is related to the operation of the first asset, and (e) transmitting to the first asset the defined at least one individualized predictive model or individualized corresponding workflow for local execution by the first asset.

As discussed above, examples provided herein are related to receiving and executing predictive models and/or workflows at an asset. In one aspect, a computing device is provided. The computing device comprises (i) an asset interface configured to couple the computing device to an asset, (ii) a network interface configured to facilitate communication between the computing device and a computing system located remote from the computing device, (iii) at least one processor, (iv) a non-transitory computer-readable medium, and (v) program instructions stored on the non-transitory computer-readable medium that are executable by the at least one processor to cause the computing device to: (a) receive, via the network interface, a predictive model that is related to the operation of the asset, wherein the predictive model is defined by the computing system based on operating data for a plurality of assets, (b) receive, via the asset interface, operating data for the asset, (c) execute the predictive model based on at least a portion of the received operating data for the asset, and (d) based on executing the predictive model, execute a workflow corresponding to the predictive model, wherein executing the workflow comprises causing the asset, via the asset interface, to perform an operation.

In another aspect, a non-transitory computer-readable medium is provided having instructions stored thereon that are executable to cause a computing device coupled to an asset via an asset interface of the computing device to: (a) receive, via a network interface of the computing device configured to facilitate communication between the computing device and a computing system located remote from the computing device, a predictive model that is related to the operation of the asset, wherein the predictive model is defined by the computing system based on operating data for a plurality of assets, (b) receive, via the asset interface, operating data for the asset, (c) execute the predictive model based on at least a portion of the received operating data for the asset, and (c) based on executing the predictive model, execute a workflow corresponding to the predictive model, wherein executing the workflow comprises causing the asset, via the asset interface, to perform an operation.

In yet another aspect, a computer-implemented method is provided. The method comprises: (a) receiving, via a network interface of a computing device that is coupled to an asset via an asset interface of the computing device, a predictive model that is related to the operation of the asset, wherein the predictive model is defined by a computing system located remote from the computing device based on operating data for a plurality of assets, (b) receiving, by the computing device via the asset interface, operating data for the asset, (b) executing, by the computing device, the predictive model based on at least a portion of the received operating data for the asset, and (c) based on executing the predictive model, executing, by the computing device, a workflow corresponding to the predictive model, wherein executing the workflow comprises causing the asset, via the asset interface, to perform an operation.

One of ordinary skill in the art will appreciate these as well as numerous other aspects in reading the following disclosure.

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 depicts an example network configuration in which example embodiments may be implemented.

Figure 2 depicts a simplified block diagram of an example asset.

Figure 3 depicts a conceptual illustration of example abnormal-condition indicators and triggering criteria.

Figure 4 depicts a simplified block diagram of an example analytics system.

Figure 5 depicts an example flow diagram of a definition phase that may be used for defining model-workflow pairs.

Figure 6A depicts a conceptual illustration of an aggregate model-workflow pair.

Figure 6B depicts a conceptual illustration of an individualized model-workflow pair.

Figure 6C depicts a conceptual illustration of another individualized model-workflow pair.

Figure 6D depicts a conceptual illustration of a modified model-workflow pair.

Figure 7 depicts an example flow diagram of a modeling phase that may be used for defining a predictive model that outputs a health metric.

Figure 8 depicts a conceptual illustration of data utilized to define a model.

Figure 9 depicts an example flow diagram of a local-execution phase that may be used for locally executing a predictive model.

Figure 10 depicts an example flow diagram of a modification phase that may be used for modifying model-workflow pairs.

Figure 11 depicts an example flow diagram of an adjustment phase that may be used for adjusting execution of model-workflow pairs.

Figure 12 depicts a flow diagram of an example method for defining and deploying an aggregate, predictive model and corresponding workflow

Figure 13 depicts a flow diagram of an example method for defining and deploying an individualized, predictive model and/or corresponding workflow

Figure 14 depicts a flow diagram of an example method for dynamically modifying the execution of model-workflow pairs.

Figure 15 depicts a flow diagram of an example method for receiving and locally executing a model-workflow pair.

DETAILED DESCRIPTION

The following disclosure makes reference to the accompanying figures and several exemplary scenarios. One of ordinary skill in the art will understand that such references are for the purpose of explanation only and are therefore not meant to be limiting. Part or all of the disclosed systems, devices, and methods may be rearranged, combined, added to, and/or removed in a variety of manners, each of which is contemplated herein.

I. EXAMPLE NETWORK CONFIGURATION

Turning now to the figures, Figure 1 depicts an example network configuration 100 in which example embodiments may be implemented. As shown, the network configuration 100 includes an asset 102, an asset 104, a communication network 106, a remote computing system 108 that may take the form of an analytics system, an output system 110, and a data source 112.

The communication network 106 may communicatively connect each of the components in the network configuration 100. For instance, the assets 102 and 104 may communicate with the analytics system 108 via the communication network 106. In some cases, the assets 102 and 104 may communicate with one or more intermediary systems, such as an asset gateway (not pictured), that in turn communicates with the analytics system 108. Likewise, the analytics system 108 may communicate with the output system 110 via the communication network 106. In some cases, the analytics system 108 may communicate with one or more intermediary systems, such as a host server (not pictured), that in turn communicates with the output system 110. Many other configurations are also possible. In example cases, the communication network 106 may facilitate secure communications between network components (e.g., via encryption or other security measures).

In general, the assets 102 and 104 may take the form of any device configured to perform one or more operations (which may be defined based on the field) and may also include equipment configured to transmit data indicative of one or more operating conditions of the given asset. In some examples, an asset may include one or more subsystems configured to perform one or more respective operations. In practice, multiple subsystems may operate in parallel or sequentially in order for an asset to operate.

Example assets may include transportation machines (e.g., locomotives, aircraft, passenger vehicles, semi-trailer trucks, ships, etc.), industrial machines (e.g., mining equipment, construction equipment, factory automation, etc.), medical machines (e.g., medical imaging equipment, surgical equipment, medical monitoring systems, medical laboratory equipment, etc.), and utility machines (e.g., turbines, solar farms, etc.), among other examples. Those of ordinary skill in the art will appreciate that these are but a few examples of assets and that numerous others are possible and contemplated herein.

In example implementations, the assets 102 and 104 may each be of the same type (e.g., a fleet of locomotives or aircrafts, a group of wind turbines, or a set of MRI machines, among other examples) and perhaps may be of the same class (e.g., same brand and/or model). In other examples, the assets 102 and 104 may differ by type, by brand, by model, etc. The assets are discussed in further detail below with reference to Figure 2.

As shown, the assets 102 and 104, and perhaps the data source 112, may communicate with the analytics system 108 via the communication network 106. In general, the communication network 106 may include one or more computing systems and network infrastructure configured to facilitate transferring data between network components. The communication network 106 may be or may include one or more Wide-Area Networks (WANs) and/or Local-Area Networks (LANs), which may be wired and/or wireless and support secure communication. In some examples, the communication network 106 may include one or more cellular networks and/or the Internet, among other networks. The communication network 106 may operate according to one or more communication protocols, such as LTE, CDMA, GSM, LPWAN, WiFi, Bluetooth, Ethernet, HTTP/S, TCP, CoAP/DTLS and the like. Although the communication network 106 is shown as a single network, it should be understood that the communication network 106 may include multiple, distinct networks that are themselves communicatively linked. The communication network 106 could take other forms as well.

As noted above, the analytics system 108 may be configured to receive data from the assets 102 and 104 and the data source 112. Broadly speaking, the analytics system 108 may include one or more computing systems, such as servers and databases, configured to receive, process, analyze, and output data. The analytics system 108 may be configured according to a given dataflow technology, such as TPL Dataflow or NiFi, among other examples. The analytics system 108 is discussed in further detail below with reference to Figure 3.

As shown, the analytics system 108 may be configured to transmit data to the assets 102 and 104 and/or to the output system 110. The particular data transmitted may take various forms and will be described in further detail below.

In general, the output system 110 may take the form of a computing system or device configured to receive data and provide some form of output. The output system 110 may take various forms. In one example, the output system 110 may be or include an output device configured to receive data and provide an audible, visual, and/or tactile output in response to the data. In general, an output device may include one or more input interfaces configured to receive user input, and the output device may be configured to transmit data through the communication network 106 based on such user input. Examples of output devices include tablets, smartphones, laptop computers, other mobile computing devices, desktop computers, smart TVs, and the like.

Another example of the output system 110 may take the form of a work-order system configured to output a request for a mechanic or the like to repair an asset. Yet another example of the output system 110 may take the form of a parts-ordering system configured to place an order for a part of an asset and output a receipt thereof. Numerous other output systems are also possible.

The data source 112 may be configured to communicate with the analytics system 108. In general, the data source 112 may be or include one or more computing systems configured to collect, store, and/or provide to other systems, such as the analytics system 108, data that may be relevant to the functions performed by the analytics system 108. The data source 112 may be configured to generate and/or obtain data independently from the assets 102 and 104. As such, the data provided by the data source 112 may be referred to herein as "external data." The data source 112 may be configured to provide current and/or historical data. In practice, the analytics system 108 may receive data from the data source 112 by "subscribing" to a service provided by the data source. However, the analytics system 108 may receive data from the data source 112 in other manners as well.

Examples of the data source 112 include environment data sources, asset-management data sources, and other data sources. In general, environment data sources provide data indicating some characteristic of the environment in which assets are operated. Examples of environment data sources include weather-data servers, global navigation satellite systems (GNSS) servers, map-data servers, and topography-data servers that provide information regarding natural and artificial features of a given area, among other examples.

In general, asset-management data sources provide data indicating events or statuses of entities (e.g., other assets) that may affect the operation or maintenance of assets (e.g., when and where an asset may operate or receive maintenance). Examples of asset-management data sources include traffic-data servers that provide information regarding air, water, and/or ground traffic, asset-schedule servers that provide information regarding expected routes and/or locations of assets on particular dates and/or at particular times, defect detector systems (also known as "hotbox" detectors) that provide information regarding one or more operating conditions of an asset that passes in proximity to the defect detector system, part-supplier servers that provide information regarding parts that particular suppliers have in stock and prices thereof, and repair-shop servers that provide information regarding repair shop capacity and the like, among other examples.

Examples of other data sources include power-grid servers that provide information regarding electricity consumption and external databases that store historical operating data for assets, among other examples. One of ordinary skill in the art will appreciate that these are but a few examples of data sources and that numerous others are possible.

It should be understood that the network configuration 100 is one example of a network in which embodiments described herein may be implemented. Numerous other arrangements are possible and contemplated herein. For instance, other network configurations may include additional components not pictured and/or more or less of the pictured components.

II. EXAMPLE ASSET

Turning to Figure 2, a simplified block diagram of an example asset 200 is depicted. Either or both of assets 102 and 104 from Figure 1 may be configured like the asset 200. As shown, the asset 200 may include one or more subsystems 202, one or more sensors 204, one or more actuators 205, a central processing unit 206, data storage 208, a network interface 210, a user interface 212, and a local analytics device 220, all of which may be communicatively linked (either directly or indirectly) by a system bus, network, or other connection mechanism. One of ordinary skill in the art will appreciate that the asset 200 may include additional components not shown and/or more or less of the depicted components.

Broadly speaking, the asset 200 may include one or more electrical, mechanical, and/or electromechanical components configured to perform one or more operations. In some cases, one or more components may be grouped into a given subsystem 202.

Generally, a subsystem 202 may include a group of related components that are part of the asset 200. A single subsystem 202 may independently perform one or more operations or the single subsystem 202 may operate along with one or more other subsystems to perform one or more operations. Typically, different types of assets, and even different classes of the same type of assets, may include different subsystems.

For instance, in the context of transportation assets, examples of subsystems 202 may include engines, transmissions, drivetrains, fuel systems, battery systems, exhaust systems, braking systems, electrical systems, signal processing systems, generators, gear boxes, rotors, and hydraulic systems, among numerous other subsystems. In the context of a medical machine, examples of subsystems 202 may include scanning systems, motors, coil and/or magnet systems, signal processing systems, rotors, and electrical systems, among numerous other subsystems.

As suggested above, the asset 200 may be outfitted with various sensors 204 that are configured to monitor operating conditions of the asset 200 and various actuators 205 that are configured to interact with the asset 200 or a component thereof and monitor operating conditions of the asset 200. In some cases, some of the sensors 204 and/or actuators 205 may be grouped based on a particular subsystem 202. In this way, the group of sensors 204 and/or actuators 205 may be configured to monitor operating conditions of the particular subsystem

202, and the actuators from that group may be configured to interact with the particular subsystem 202 in some way that may alter the subsystem's behavior based on those operating conditions.

In general, a sensor 204 may be configured to detect a physical property, which may be indicative of one or more operating conditions of the asset 200, and provide an indication, such as an electrical signal, of the detected physical property. In operation, the sensors 204 may be configured to obtain measurements continuously, periodically (e.g., based on a sampling frequency), and/or in response to some triggering event. In some examples, the sensors 204 may be preconfigured with operating parameters for performing measurements and/or may perform measurements in accordance with operating parameters provided by the central processing unit 206 (e.g., sampling signals that instruct the sensors 204 to obtain measurements). In examples, different sensors 204 may have different operating parameters (e.g., some sensors may sample based on a first frequency, while other sensors sample based on a second, different frequency). In any event, the sensors 204 may be configured to transmit electrical signals indicative of a measured physical property to the central processing unit 206. The sensors 204 may continuously or periodically provide such signals to the central processing unit 206.

For instance, sensors 204 may be configured to measure physical properties such as the location and/or movement of the asset 200, in which case the sensors may take the form of GNSS sensors, dead-reckoning-based sensors, accelerometers, gyroscopes, pedometers, magnetometers, or the like.

Additionally, various sensors 204 may be configured to measure other operating conditions of the asset 200, examples of which may include temperatures, pressures, speeds, acceleration or deceleration rates, friction, power usages, fuel usages, fluid levels, runtimes, voltages and currents, magnetic fields, electric fields, presence or absence of objects, positions of components, and power generation, among other examples. One of ordinary skill in the art will appreciate that these are but a few example operating conditions that sensors may be configured to measure. Additional or fewer sensors may be used depending on the industrial application or specific asset.

As suggested above, an actuator 205 may be configured similar in some respects to a sensor 204. Specifically, an actuator 205 may be configured to detect a physical property indicative of an operating condition of the asset 200 and provide an indication thereof in a manner similar to the sensor 204.

Moreover, an actuator 205 may be configured to interact with the asset 200, one or more subsystems 202, and/or some component thereof. As such, an actuator 205 may include a motor or the like that is configured to perform a mechanical operation (e.g., move) or otherwise control a component, subsystem, or system. In a particular example, an actuator may be configured to measure a fuel flow and alter the fuel flow (e.g., restrict the fuel flow), or an actuator may be configured to measure a hydraulic pressure and alter the hydraulic pressure (e.g., increase or decrease the hydraulic pressure). Numerous other example interactions of an actuator are also possible and contemplated herein.

Generally, the central processing unit 206 may include one or more processors and/or controllers, which may take the form of a general- or special-purpose processor or controller. In particular, in example implementations, the central processing unit 206 may be or include microprocessors, microcontrollers, application specific integrated circuits, digital signal processors, and the like. In turn, the data storage 208 may be or include one or more non-transitory computer-readable storage media, such as optical, magnetic, organic, or flash memory, among other examples.

The central processing unit 206 may be configured to store, access, and execute computer-readable program instructions stored in the data storage 208 to perform the operations of an asset described herein. For instance, as suggested above, the central processing unit 206 may be configured to receive respective sensor signals from the sensors 204 and/or actuators 205. The central processing unit 206 may be configured to store sensor and/or actuator data in and later access it from the data storage 208.

The central processing unit 206 may also be configured to determine whether received sensor and/or actuator signals trigger any abnormal-condition indicators, such as fault codes. For instance, the central processing unit 206 may be configured to store in the data storage 208 abnormal-condition rules, each of which include a given abnormal-condition indicator representing a particular abnormal condition and respective triggering criteria that trigger the abnormal-condition indicator. That is, each abnormal-condition indicator corresponds with one or more sensor and/or actuator measurement values that must be satisfied before the abnormal-condition indicator is triggered. In practice, the asset 200 may be pre-programmed with the abnormal-condition rules and/or may receive new abnormal-condition rules or updates to existing rules from a computing system, such as the analytics system 108.

In any event, the central processing unit 206 may be configured to determine whether received sensor and/or actuator signals trigger any abnormal-condition indicators. That is, the central processing unit 206 may determine whether received sensor and/or actuator signals satisfy any triggering criteria. When such a determination is affirmative, the central processing unit 206 may generate abnormal-condition data and may also cause the asset's user interface 212 to output an indication of the abnormal condition, such as a visual and/or audible alert.

Additionally, the central processing unit 206 may log the occurrence of the abnormal-condition indicator being triggered in the data storage 208, perhaps with a timestamp.

Figure 3 depicts a conceptual illustration of example abnormal-condition indicators and respective triggering criteria for an asset. In particular, Figure 3 depicts a conceptual illustration of example fault codes. As shown, table 300 includes columns 302, 304, and 306 that correspond to Sensor A, Actuator B, and Sensor C, respectively, and rows 308, 310, and 312 that correspond to Fault Codes 1, 2, and 3, respectively. Entries 314 then specify sensor criteria (e.g., sensor value thresholds) that correspond to the given fault codes.

For example, Fault Code 1 will be triggered when Sensor A detects a rotational measurement greater than 135 revolutions per minute (RPM) and Sensor C detects a temperature measurement greater than 65° Celsius (C), Fault Code 2 will be triggered when Actuator B detects a voltage measurement greater than 1000 Volts (V) and Sensor C detects a temperature measurement less than 55°C, and Fault Code 3 will be triggered when Sensor A detects a rotational measurement greater than 100 RPM, Actuator B detects a voltage measurement greater than 750 V, and Sensor C detects a temperature measurement greater than 60°C. One of ordinary skill in the art will appreciate that Figure 3 is provided for purposes of example and explanation only and that numerous other fault codes and/or triggering criteria are possible and contemplated herein.

Referring back to Figure 2, the central processing unit 206 may be configured to carry out various additional functions for managing and/or controlling operations of the asset 200 as well. For example, the central processing unit 206 may be configured to provide instruction signals to the subsystems 202 and/or the actuators 205 that cause the subsystems 202 and/or the actuators 205 to perform some operation, such as modifying a throttle position. Additionally, the central processing unit 206 may be configured to modify the rate at which it processes data from the sensors 204 and/or the actuators 205, or the central processing unit 206 may be configured to provide instruction signals to the sensors 204 and/or actuators 205 that cause the sensors 204 and/or actuators 205 to, for example, modify a sampling rate. Moreover, the central processing unit 206 may be configured to receive signals from the subsystems 202, the sensors 204, the actuators 205, the network interfaces 210, and/or the user interfaces 212 and based on such signals, cause an operation to occur. Further still, the central processing unit 206 may be configured to receive signals from a computing device, such as a diagnostic device, that cause the central processing unit 206 to execute one or more diagnostic tools in accordance with diagnostic rules stored in the data storage 208. Other functionalities of the central processing unit 206 are discussed below.

The network interface 210 may be configured to provide for communication between the asset 200 and various network components connected to communication network 106. For example, the network interface 210 may be configured to facilitate wireless communications to and from the communication network 106 and may thus take the form of an antenna structure and associated equipment for transmitting and receiving various over-the-air signals. Other examples are possible as well. In practice, the network interface 210 may be configured according to a communication protocol, such as but not limited to any of those described above.

The user interface 212 may be configured to facilitate user interaction with the asset 200 and may also be configured to facilitate causing the asset 200 to perform an operation in response to user interaction. Examples of user interfaces 212 include touch-sensitive interfaces, mechanical interfaces (e.g., levers, buttons, wheels, dials, keyboards, etc.), and other input interfaces (e.g., microphones), among other examples. In some cases, the user interface 212 may include or provide connectivity to output components, such as display screens, speakers, headphone jacks, and the like.

The local analytics device 220 may generally be configured to receive and analyze data related to the asset 200 and based on such analysis, may cause one or more operations to occur at the asset 200. For instance, the local analytics device 220 may receive operating data for the asset 200 (e.g., data generated by the sensors 204 and/or actuators 205) and based on such data, may provide instructions to the central processing unit 206, the sensors 204, and/or the actuators 205 that cause the asset 200 to perform an operation.

To facilitate this operation, the local analytics device 220 may include one or more asset interfaces that are configured to couple the local analytics device 220 to one or more of the asset's on-board systems. For instance, as shown in Figure 2, the local analytics device 220 may have an interface to the asset's central processing unit 206, which may enable the local analytics device 220 to receive operating data from the central processing unit 206 (e.g., operating data that is generated by sensors 204 and/or actuators 205 and sent to the central processing unit 206) and then provide instructions to the central processing unit 206. In this way, the local analytics device 220 may indirectly interface with and receive data from other on-board systems of the asset 200 (e.g., the sensors 204 and/or actuators 205) via the central processing unit 206. Additionally or alternatively, as shown in Figure 2, the local analytics device 220 could have an interface to one or more sensors 204 and/or actuators 205, which may enable the local analytics device 220 to communicate directly with the sensors 204 and/or actuators 205. The local analytics device 220 may interface with the on-board systems of the asset 200 in other manners as well, including the possibility that the interfaces illustrated in Figure 2 are facilitated by one or more intermediary systems that are not shown.

In practice, the local analytics device 220 may enable the asset 200 to locally perform advanced analytics and associated operations, such as executing a predictive model and corresponding workflow, that may otherwise not be able to be performed with the other on-asset components. As such, the local analytics device 220 may help provide additional processing power and/or intelligence to the asset 200.

It should be understood that the local analytics device 220 may also be configured to cause the asset 200 to perform operations that are not related a predictive model. For example, the local analytics device 220 may receive data from a remote source, such as the analytics system 108 or the output system 110, and based on the received data cause the asset 200 to perform one or more operations. One particular example may involve the local analytics device 220 receiving a firmware update for the asset 200 from a remote source and then causing the asset 200 to update its firmware. Another particular example may involve the local analytics device 220 receiving a diagnosis instruction from a remote source and then causing the asset 200 to execute a local diagnostic tool in accordance with the received instruction. Numerous other examples are also possible.

As shown, in addition to the one or more asset interfaces discussed above, the local analytics device 220 may also include a processing unit 222, a data storage 224, and a network interface 226, all of which may be communicatively linked by a system bus, network, or other connection mechanism. The processing unit 222 may include any of the components discussed above with respect to the central processing unit 206. In turn, the data storage 224 may be or include one or more non-transitory computer-readable storage media, which may take any of the forms of computer-readable storage media discussed above.

The processing unit 222 may be configured to store, access, and execute computer-readable program instructions stored in the data storage 224 to perform the operations of a local analytics device described herein. For instance, the processing unit 222 may be configured to receive respective sensor and/or actuator signals generated by the sensors 204 and/or actuators 205 and may execute a predictive model-workflow pair based on such signals. Other functions are described below.

The network interface 226 may be the same or similar to the network interfaces described above. In practice, the network interface 226 may facilitate communication between the local analytics device 220 and the analytics system 108.

In some example implementations, the local analytics device 220 may include and/or communicate with a user interface that may be similar to the user interface 212. In practice, the user interface may be located remote from the local analytics device 220 (and the asset 200). Other examples are also possible.

While Figure 2 shows the local analytics device 220 physically and communicatively coupled to its associated asset (e.g., the asset 200) via one or more asset interfaces, it should also be understood that this might not always be the case. For example, in some implementations, the local analytics device 220 may not be physically coupled to its associated asset and instead may be located remote from the asset 220. In an example of such an implementation, the local analytics device 220 may be wirelessly, communicatively coupled to the asset 200. Other arrangements and configurations are also possible.

One of ordinary skill in the art will appreciate that the asset 200 shown in Figure 2 is but one example of a simplified representation of an asset and that numerous others are also possible. For instance, other assets may include additional components not pictured and/or more or less of the pictured components. Moreover, a given asset may include multiple, individual assets that are operated in concert to perform operations of the given asset. Other examples are also possible.

III. EXAMPLE ANALYTICS SYSTEM

Referring now to Figure 4, a simplified block diagram of an example analytics system 400 is depicted. As suggested above, the analytics system 400 may include one or more computing systems communicatively linked and arranged to carry out various operations described herein. Specifically, as shown, the analytics system 400 may include a data intake system 402, a data science system 404, and one or more databases 406. These system components may be communicatively coupled via one or more wireless and/or wired connections, which may be configured to facilitate secure communications.

The data intake system 402 may generally function to receive and process data and output data to the data science system 404. As such, the data intake system 402 may include one or more network interfaces configured to receive data from various network components of the network configuration 100, such as the assets 102 and 104, the output system 110, and/or the data source 112. Specifically, the data intake system 402 may be configured to receive analog signals, data streams, and/or network packets, among other examples. As such, the network interfaces may include one or more wired network interfaces, such as a port or the like, and/or wireless network interfaces, similar to those described above. In some examples, the data intake system 402 may be or include components configured according to a given dataflow technology, such as a NiFi receiver or the like.

The data intake system 402 may include one or more processing components configured to perform one or more operations. Example operations may include compression and/or decompression, encryption and/or de-encryption, analog-to-digital and/or digital-to-analog conversion, filtration, and amplification, among other operations. Moreover, the data intake system 402 may be configured to parse, sort, organize, and/or route data based on data type and/or characteristics of the data. In some examples, the data intake system 402 may be configured to format, package, and/or route data based on one or more characteristics or operating parameters of the data science system 404.

In general, the data received by the data intake system 402 may take various forms. For example, the payload of the data may include a single sensor or actuator measurement, multiple sensor and/or actuator measurements and/or one or more abnormal-condition data. Other examples are also possible.

Moreover, the received data may include certain characteristics, such as a source identifier and a timestamp (e.g., a date and/or time at which the information was obtained). For instance, a unique identifier (e.g., a computer generated alphabetic, numeric, alphanumeric, or the like identifier) may be assigned to each asset, and perhaps to each sensor and actuator. Such identifiers may be operable to identify the asset, sensor, or actuator from which data originates. In some cases, another characteristic may include the location (e.g., GPS coordinates) at which the information was obtained. Data characteristics may come in the form of signal signatures or metadata, among other examples.

The data science system 404 may generally function to receive (e.g., from the data intake system 402) and analyze data and based on such analysis, cause one or more operations to occur. As such, the data science system 404 may include one or more network interfaces 408, a processing unit 410, and data storage 412, all of which may be communicatively linked by a system bus, network, or other connection mechanism. In some cases, the data science system 404 may be configured to store and/or access one or more application program interfaces (APIs) that facilitate carrying out some of the functionality disclosed herein.

The network interfaces 408 may be the same or similar to any network interface described above. In practice, the network interfaces 408 may facilitate communication (e.g., with some level of security) between the data science system 404 and various other entities, such as the data intake system 402, the databases 406, the assets 102, the output system 110, etc.

The processing unit 410 may include one or more processors, which may take any of the processor forms described above. In turn, the data storage 412 may be or include one or more non-transitory computer-readable storage media, which may take any of the forms of computer-readable storage media discussed above. The processing unit 410 may be configured to store, access, and execute computer-readable program instructions stored in the data storage 412 to perform the operations of an analytics system described herein.

In general, the processing unit 410 may be configured to perform analytics on data received from the data intake system 402. To that end, the processing unit 410 may be configured to execute one or more modules, which may each take the form of one or more sets of program instructions that are stored in the data storage 412. The modules may be configured to facilitate causing an outcome to occur based on the execution of the respective program instructions. An example outcome from a given module may include outputting data into another module, updating the program instructions of the given module and/or of another module, and outputting data to a network interface 408 for transmission to an asset and/or the output system 110, among other examples.

The databases 406 may generally function to receive (e.g., from the data science system 404) and store data. As such, each database 406 may include one or more non-transitory computer-readable storage media, such as any of the examples provided above. In practice, the databases 406 may be separate from or integrated with the data storage 412.

The databases 406 may be configured to store numerous types of data, some of which is discussed below. In practice, some of the data stored in the databases 406 may include a timestamp indicating a date and time at which the data was generated or added to the database. Moreover, data may be stored in a number of manners in the databases 406. For instance, data may be stored in time sequence, in a tabular manner, and/or organized based on data source type (e.g., based on asset, asset type, sensor, sensor type, actuator, or actuator type) or abnormal-condition indicator, among other examples.

IV. EXAMPLE OPERATIONS

The operations of the example network configuration 100 depicted in Figure 1 will now be discussed in further detail below. To help describe some of these operations, flow diagrams may be referenced to describe combinations of operations that may be performed. In some cases, each block may represent a module or portion of program code that includes instructions that are executable by a processor to implement specific logical functions or steps in a process. The program code may be stored on any type of computer-readable medium, such as non-transitory computer-readable media. In other cases, each block may represent circuitry that is wired to perform specific logical functions or steps in a process. Moreover, the blocks shown in the flow diagrams may be rearranged into different orders, combined into fewer blocks, separated into additional blocks, and/or removed based upon the particular embodiment.

The following description may reference examples where a single data source, such as the asset 102, provides data to the analytics system 108 that then performs one or more functions. It should be understood that this is done merely for sake of clarity and explanation and is not meant to be limiting. In practice, the analytics system 108 generally receives data from multiple sources, perhaps simultaneously, and performs operations based on such aggregate received data.

A. COLLECTION OF OPERATING DATA

As mentioned above, the representative asset 102 may take various forms and may be configured to perform a number of operations. In a non-limiting example, the asset 102 may take the form of a locomotive that is operable to transfer cargo across the United States. While in transit, the sensors and/or actuators of the asset 102 may obtain data that reflects one or more operating conditions of the asset 102. The sensors and/or actuators may transmit the data to a processing unit of the asset 102.

The processing unit may be configured to receive the data from the sensors and/or actuators. In practice, the processing unit may receive sensor data from multiple sensors and/or actuator data from multiple actuators simultaneously or sequentially. As discussed above, while receiving this data, the processing unit may also be configured to determine whether the data satisfies triggering criteria that trigger any abnormal-condition indicators, such as fault codes. In the event the processing unit determines that one or more abnormal-condition indicators are triggered, the processing unit may be configured to perform one or more local operations, such as outputting an indication of the triggered indicator via a user interface.

The asset 102 may then transmit operating data to the analytics system 108 via a network interface of the asset 102 and the communication network 106. In operation, the asset 102 may transmit operating data to the analytics system 108 continuously, periodically, and/or in response to triggering events (e.g., abnormal conditions). Specifically, the asset 102 may transmit operating data periodically based on a particular frequency (e.g., daily, hourly, every fifteen minutes, once per minute, once per second, etc.), or the asset 102 may be configured to transmit a continuous, real-time feed of operating data. Additionally or alternatively, the asset 102 may be configured to transmit operating data based on certain triggers, such as when sensor and/or actuator measurements satisfy triggering criteria for any abnormal-condition indicators. The asset 102 may transmit operating data in other manners as well.

In practice, operating data for the asset 102 may include sensor data, actuator data, and/or abnormal-condition data. In some implementations, the asset 102 may be configured to provide the operating data in a single data stream, while in other implementations the asset 102 may be configured to provide the operating data in multiple, distinct data streams. For example, the asset 102 may provide to the analytics system 108 a first data stream of sensor and/or actuator data and a second data stream of abnormal-condition data. Other possibilities also exist.

Sensor and actuator data may take various forms. For example, at times, sensor data (or actuator data) may include measurements obtained by each of the sensors (or actuators) of the asset 102. While at other times, sensor data (or actuator data) may include measurements obtained by a subset of the sensors (or actuators) of the asset 102.

Specifically, the sensor and/or actuator data may include measurements obtained by the sensors and/or actuators associated with a given triggered abnormal-condition indicator. For example, if a triggered fault code is Fault Code 1 from Figure 3, then sensor data may include raw measurements obtained by Sensors A and C. Additionally or alternatively, the data may include measurements obtained by one or more sensors or actuators not directly associated with the triggered fault code. Continuing off the last example, the data may additionally include measurements obtained by Actuator B and/or other sensors or actuators. In some examples, the asset 102 may include particular sensor data in the operating data based on a fault-code rule or instruction provided by the analytics system 108, which may have, for example, determined that there is a correlation between that which Actuator B is measuring and that which caused the Fault Code 1 to be triggered in the first place. Other examples are also possible.

Further still, the data may include one or more sensor and/or actuator measurements from each sensor and/or actuator of interest based on a particular time of interest, which may be selected based on a number of factors. In some examples, the particular time of interest may be based on a sampling rate. In other examples, the particular time of interest may be based on the time at which an abnormal-condition indicator is triggered.

In particular, based on the time at which an abnormal-condition indicator is triggered, the data may include one or more respective sensor and/or actuator measurements from each sensor and/or actuator of interest (e.g., sensors and/or actuators directly and indirectly associated with the triggered indicator). The one or more measurements may be based on a particular number of measurements or particular duration of time around the time of the triggered abnormal-condition indicator.

For example, if a triggered fault code is Fault Code 2 from Figure 3, the sensors and actuators of interest might include Actuator B and Sensor C. The one or more measurements may include the most recent respective measurements obtained by Actuator B and Sensor C prior to the triggering of the fault code (e.g., triggering measurements) or a respective set of measurements before, after, or about the triggering measurements. For example, a set of five measurements may include the five measurements before or after the triggering measurement (e.g., excluding the triggering measurement), the four measurements before or after the triggering measurement and the triggering measurement, or the two measurements before and the two after as well as the triggering measurement, among other possibilities.

Similar to sensor and actuator data, the abnormal-condition data may take various forms. In general, the abnormal-condition data may include or take the form of an indicator that is operable to uniquely identify a particular abnormal condition that occurred at the asset 102 from all other abnormal conditions that may occur at the asset 102. The abnormal-condition indicator may take the form of an alphabetic, numeric, or alphanumeric identifier, among other examples. Moreover, the abnormal-condition indicator may take the form of a string of words that is descriptive of the abnormal condition, such as "Overheated Engine" or "Out of Fuel", among other examples.

The analytics system 108, and in particular, the data intake system of the analytics system 108, may be configured to receive operating data from one or more assets and/or data sources. The data intake system may be configured to perform one or more operations to the received data and then relay the data to the data science system of the analytics system 108. In turn, the data science system may analyze the received data and based on such analysis, perform one or more operations.

B. DEFINING PREDICTIVE MODELS & WORKFLOWS

As one example, the analytics system 108 may be configured to define predictive models and corresponding workflows based on received operating data for one or more assets and/or received external data related to the one or more assets. The analytics system 108 may define model-workflow pairs based on various other data as well.

In general, a model-workflow pair may include a set of program instructions that cause an asset to monitor certain operating conditions and carry out certain operations that help facilitate preventing the occurrence of a particular event suggested by the monitored operating conditions. Specifically, a predictive model may include one or more algorithms whose inputs are sensor and/or actuator data from one or more sensors and/or actuators of an asset and whose outputs are utilized to determine a probability that a particular event may occur at the asset within a particular period of time in the future. In turn, a workflow may include one or more triggers (e.g., model output values) and corresponding operations that the asset carries out based on the triggers.

As suggested above, the analytics system 108 may be configured to define aggregate and/or individualized predictive models and/or workflows. An "aggregate" model/workflow may refer to a model/workflow that is generic for a group of assets and defined without taking into consideration particular characteristics of the assets to which the model/workflow is deployed. On the other hand, an "individualized" model/workflow may refer to a model/workflow that is specifically tailored for a single asset or a subgroup of assets from the group of assets and defined based on particular characteristics of the single asset or subgroup of assets to which the model/workflow is deployed. These different types of models/workflows and the operations performed by the analytics system 108 to define them are discussed in further detail below.

1. Aggregate Models & Workflows

In example implementations, the analytics system 108 may be configured to define an aggregate model-workflow pair based on aggregated data for a plurality of assets. Defining aggregate model-workflow pairs may be performed in a variety of manners.

Figure 5 is a flow diagram 500 depicting one possible example of a definition phase that may be used for defining model-workflow pairs. For purposes of illustration, the example definition phase is described as being carried out by the analytics system 108, but this definition phase may be carried out by other systems as well. One of ordinary skill in the art will appreciate that the flow diagram 500 is provided for sake of clarity and explanation and that numerous other combinations of operations may be utilized to define a model-workflow pair.

As shown in Figure 5, at block 502, the analytics system 108 may begin by defining a set of data that forms the basis for a given predictive model (e.g., the data of interest). The data of interest may derive from a number of sources, such as the assets 102 and 104 and the data source 112, and may be stored in a database of the analytics system 108.

The data of interest may include historical data for a particular set of assets from a group of assets or all of the assets from a group of assets (e.g., the assets of interest). Moreover, the data of interest may include measurements from a particular set of sensors and/or actuators from each of the assets of interest or from all of the sensors and/or actuators from each of the assets of interest. Further still, the data of interest may include data from a particular period of time in the past, such as two week's worth of historical data.

The data of interest may include a variety of types data, which may depend on the given predictive model. In some instances, the data of interest may include at least operating data indicating operating conditions of assets, where the operating data is as discussed above in the Collection of Operating Data section. Additionally, the data of interest may include environment data indicating environments in which assets are typically operated and/or scheduling data indicating planned dates and times during which assets are to carry out certain tasks. Other types of data may also be included in the data of interest.

In practice, the data of interest may be defined in a number of manners. In one example, the data of interest may be user-defined. In particular, a user may operate an output system 110 that receives user inputs indicating a selection of certain data of interest, and the output system 110 may provide to the analytics system 108 data indicating such selections. Based on the received data, the analytics system 108 may then define the data of interest.

In another example, the data of interest may be machine-defined. In particular, the analytics system 108 may perform various operations, such as simulations, to determine the data of interest that generates the most accurate predictive model. Other examples are also possible.

Returning to Figure 5, at block 504, the analytics system 108 may be configured to, based on the data of interest, define an aggregate, predictive model that is related to the operation of assets. In general, an aggregate, predictive model may define a relationship between operating conditions of assets and a likelihood of an event occurring at the assets. Specifically, an aggregate, predictive model may receive as inputs sensor data from sensors of an asset and/or actuator data from actuators of the asset and output a probability that an event will occur at the asset within a certain amount of time into the future.

The event that the predictive model predicts may vary depending on the particular implementation. For example, the event may be a failure and so, the predictive model may be a failure model that predicts whether a failure will occur within a certain period of time in the future (failure models are discussed in detail below in the Health-Score Models & Workflows section). In another example, the event may be an asset completing a task and so, the predictive model may predict the likelihood that an asset will complete a task on time. In other examples, the event may be a fluid or component replacement, and so, the predictive model may predict an amount of time before a particular asset fluid or component needs to be replaced. In yet other examples, the event may be a change in asset productivity, and so, the predictive model may predict the productivity of an asset during a particular period of time in the future. In one other example, the event may be the occurrence of a "leading indicator" event, which may indicate an asset behavior that differs from expected asset behaviors, and so, the predictive model may predict the likelihood of one or more leading indicator events occurring in the future. Other examples of predictive models are also possible.

In any event, the analytics system 108 may define the aggregate, predictive model in a variety of manners. In general, this operation may involve utilizing one or more modeling techniques to generate a model that returns a probability between zero and one, such as a random forest technique, logistic regression technique, or other regression technique, among other modeling techniques. In a particular example implementation, the analytics system 108 may define the aggregate, predictive model in line with the below discussion referencing Figure 7. The analytics system 108 may define the aggregate model in other manners as well.

At block 506, the analytics system 108 may be configured to define an aggregate workflow that corresponds to the defined model from block 504. In general, a workflow may take the form of an action that is carried out based on a particular output of a predictive model. In example implementations, a workflow may include one or more operations that an asset performs based on the output of the defined predictive model. Examples of operations that may be part of a workflow include an asset acquiring data according to a particular data-acquisition scheme, transmitting data to the analytics system 108 according to a particular data-transmission scheme, executing a local diagnostic tool, and/or modifying an operating condition of the asset, among other example workflow operations.

A particular data-acquisition scheme may indicate how an asset acquires data. In particular, a data-acquisition scheme may indicate certain sensors and/or actuators from which the asset obtains data, such as a subset of sensors and/or actuators of the asset's plurality of sensors and actuators (e.g., sensors/actuators of interest). Further, a data-acquisition scheme may indicate an amount of data that the asset obtains from the sensors/actuators of interest and/or a sampling frequency at which the asset acquires such data. Data-acquisition schemes may include various other attributes as well. In a particular example implementation, a particular data-acquisition scheme may correspond to a predictive model for asset health and may be adjusted to acquire more data and/or particular data (e.g., from particular sensors) based on a decreasing asset health. Or a particular data-acquisition scheme may correspond to a leading-indicators predictive model and may be adjusted to a modify data acquired by asset sensors and/or actuators based on an increased likelihood of an occurrence of a leading indicator event that may signal that a subsystem failure might occur.

A particular data-transmission scheme may indicate how an asset transmits data to the analytics system 108. Specifically, a data-transmission scheme may indicate a type of data (and may also indicate a format and/or structure of the data) that the asset should transmit, such as data from certain sensors or actuators, a number of data samples that the asset should transmit, a transmission frequency, and/or a priority-scheme for the data that the asset should include in its data transmission. In some cases, a particular data-acquisition scheme may include a data-transmission scheme or a data-acquisition scheme may be paired with a data-transmission scheme. In some example implementations, a particular data-transmission scheme may correspond to a predictive model for asset health and may be adjusted to transmit data less frequently based on an asset health that is above a threshold value. Other examples are also possible.

As suggested above, a local diagnostic tool may be a set of procedures or the like that are stored locally at an asset. The local diagnostic tool may generally facilitate diagnosing a cause of a fault or failure at an asset. In some cases, when executed, a local diagnostic tool may pass test inputs into a subsystem of an asset or a portion thereof to obtain test results, which may facilitate diagnosing the cause of a fault or failure. These local diagnostic tools are typically dormant on an asset and will not be executed unless the asset receives particular diagnostic instructions. Other local diagnostic tools are also possible. In one example implementation, a particular local diagnostic tool may correspond to a predictive model for health of a subsystem of an asset and may be executed based on a subsystem health that is at or below a threshold value.

Lastly, a workflow may involve modifying an operating condition of an asset. For instance, one or more actuators of an asset may be controlled to facilitate modifying an operating condition of the asset. Various operating conditions may be modified, such as a speed, temperature, pressure, fluid level, current draw, and power distribution, among other examples. In a particular example implementation, an operating-condition modification workflow may correspond to a predictive model for predicting whether an asset will complete a task on time and may cause the asset to increase its speed of travel based on a predicted completion percentage that is below a threshold value.

In any event, the aggregate workflow may be defined in a variety of manners. In one example, the aggregate workflow may be user defined. Specifically, a user may operate a computing device that receives user inputs indicating selection of certain workflow operations, and the computing device may provide to the analytics system 108 data indicating such selections. Based on this data, the analytics system 108 may then define the aggregate workflow.

In another example, the aggregate workflow may be machine-defined. In particular, the analytics system 108 may perform various operations, such as simulations, to determine a workflow that may facilitate determining a cause of the probability output by the predictive model and/or preventing an occurrence of an event predicted by the model. Other examples of defining the aggregate workflow are also possible.

In defining the workflow corresponding to the predictive model, the analytics system 108 may define the triggers of the workflow. In example implementations, a workflow trigger may be a value of the probability output by the predictive model or a range of values output by the predictive model. In some cases, a workflow may have multiple triggers, each of which may cause a different operation or operations to occur.

To illustrate, Figure 6A is a conceptual illustration of an aggregate model-workflow pair 600. As shown, the aggregate model-workflow pair illustration 600 includes a column for model inputs 602, model calculations 604, model output ranges 606, and corresponding workflow operations 608. In this example, the predictive model has a single input, data from Sensor A, and has two calculations, Calculations I and II. The output of this predictive model affects the workflow operation that is performed. If the output probability is less than or equal to 80%, then workflow Operation 1 is performed. Otherwise, the workflow Operation 2 is performed. Other example model-workflow pairs are possible and contemplated herein.

2. Individualized Models & Workflows

In another aspect, the analytics system 108 may be configured to define individualized predictive models and/or workflows for assets, which may involve utilizing the aggregate model-workflow pair as a baseline. The individualization may be based on certain characteristics of assets. In this way, the analytics system 108 may provide a given asset a more accurate and robust model-workflow pair compared to the aggregate model-workflow pair.

In particular, returning to Figure 5, at block 508, the analytics system 108 may be configured to decide whether to individualize the aggregate model defined at block 504 for a given asset, such as the asset 102. The analytics system 108 may carry out this decision in a number of manners.

In some cases, the analytics system 108 may be configured to define individualized predictive models by default. In other cases, the analytics system 108 may be configured to decide whether to define an individualized predictive model based on certain characteristics of the asset 102. For example, in some cases, only assets of certain types or classes, or operated in certain environments, or that have certain health scores may receive an individualized predictive model. In yet other cases, a user may define whether an individualized model is defined for the asset 102. Other examples are also possible.

In any event, if the analytics system 108 decides to define an individualized predictive model for the asset 102, the analytics system 108 may do so at block 510. Otherwise, the analytics system 108 may proceed to block 512.

At block 510, the analytics system 108 may be configured to define an individualized predictive model in a number of manners. In example implementations, the analytics system 108 may define an individualized predictive model based at least in part on one or more characteristics of the asset 102.

Before defining the individualized predictive model for the asset 102, the analytics system 108 may have determined one or more asset characteristics of interest that form the basis of individualized models. In practice, different predictive models may have different corresponding characteristics of interest.

In general, the characteristics of interest may be characteristics that are related to the aggregate model-workflow pair. For instance, the characteristics of interest may be characteristics that the analytics system 108 has determined influence the accuracy of the aggregate model-workflow pair. Examples of such characteristics may include asset age, asset usage, asset capacity, asset load, asset health (perhaps indicated by an asset health metric, discussed below), asset class (e.g., brand and/or model), and environment in which an asset is operated, among other characteristics.

The analytics system 108 may have determined the characteristics of interest in a number of manners. In one example, the analytics system 108 may have done so by performing one or more modeling simulations that facilitate identifying the characteristics of interest. In another example, the characteristics of interest may have been predefined and stored in the data storage of the analytics system 108. In yet another example, characteristics of interest may have been define by a user and provided to the analytics system 108 via the output system 110. Other examples are also possible.

In any event, after determining the characteristics of interest, the analytics system 108 may determine characteristics of the asset 102 that correspond to the determined characteristics of interest. That is, the analytics system 108 may determine a type, value, existence or lack thereof, etc. of the asset 102's characteristics that correspond to the characteristics of interest. The analytics system 108 may perform this operation in a number of manners.

For examples, the analytics system 108 may be configured to perform this operation based on data originating from the asset 102 and/or the data source 112. In particular, the analytics system108 may utilize operating data for the asset 102 and/or external data from the data source 112 to determine one or more characteristics of the asset 102. Other examples are also possible.

Based on the determined one or more characteristics of the asset 102, the analytics system 108 may define an individualized, predictive model by modifying the aggregate model. The aggregate model may be modified in a number of manners. For example, the aggregate model may be modified by changing (e.g., adding, removing, re-ordering, etc.) one or more model inputs, changing one or more sensor and/or actuator measurement ranges that correspond to asset-operating limits (e.g., changing operating limits that correspond to "leading indicator" events), changing one or more model calculations, weighting (or changing a weight of) a variable or output of a calculation, utilizing a modeling technique that differs from that which was utilized to define the aggregate model, and/or utilizing a response variable that differs from that which was utilized to define the aggregate model, among other examples.

To illustrate, Figure 6B is a conceptual illustration of an individualized model-workflow pair 610. Specifically, the individualized model-workflow pair illustration 610 is a modified version of the aggregate model-workflow pair from Figure 6A. As shown, the individualized model-workflow pair illustration 610 includes a modified column for model inputs 612 and model calculations 614 and includes the original columns for model output ranges 606 and workflow operations 608 from Figure 6A. In this example, the individualized model has two inputs, data from Sensor A and Actuator B, and has two calculations, Calculations II and III. The output ranges and corresponding workflow operations are the same as those of Figure 6A. The analytics system 108 may have defined the individualized model in this way based on determining that the asset 102 is, for example, relatively old and has relatively poor health, among other reasons.

In practice, individualizing the aggregate model may depend on the one or more characteristics of the given asset. In particular, certain characteristics may affect the modification of the aggregate model differently than other characteristics. Further, the type, value, existence, or the like of a characteristic may affect the modification as well. For example, the asset age may affect a first part of the aggregate model, while an asset class may affect a second, different part of the aggregate model. And an asset age within a first range of ages may affect the first part of the aggregate model in a first manner, while an asset age within a second range of ages, different from the first range, may affect the first part of the aggregate model in a second, different manner. Other examples are also possible.

In some implementations, individualizing the aggregate model may depend on considerations in addition to or alternatively to asset characteristics. For instance, the aggregate model may be individualized based on sensor and/or actuator readings of an asset when the asset is known to be in a relatively good operating state (e.g., as defined by a mechanic or the like). More particularly, in an example of a leading-indicator predictive model, the analytics system 108 may be configured to receive an indication that the asset is in a good operating state (e.g., from a computing device operated by a mechanic) along with operating data from the asset. Based at least on the operating data, the analytics system 108 may then individualize the leading-indicator predictive model for the asset by modifying respective operating limits corresponding to “leading indicator” events. Other examples are also possible.

Returning to Figure 5, at block 512, the analytics system 108 may also be configured to decide whether to individualize a workflow for the asset 102. The analytics system 108 may carry out this decision in a number of manners. In some implementations, the analytics system 108 may perform this operation in line with block 508. In other implementations, the analytics system 108 may decide whether to define an individualized workflow based on the individualized predictive model. In yet another implementation, the analytics system 108 may decide to define an individualized workflow if an individualized predictive model was defined. Other examples are also possible.

In any event, if the analytics system 108 decides to define an individualized workflow for the asset 102, the analytics system 108 may do so at block 514. Otherwise, the analytics system 108 may end the definition phase.

At block 514, the analytics system 108 may be configured to define an individualized workflow in a number of manners. In example implementations, the analytics system 108 may define an individualized workflow based at least in part on one or more characteristics of the asset 102.

Before defining the individualized workflow for the asset 102, similar to defining the individualized predictive model, the analytics system 108 may have determined one or more asset characteristics of interest that form the basis of an individualized workflow, which may have been determined in line with the discussion of block 510. In general, these characteristics of interest may be characteristics that affect the efficacy of the aggregate workflow. Such characteristics may include any of the example characteristics discussed above. Other characteristics are possible as well.

Similar again to block 510, the analytics system 108 may determine characteristics of the asset 102 that correspond to the determined characteristics of interest for an individualized workflow. In example implementations, the analytics system 108 may determine characteristics of the asset 102 in a manner similar to the characteristic determination discussed with reference to block 510 and in fact, may utilize some or all of that determination.

In any event, based on the determined one or more characteristics of the asset 102, the analytics system 108 may individualize a workflow for the asset 102 by modifying the aggregate workflow. The aggregate workflow may be modified in a number of manners. For example, the aggregate workflow may be modified by changing (e.g., adding, removing, re-ordering, replacing, etc.) one or more workflow operations (e.g., changing from a first data-acquisition scheme to a second scheme or changing from a particular data-acquisition scheme to a particular local diagnostic tool) and/or changing (e.g., increasing, decreasing, adding to, removing from, etc.) the corresponding model output value or range of values that triggers particular workflow operations, among other examples. In practice, modification to the aggregate workflow may depend on the one or more characteristics of the asset 102 in a manner similar to the modification to the aggregate model.

To illustrate, Figure 6C is a conceptual illustration of an individualized model-workflow pair 620. Specifically, the individualized model-workflow pair illustration 620 is a modified version of the aggregate model-workflow pair from Figure 6A. As shown, the individualized model-workflow pair illustration 620 includes the original columns for model inputs 602, model calculations 604, and model output ranges 606 from Figure 6A, but includes a modified column for workflow operations 628. In this example, the individualized model-workflow pair is similar to the aggregate model-workflow pair from Figure 6A, except that when the output of the aggregate model is greater than 80% workflow Operation 3 is triggered instead of Operation 1. The analytics system 108 may have defined this individual workflow based on determining that the asset 102, for example, operates in an environment that historically increases the occurrence of asset failures, among other reasons.

After defining the individualized workflow, the analytics system 108 may end the definition phase. At that point, the analytics system 108 may then have an individualized model-workflow pair for the asset 102.

In some example implementations, the analytics system 108 may be configured to define an individualized predictive model and/or corresponding workflow for a given asset without first defining an aggregate predictive model and/or corresponding workflow. Other examples are also possible.

While the above discussed the analytics system 108 individualizing predictive models and/or workflows, other devices and/or systems may perform the individualization. For example, the local analytics device of the asset 102 may individualize a predictive model and/or workflow or may work with the analytics system 108 to perform such operations. The local analytics device performing such operations is discussed in further detail below.

3. Health-Score Models & Workflows

In a particular implementation, as mentioned above, the analytics system 108 may be configured to define predictive models and corresponding workflows associated with the health of assets. In example implementations, one or more predictive models for monitoring the health of an asset may be utilized to output a health metric (e.g., "health score") for an asset, which is a single, aggregated metric that indicates whether a failure will occur at a given asset within a given timeframe into the future (e.g., the next two weeks). In particular, a health metric may indicate a likelihood that no failures from a group of failures will occur at an asset within a given timeframe into the future, or a health metric may indicate a likelihood that at least one failure from a group of failures will occur at an asset within a given timeframe into the future.

In practice, the predictive models utilized to output a health metric and the corresponding workflows may be defined as aggregate or individualized models and/or workflows, in line with the above discussion.

Moreover, depending on the desired granularity of the health metric, the analytics system 108 may be configured to define different predictive models that output different levels of health metrics and to define different corresponding workflows. For example, the analytics system 108 may define a predictive model that outputs a health metric for the asset as a whole (i.e., an asset-level health metric). As another example, the analytics system 108 may define a respective predictive model that outputs a respective health metric for one or more subsystems of the asset (i.e., subsystem-level health metrics). In some cases, the outputs of each subsystem-level predictive model may be combined to generate an asset-level health metric. Other examples are also possible.

In general, defining a predictive model that outputs a health metric may be performed in a variety of manners. Figure 7 is a flow diagram 700 depicting one possible example of a modeling phase that may be used for defining a model that outputs a health metric. For purposes of illustration, the example modeling phase is described as being carried out by the analytics system 108, but this modeling phase may be carried out by other systems as well. One of ordinary skill in the art will appreciate that the flow diagram 700 is provided for sake of clarity and explanation and that numerous other combinations of operations may be utilized to determine a health metric.

As shown in Figure 7, at block 702, the analytics system 108 may begin by defining a set of the one or more failures that form the basis for the health metric (i.e., the failures of interest). In practice, the one or more failures may be those failures that could render an asset (or a subsystem thereof) inoperable if they were to occur. Based on the defined set of failures, the analytics system 108 may take steps to define a model for predicting a likelihood of any of the failures occurring within a given timeframe in the future (e.g., the next two weeks).

In particular, at block 704, the analytics system 108 may analyze historical operating data for a group of one or more assets to identify past occurrences of a given failure from the set of failures. At block 706, the analytics system 108 may identify a respective set of operating data that is associated with each identified past occurrence of the given failure (e.g., sensor and/or actuator data from a given timeframe prior to the occurrence of the given failure). At block 708, the analytics system 108 may analyze the identified sets of operating data associated with past occurrences of the given failure to define a relationship (e.g., a failure model) between (1) the values for a given set of operating metrics and (2) the likelihood of the given failure occurring within a given timeframe in the future (e.g., the next two weeks). Lastly, at block 710, the defined relationship for each failure in the defined set (e.g., the individual failure models) may then be combined into a model for predicting the overall likelihood of a failure occurring.

As the analytics system 108 continues to receive updated operating data for the group of one or more assets, the analytics system 108 may also continue to refine the predictive model for the defined set of one or more failures by repeating steps 704-710 on the updated operating data.

The functions of the example modeling phase illustrated in Figure 7 will now be described in further detail. Starting with block 702, as noted above, the analytics system 108 may begin by defining a set of the one or more failures that form the basis for the health metric. The analytics system 108 may perform this function in various manners.

In one example, the set of the one or more failures may be based on one or more user inputs. Specifically, the analytics system 108 may receive from a computing system operated by a user, such as the output system 110, input data indicating a user selection of the one or more failures. As such, the set of one or more failures may be user-defined.

In other examples, the set of the one or more failures may be based on a determination made by the analytics system 108 (e.g., machine-defined). In particular, the analytics system 108 may be configured to define the set of one or more failures, which may occur in a number of manners.

For instance, the analytics system 108 may be configured to define the set of failures based on one or more characteristics of the asset 102. That is, certain failures may correspond to certain characteristics, such as asset type, class, etc., of an asset. For example, each type and/or class of asset may have respective failures of interest.

In another instance, the analytics system 108 may be configured to define the set of failures based on historical data stored in the databases of the analytics system 108 and/or external data provided by the data source 112. For example, the analytics system 108 may utilize such data to determine which failures result in the longest repair-time and/or which failures are historically followed by additional failures, among other examples.

In yet other examples, the set of one or more failures may be defined based on a combination of user inputs and determinations made by the analytics system 108. Other examples are also possible.

At block 704, for each of the failures from the set of failures, the analytics system 108 may analyze historical operating data for a group of one or more assets (e.g., abnormal-behavior data) to identify past occurrences of a given failure. The group of the one or more assets may include a single asset, such as asset 102, or multiple assets of a same or similar type, such as fleet of assets that includes the assets 102 and 104. The analytics system 108 may analyze a particular amount of historical operating data, such as a certain amount of time's worth of data (e.g., a month's worth) or a certain number of data-points (e.g., the most recent thousand data-points), among other examples.

In practice, identifying past occurrences of the given failure may involve the analytics system 108 identifying the type of operating data, such as abnormal-condition data, that indicates the given failure. In general, a given failure may be associated with one or multiple abnormal-condition indicators, such as fault codes. That is, when the given failure occurs, one or multiple abnormal-condition indicators may be triggered. As such, abnormal-condition indicators may be reflective of an underlying symptom of a given failure.

After identifying the type of operating data that indicates the given failure, the analytics system 108 may identify the past occurrences of the given failure in a number of manners. For instance, the analytics system 108 may locate, from historical operating data stored in the databases of the analytics system 108, abnormal-condition data corresponding to the abnormal-condition indicators associated with the given failure. Each located abnormal-condition data would indicate an occurrence of the given failure. Based on this located abnormal-condition data, the analytics system 108 may identify a time at which a past failure occurred.

At block 706, the analytics system 108 may identify a respective set of operating data that is associated with each identified past occurrence of the given failure. In particular, the analytics system 108 may identify a set of sensor and/or actuator data from a certain timeframe around the time of the given occurrence of the given failure. For example, the set of data may be from a particular timeframe (e.g., two weeks) before, after, or around the given occurrence of the failure. In other cases, the set of data may be identified from a certain number of data-points before, after, or around the given occurrence of the failure.

In example implementations, the set of operating data may include sensor and/or actuator data from some or all of the sensors and actuators of the asset 102. For example, the set of operating data may include data from sensors and/or actuators associated with an abnormal-condition indicator corresponding to the given failure.

To illustrate, Figure 8 depicts a conceptual illustration of historical operating data that the analytics system 108 may analyze to facilitate defining a model. Plot 800 may correspond to a segment of historical data that originated from some (e.g., Sensor A and Actuator B) or all of the sensors and actuators of the asset 102. As shown, the plot 800 includes time on the x-axis 802, measurement values on the y-axis 804, and sensor data 806 corresponding to Sensor A and actuator data 808 corresponding to Actuator B, each of which includes various data-points representing measurements at particular points in time, T_i. Moreover, the plot 800 includes an indication of an occurrence of a failure 810 that occurred at a past time, T_f (e.g., "time of failure"), and an indication of an amount of time 812 before the occurrence of the failure, ΔT, from which sets of operating data are identified. As such, $T_f - \Delta T$ defines a timeframe 814 of data-points of interest.

Returning to Figure 7, after the analytics system 108 identifies the set of operating data for the given occurrence of the given failure (e.g., the occurrence at T_f), the analytics system 108 may determine whether there are any remaining occurrences for which a set of operating data should be identified. In the event that there is a remaining occurrence, block 706 would be repeated for each remaining occurrence.

Thereafter, at block 708, the analytics system 108 may analyze the identified sets of operating data associated with the past occurrences of the given failure to define a relationship (e.g., a failure model) between (1) a given set of operating metrics (e.g., a given set of sensor and/or actuator measurements) and (2) the likelihood of the given failure occurring within a given timeframe in the future (e.g., the next two weeks). That is, a given failure model may take as inputs sensor and/or actuator measurements from one or more sensors and/or actuators and output a probability that the given failure will occur within the given timeframe in the future.

In general, a failure model may define a relationship between operating conditions of the asset 102 and the likelihood of a failure occurring. In some implementations, in addition to raw data signals from sensors and/or actuators of the asset 102, a failure model may receive a number of other data inputs, also known as features, which are derived from the sensor and/or actuator signals. Such features may include an average or range of values that were historically measured when a failure occurred, an average or range of value gradients (e.g., a rate of change in measurements) that were historically measured prior to an occurrence of a failure, a duration of time between failures (e.g., an amount of time or number of data-points between a first occurrence of a failure and a second occurrence of a failure), and/or one or more failure patterns indicating sensor and/or actuator measurement trends around the occurrence of a failure. One of ordinary skill in the art will appreciate that these are but a few example features that can be derived from sensor and/or actuator signals and that numerous other features are possible.

In practice, a failure model may be defined in a number of manners. In example implementations, the analytics system 108 may define a failure model by utilizing one or more modeling techniques that return a probability between zero and one, which may take the form of any modeling techniques described above.

In a particular example, defining a failure model may involve the analytics system 108 generating a response variable based on the historical operating data identified at block 706. Specifically, the analytics system 108 may determine an associated response variable for each set of sensor and/or actuator measurements received at a particular point in time. As such, the response variable may take the form of a data set associated with the failure model.

The response variable may indicate whether the given set of measurements is within any of the timeframes determined at block 706. That is, a response variable may reflect whether a given set of data is from a time of interest about the occurrence of a failure. The response variable may be a binary-valued response variable such that, if the given set of measurements is within any of determined timeframes, the associated response variable is assigned a value of one, and otherwise, the associated response variable is assigned a value of zero.

Returning to Figure 8, a conceptual illustration of a response variable vector, Y_{res}, is shown on the plot 800. As shown, response variables associated with sets of measurements that are within the timeframe 814 have a value of one (e.g., Y_{res} at times T_{i+3} -T_{i+8}), while response variables associated with sets of measurements outside the timeframe 814 have a value of zero (e.g., Y_{res} at times T_i -T_{i+2} and T_{i+9} -T_{i+10}). Other response variables are also possible.

Continuing in the particular example of defining a failure model based on a response variable, the analytics system 108 may train the failure model with the historical operating data identified at block 706 and the generated response variable. Based on this training process, the analytics system 108 may then define the failure model that receives as inputs various sensor and/or actuator data and outputs a probability between zero and one that a failure will occur within a period of time equivalent to the timeframe used to generate the response variable.

In some cases, training with the historical operating data identified at block 706 and the generated response variable may result in variable importance statistics for each sensor and/or actuator. A given variable importance statistic may indicate the sensor's or actuator's relative effect on the probability that a given failure will occur within the period of time into the future.

Additionally or alternatively, the analytics system 108 may be configured to define a failure model based on one or more survival analysis techniques, such as a Cox proportional hazard technique. The analytics system 108 may utilize a survival analysis technique in a manner similar in some respects to the above-discussed modeling technique, but the analytics system 108 may determine a survival time-response variable that indicates an amount of time from the last failure to a next expected event. A next expected event may be either reception of senor and/or actuator measurements or an occurrence of a failure, whichever occurs first. This response variable may include a pair of values that are associated with each of the particular points in time at which measurements are received. The response variable may then be utilized to determine a probability that a failure will occur within the given timeframe in the future.

In some example implementations, the failure model may be defined based in part on external data, such as weather data, and "hotbox" data, among other data. For instance, based on such data, the failure model may increase or decrease an output failure probability.

In practice, external data may be observed at points in time that do not coincide with times at which asset sensors and/or actuators obtain measurements. For example, the times at which "hotbox" data is collected (e.g., times at which a locomotive passes along a section of railroad track that is outfitted with hot box sensors) may be in disagreement with sensor and/or actuator measurement times. In such cases, the analytics system 108 may be configured to perform one or more operations to determine external data observations that would have been observed at times that correspond to the sensor measurement times.

Specifically, the analytics system 108 may utilize the times of the external data observations and times of the measurements to interpolate the external data observations to produce external data values for times corresponding to the measurement times. Interpolation of the external data may allow external data observations or features derived therefrom to be included as inputs into the failure model. In practice, various techniques may be used to interpolate the external data with the sensor and/or actuator data, such as nearest-neighbor interpolation, linear interpolation, polynomial interpolation, and spline interpolation, among other examples.

Returning to Figure 7, after the analytics system 108 determines a failure model for a given failure from the set of failures defined at block 702, the analytics system 108 may determine whether there are any remaining failures for which a failure model should be determined. In the event that there remains a failure for which a failure model should be determined, the analytics system 108 may repeat the loop of blocks 704-708. In some implementations, the analytics system 108 may determine a single failure model that encompasses all of the failures defined at block 702. In other implementations, the analytics system 108 may determine a failure model for each subsystem of the asset 102, which may then be utilized to determine an asset-level failure model. Other examples are also possible.

Lastly, at block 710, the defined relationship for each failure in the defined set (e.g., the individual failure models) may then be combined into the model (e.g., the health-metric model) for predicting the overall likelihood of a failure occurring within the given timeframe in the future (e.g., the next two weeks). That is, the model receives as inputs sensor and/or actuator measurements from one or more sensors and/or actuators and outputs a single probability that at least one failure from the set of failures will occur within the given timeframe in the future.

The analytics system 108 may define the health-metric model in a number of manners, which may depend on the desired granularity of the health metric. That is, in instances where there are multiple failure models, the outcomes of the failure models may be utilized in a number of manners to obtain the output of the health-metric model. For example, the analytics system 108 may determine a maximum, median, or average from the multiple failure models and utilize that determined value as the output of the health-metric model.

In other examples, determining the health-metric model may involve the analytics system 108 attributing a weight to individual probabilities output by the individual failure models. For instance, each failure from the set of failures may be considered equally undesirable, and so each probability may likewise be weighted the same in determining the health-metric model. In other instances, some failures may be considered more undesirable than others (e.g., more catastrophic or require longer repair time, etc.), and so those corresponding probabilities may be weighted more than others.

In yet other examples, determining the health-metric model may involve the analytics system 108 utilizing one or more modeling techniques, such as a regression technique. An aggregate response variable may take the form of the logical disjunction (logical OR) of the response variables (e.g., Y_{res} in Figure 8) from each of the individual failure models. For example, aggregate response variables associated with any set of measurements that occur within any timeframe determined at block 706 (e.g., the timeframe 814 of Figure 8) may have a value of one, while aggregate response variables associated with sets of measurements that occur outside any of the timeframes may have a value of zero. Other manners of defining the health-metric model are also possible.

In some implementations, block 710 may be unnecessary. For example, as discussed above, the analytics system 108 may determine a single failure model, in which case the health-metric model may be the single failure model.

In practice, the analytics system 108 may be configured to update the individual failure models and/or the overall health-metric model. The analytics system 108 may update a model daily, weekly, monthly, etc. and may do so based on a new portion of historical operating data from the asset 102 or from other assets (e.g., from other assets in the same fleet as the asset 102). Other examples are also possible.

C. DEPLOYING MODELS & WORKFLOWS

After the analytics system 108 defines a model-workflow pair, the analytics system 108 may deploy the defined model-workflow pair to one or more assets. Specifically, the analytics system 108 may transmit the defined predictive model and/or corresponding workflow to at least one asset, such as the asset 102. The analytics system 108 may transmit model-workflow pairs periodically or based on triggering events, such as any modifications or updates to a given model-workflow pair.

In some cases, the analytics system 108 may transmit only one of an individualized model or an individualized workflow. For example, in scenarios where the analytics system 108 defined only an individualized model or workflow, the analytics system 108 may transmit an aggregate version of the workflow or model along with the individualized model or workflow, or the analytics system 108 may not need to transmit an aggregate version if the asset 102 already has the aggregate version stored in data storage. In sum, the analytics system 108 may transmit (1) an individualized model and/or individualized workflow, (2) an individualized model and the aggregate workflow, (3) the aggregate model and an individualized workflow, or (4) the aggregate model and the aggregate workflow.

In practice, the analytics system 108 may have carried out some or all of the operations of blocks 702-710 of Figure 7 for multiple assets to define model-workflow pairs for each asset. For example, the analytics system 108 may have additionally defined a model-workflow pair for the asset 104. The analytics system 108 may be configured to transmit respective model-workflow pairs to the assets 102 and 104 simultaneously or sequentially.

D. LOCAL EXECUTION BY ASSET

A given asset, such as the asset 102, may be configured to receive a model-workflow pair or a portion thereof and operate in accordance with the received model-workflow pair. That is, the asset 102 may store in data storage the model-workflow pair and input into the predictive model data obtained by sensors and/or actuators of the asset 102 and at times, execute the corresponding workflow based on the output of the predictive model.

In practice, various components of the asset 102 may execute the predictive model and/or corresponding workflow. For example, as discussed above, each asset may include a local analytics device configured to store and run model-workflow pairs provided by the analytics system 108. When the local analytics device receives particular sensor and/or actuator data, it may input the received data into the predictive model and depending on the output of the model, may execute one or more operations of the corresponding workflow.

In another example, a central processing unit of the asset 102 that is separate from the local analytics device may execute the predictive model and/or corresponding workflow. In yet other examples, the local analytics device and central processing unit of the asset 102 may collaboratively execute the model-workflow pair. For instance, the local analytics device may execute the predictive model and the central processing unit may execute the workflow or *vice versa*.

In example implementations, before the model-workflow pair is locally executed (or perhaps when the model-workflow is first locally executed), the local analytics device may individualize the predictive model and/or corresponding workflow for the asset 102. This may occur whether the model-workflow pair takes the form of an aggregate model-workflow pair or an individualized model-workflow pair.

As suggested above, the analytics system 108 may define a model-workflow pair based on certain predictions, assumptions, and/or generalizations about a group of assets or a particular asset. For instance, in defining a model-workflow pair, the analytics system 108 may predict, assume, and/or generalize regarding characteristics of assets and/or operating conditions of assets, among other considerations.

In any event, the local analytics device individualizing a predictive model and/or corresponding workflow may involve the local analytics device confirming or refuting one or more of the predictions, assumptions, and/or generalizations made by the analytics system 108 when the model-workflow pair was defined. The local analytics device may thereafter modify (or further modify, in the case of an already-individualized model and/or workflow) the predictive model and/or workflow in accordance with its evaluation of the predictions, assumptions, and/or generalizations. In this way, the local analytics device may help define a more realistic and/or accurate model-workflow pair, which may result in more efficacious asset monitoring.

In practice, the local analytics device may individualize a predictive model and/or workflow based on a number of considerations. For example, the local analytics device may do so based on operating data generated by one or more sensors and/or actuators of the asset 102. Specifically, the local analytics device may individualize by (1) obtaining operating data generated by a particular group of one or more sensors and/or actuators (e.g., by obtaining such data indirectly via the asset's central processing unit or perhaps directly from certain of the sensor(s) and/or actuator(s) themselves), (2) evaluating one or more predictions, assumptions, and/or generalizations associated with the model-workflow pair based on the obtained operating data, and (3) if the evaluation indicates that any prediction, assumption, and/or generalization was incorrect, then modify the model and/or workflow accordingly. These operations may be performed in a variety of manners.

In one example, the local analytics device obtaining operating data generated by a particular group of sensors and/or actuators (e.g., via the asset's central processing unit) may be based on instructions included as part of or along with the model-workflow pair. In particular, the instructions may identify one or more tests for the local analytics device to execute that evaluate some or all predictions, assumptions, and/or generalizations that were involved in defining the model-workflow pair. Each test may identify one or more sensors and/or actuators of interest for which the local analytics device is to obtain operating data, an amount of operating data to obtain, and/or other test considerations. Therefore, the local analytics device obtaining operating data generated by the particular group of sensors and/or actuators may involve the local analytics device obtaining such operating data in accordance with test instructions or the like. Other examples of the local analytics device obtaining operating data for use in individualizing a model-workflow pair are also possible.

As noted above, after obtaining the operating data, the local analytics device may utilize the data to evaluate some or all predictions, assumptions, and/or generalizations that were involved in defining the model-workflow pair. This operation may be performed in a variety of manners. In one example, the local analytics device may compare the obtained operating data to one or more thresholds (e.g., threshold values and/or threshold ranges of values). Generally, a given threshold value or range may correspond to one or more predictions, assumptions, and/or generalizations used to define the model-workflow pair. Specifically, each sensor or actuator (or a combination of sensors and/or actuators) identified in the test instructions may have a corresponding threshold value or range. The local analytics device may then determine whether the operating data generated by a given sensor or actuator is above or below the corresponding threshold value or range. Other examples of the local analytics device evaluating predictions, assumptions, and/or generalizations are also possible.

Thereafter, the local analytics device may modify (or not) the predictive model and/or workflow based on the evaluation. That is, if the evaluation indicates that any predictions, assumptions, and/or generalizations were incorrect, then the local analytics device may modify the predictive model and/or workflow accordingly. Otherwise, the local analytics device may execute the model-workflow pair without modifications.

In practice, the local analytics device may modify a predictive model and/or workflow in a number of manners. For example, the local analytics device may modify one or more parameters of the predictive model and/or corresponding workflow and/or trigger points of the predictive model and/or workflow (e.g., by modifying a value or range of values), among other examples.

As one non-limiting example, the analytics system 108 may have defined a model-workflow pair for the asset 102 assuming that the asset 102's engine operating temperature does not exceed a particular temperature. As a result, part of the predictive model for the asset 102 may involve determining a first calculation and then a second calculation only if the first calculation exceeds a threshold value, which was determined based on the assumed engine operating temperature. When individualizing the model-workflow pair, the local analytics device may obtain data generated by one or more sensors and/or actuators that measure operating data of the asset 102's engine. The local analytics device may then use this data to determine whether the assumption regarding the engine operating temperate is true in practice (e.g., whether the engine operating temperate exceeds the threshold value). If the data indicates that the engine operating temperate has a value that exceeds, or is a threshold amount above, the assumed particular temperature, then the local analytics device may, for example, modify the threshold value that triggers determining the second calculation. Other examples of the local analytics device individualizing a predictive model and/or workflow are also possible.

The local analytics device may individualize a model-workflow pair based on additional or alternative considerations. For example, the local analytics device may do so based on one or more asset characteristics, such as any of those discussed above, which may be determined by the local analytics device or provided to the local analytics device. Other examples are also possible.

In example implementations, after the local analytics device individualizes a predictive model and/or workflow, the local analytics device may provide to the analytics system 108 an indication that the predictive model and/or workflow has been individualized. Such an indication may take various forms. For example, the indication may identify an aspect or part of the predictive model and/or workflow that the local analytics device modified (e.g., the parameter that was modified and/or to what the parameter was modified to) and/or may identify the cause of the modification (e.g., the underlying operating data or other asset data that caused the local analytics device to modify and/or a description of the cause). Other examples are also possible.

In some example implementations, a local analytics device and the analytics system 108 may both be involved in individualizing a model-workflow pair, which may be performed in a variety of manners. For example, the analytics system 108 may provide to the local analytics device an instruction to test certain conditions and/or characteristics of the asset 102. Based on the instruction, the local analytics device may execute the tests at the asset 102. For instance, the local analytics device may obtain operating data generated by particular asset sensors and/or actuators. Thereafter, the local analytics device may provide to the analytics system 108 the results from the tested conditions. Based on such results, the analytics system 108 may accordingly define a predictive model and/or workflow for the asset 102 and transmit it to local analytics device for local execution.

In other examples, the local analytics device may perform the same or similar test operations as part of executing a workflow. That is, a particular workflow corresponding to a predictive model may cause the local analytics device to execute certain tests and transmit results to the analytics system 108.

In example implementations, after the local analytics device individualizes a predictive model and/or workflow (or works with the analytics system 108 to do the same), the local analytics device may execute the individualized predictive model and/or workflow instead of the original model and/or workflow (e.g., that which the local analytics device originally received from the analytics system 108). In some cases, although the local analytics device executes the individualized version, the local analytics device may retain the original version of the model and/or workflow in data storage.

In general, an asset executing a predictive model and, based on the resulting output, executing operations of the workflow may facilitate determining a cause or causes of the likelihood of a particular event occurring that is output by the model and/or may facilitate preventing a particular event from occurring in the future. In executing a workflow, an asset may locally determine and take actions to help prevent an event from occurring, which may be beneficial in situations when reliance on the analytics system 108 to make such determinations and provide recommended actions is not efficient or feasible (e.g., when there is network latency, when network connection is poor, when the asset moves out of coverage of the communication network 106, etc.).

In practice, an asset may execute a predictive model in a variety of manners, which may be dependent on the particular predictive model. Figure 9 is a flow diagram 900 depicting one possible example of a local-execution phase that may be used for locally executing a predictive model. The example local-execution phase will be discussed in the context of a health-metric model that outputs a health metric of an asset, but it should be understood that a same or similar local-execution phase may be utilized for other types of predictive models. Moreover, for purposes of illustration, the example local-execution phase is described as being carried out by a local analytics device of the asset 102, but this phase may be carried out by other devices and/or systems as well. One of ordinary skill in the art will appreciate that the flow diagram 900 is provided for sake of clarity and explanation and that numerous other combinations of operations and functions may be utilized to locally execute a predictive model.

As shown in Figure 9, at block 902, the local analytics device may receive data that reflects the current operating conditions of the asset 102. At block 904, the local analytics device may identify, from the received data, the set of operating data that is to be input into the model provided by the analytics system 108. At block 906, the local analytics device may then input the identified set of operating data into the model and run the model to obtain a health metric for the asset 102.

As the local analytics device continues to receive updated operating data for the asset 102, the local analytics device may also continue to update the health metric for the asset 102 by repeating the operations of blocks 902-906 based on the updated operating data. In some cases, the operations of blocks 902-906 may be repeated each time the local analytics device receives new data from sensors and/or actuators of the asset 102 or periodically (e.g., hourly, daily, weekly, monthly, etc.). In this way, local analytics devices may be configured to dynamically update health metrics, perhaps in real-time, as assets are used in operation.

The functions of the example local-execution phase illustrated in Figure 9 will now be described in further detail. At block 902, the local analytics device may receive data that reflects the current operating conditions of the asset 102. Such data may include sensor data from one or more of the sensors of the asset 102, actuator data from one or more actuators of the asset 102, and/or it may include abnormal-condition data, among other types of data.

At block 904, the local analytics device may identify, from the received data, the set of operating data that is to be input into the health-metric model provided by the analytics system 108. This operation may be performed in a number of manners.

In one example, the local analytics device may identify the set of operating data inputs (e.g., data from particular sensors and/or actuators of interest) for the model based on a characteristic of the asset 102, such as asset type or asset class, for which the health metric is being determined. In some cases, the identified set of operating data inputs may be sensor data from some or all of the sensors of the asset 102 and/or actuator data from some of all of the actuators of the asset 102.

In another example, the local analytics device may identify the set of operating data inputs based on the predictive model provided by the analytics system 108. That is, the analytics system 108 may provide some indication to the asset 102 (e.g., either in the predictive model or in a separate data transmission) of the particular inputs for the model. Other examples of identifying the set of operating data inputs are also possible.

At block 906, the local analytics device may then run the health-metric model. Specifically, the local analytics device may input the identified set of operating data into the model, which in turn determines and outputs an overall likelihood of at least one failure occurring within the given timeframe in the future (e.g., the next two weeks).

In some implementations, this operation may involve the local analytics device inputting particular operating data (e.g., sensor and/or actuator data) into one or more individual failure models of the health-metric model, which each may output an individual probability. The local analytics device may then use these individual probabilities, perhaps weighting some more than others in accordance with the health-metric model, to determine the overall likelihood of a failure occurring within the given timeframe in the future.

After determining the overall likelihood of a failure occurring, the local analytics device may convert the probability of a failure occurring into the health metric that may take the form of a single, aggregated parameter that reflects the likelihood that no failures will occur at the asset 102 within the give timeframe in the future (e.g., two weeks). In example implementations, converting the failure probability into the health metric may involve the local analytics device determining the complement of the failure probability. Specifically, the overall failure probability may take the form of a value ranging from zero to one; the health metric may be determined by subtracting one by that number. Other examples of converting the failure probability into the health metric are also possible.

After an asset locally executes a predictive model, the asset may then execute a corresponding workflow based on the resulting output of the executed predictive model. Generally, the asset executing a workflow may involve the local analytics device causing the performance of an operation at the asset (e.g., by sending an instruction to one or more of the asset's on-board systems) and/or the local analytics device causing a computing system, such as the analytics system 108 and/or the output system 110, to execute an operation remote from the asset. As mentioned above, workflows may take various forms and so, workflows may be executed in a variety of manners.

For example, the asset 102 may be caused to internally execute one or more operations that modify some behavior of the asset 102, such as modifying a data-acquisition and/or -transmission scheme, executing a local diagnostic tool, modifying an operating condition of the asset 102 (e.g., modifying a velocity, acceleration, fan speed, propeller angle, air intake, etc. or performing other mechanical operations via one or more actuators of the asset 102), or outputting an indication, perhaps of a relatively low health metric or of recommended preventative actions that should be executed in relation to the asset 102, at a user interface of the asset 102 or to an external computing system.

In another example, the asset 102 may transmit to a system on the communication network 106, such as the output system 110, an instruction to cause the system to carry out an operation, such as generating a work-order or ordering a particular part for a repair of the asset 102. In yet another example, the asset 102 may communicate with a remote system, such as the analytics system 108, that then facilitates causing an operation to occur remote from the asset 102. Other examples of the asset 102 locally executing a workflow are also possible.

E. MODEL/WORKFLOW MODIFICATION PHASE

In another aspect, the analytics system 108 may carry out a modification phase during which the analytics system 108 modifies a deployed model and/or workflow based on new asset data. This phase may be performed for both aggregate and individualized models and workflows.

In particular, as a given asset (e.g., the asset 102) operates in accordance with a model-workflow pair, the asset 102 may provide operating data to the analytics system 108 and/or the data source 112 may provide to the analytics system 108 external data related to the asset 102. Based at least on this data, the analytics system 108 may modify the model and/or workflow for the asset 102 and/or the model and/or workflow for other assets, such as the asset 104. In modifying models and/or workflows for other assets, the analytics system 108 may share information learned from the behavior of the asset 102.

In practice, the analytics system 108 may make modifications in a number of manners. Figure 10 is a flow diagram 1000 depicting one possible example of a modification phase that may be used for modifying model-workflow pairs. For purposes of illustration, the example modification phase is described as being carried out by the analytics system 108, but this modification phase may be carried out by other systems as well. One of ordinary skill in the art will appreciate that the flow diagram 1000 is provided for sake of clarity and explanation and that numerous other combinations of operations may be utilized to modify model-workflow pairs.

As shown in Figure 10, at block 1002, the analytics system 108 may receive data from which the analytics system 108 identifies an occurrence of a particular event. The data may be operating data originating from the asset 102 or external data related to the asset 102 from the data source 112, among other data. The event may take the form of any of the events discussed above, such as a failure at the asset 102.

In other example implementations, the event may take the form of a new component or subsystem being added to the asset 102. Another event may take the form of a "leading indicator" event, which may involve sensors and/or actuators of the asset 102 generating data that differs, perhaps by a threshold differential, from the data identified at block 706 of Figure 7 during the model-definition phase. This difference may indicate that the asset 102 has operating conditions that are above or below normal operating conditions for assets similar to the asset 102. Yet another event may take the form of an event that is followed by one or more leading indicator events.

Based on the identified occurrence of the particular event and/or the underlying data (e.g., operating data and/or external data related to the asset 102), the analytics system 108 may then modify the aggregate, predictive model and/or workflow and/or one or more individualized predictive models and/or workflows. In particular, at block 1004, the analytics system 108 may determine whether to modify the aggregate, predictive model. The analytics system 108 may determine to modify the aggregate, predictive model for a number of reasons.

For example, the analytics system 108 may modify the aggregate, predictive model if the identified occurrence of the particular event was the first occurrence of this particular event for a plurality of assets including the asset 102, such as the first time a particular failure occurred at an asset from a fleet of assets or the first time a particular new component was added to an asset from a fleet of assets.

In another example, the analytics system 108 may make a modification if data associated with the identified occurrence of the particular event is different from data that was utilized to originally define the aggregate model. For instance, the identified occurrence of the particular event may have occurred under operating conditions that had not previously been associated with an occurrence of the particular event (e.g., a particular failure might have occurred with associated sensor values not previously measured before with the particular failure). Other reasons for modifying the aggregate model are also possible.

If the analytics system 108 determines to modify the aggregate, predictive model, the analytics system 108 may do so at block 1006. Otherwise, the analytics system 108 may proceed to block 1008.

At block 1006, the analytics system 108 may modify the aggregate model based at least in part on the data related to the asset 102 that was received at block 1002. In example implementations, the aggregate model may be modified in various manners, such as any manner discussed above with reference to block 510 of Figure 5. In other implementations, the aggregate model may be modified in other manners as well.

At block 1008, the analytics system 108 may then determine whether to modify the aggregate workflow. The analytics system 108 may modify the aggregate workflow for a number of reasons.

For example, the analytics system 108 may modify the aggregate workflow based on whether the aggregate model was modified at block 1004 and/or if there was some other change at the analytics system 108. In other examples, the analytics system 108 may modify the aggregate workflow if the identified occurrence of the event at block 1002 occurred despite the asset 102 executing the aggregate workflow. For instance, if the workflow was aimed to help facilitate preventing the occurrence of the event (e.g., a failure) and the workflow was executed properly but the event still occurred nonetheless, then the analytics system 108 may modify the aggregate workflow. Other reasons for modifying the aggregate workflow are also possible.

If the analytics system 108 determines to modify the aggregate workflow, the analytics system 108 may do so at block 1010. Otherwise, the analytics system 108 may proceed to block 1012.

At block 1010, the analytics system 108 may modify the aggregate workflow based at least in part on the data related to the asset 102 that was received at block 1002. In example implementations, the aggregate workflow may be modified in various manners, such as any manner discussed above with reference to block 514 of Figure 5. In other implementations, the aggregate model may be modified in other manners as well.

At blocks 1012 through blocks 1018, the analytics system 108 may be configured to modify one or more individualized models (e.g., for each of assets 102 and 104) and/or one or more individualized workflows (e.g., for one of asset 102 or asset 104) based at least in part on the data related to the asset 102 that was received at block 1002. The analytics system 108 may do so in a manner similar to blocks 1004-1010.

However, the reasons for modifying an individualized model or workflow may differ from the reasons for the aggregate case. For instance, the analytics system 108 may further consider the underlying asset characteristics that were utilized to define the individualized model and/or workflow in the first place. In a particular example, the analytics system 108 may modify an individualized model and/or workflow if the identified occurrence of the particular event was the first occurrence of this particular event for assets with asset characteristics of the asset 102. Other reasons for modifying an individualized model and/or workflow are also possible.

To illustrate, Figure 6D is a conceptual illustration of a modified model-workflow pair 630. Specifically, the model-workflow pair illustration 630 is a modified version of the aggregate model-workflow pair from Figure 6A. As shown, the modified model-workflow pair illustration 630 includes the original column for model inputs 602 from Figure 6A and includes modified columns for model calculations 634, model output ranges 636, and workflow operations 638. In this example, the modified predictive model has a single input, data from Sensor A, and has two calculations, Calculations I and III. If the output probability of the modified model is less than 75%, then workflow Operation 1 is performed. If the output probability is between 75% and 85%, then workflow Operation 2 is performed. And if the output probability is greater than 85%, then workflow Operation 3 is performed. Other example modified model-workflow pairs are possible and contemplated herein.

Returning to Figure 10, at block 1020, the analytics system 108 may then transmit any model and/or workflow modifications to one or more assets. For example, the analytics system 108 may transmit a modified individualized model-workflow pair to the asset 102 (e.g., the asset whose data caused the modification) and a modified aggregate model to the asset 104. In this way, the analytics system 108 may dynamically modify models and/or workflows based on data associated with the operation of the asset 102 and distribute such modifications to multiple assets, such as the fleet to which the asset 102 belongs. Accordingly, other assets may benefit from the data originating from the asset 102 in that the other assets' local model-workflow pairs may be refined based on such data, thereby helping to create more accurate and robust model-workflow pairs.

While the above modification phase was discussed as being performed by the analytics system 108, in example implementations, the local analytics device of the asset 102 may additionally or alternatively carry out the modification phase in a similar manner as discussed above. For instance, in one example, the local analytics device may modify a model-workflow pair as the asset 102 operates by utilizing operating data generated by one or more sensors and/or actuators. Therefore, the local analytics device of the asset 102, the analytics system 108, or some combination thereof may modify a predictive model and/or workflow as asset-related conditions change. In this way, the local analytics device and/or the analytics system 108 may continuously adapt model-workflow pairs based on the most recent data available to them.

F. DYNAMIC EXECUTION OF MODEL/WORKFLOW

In another aspect, the asset 102 and/or the analytics system 108 may be configured to dynamically adjust executing a model-workflow pair. In particular, the asset 102 and/or the analytics system 108 may be configured to detect certain events that trigger a change in responsibilities with respect to whether the asset 102 and/or the analytics system 108 should be executing the predictive model and/or workflow.

In operation, both the asset 102 and the analytics system 108 may execute all or a part of a model-workflow pair on behalf of the asset 102. For example, after the asset 102 receives a model-workflow pair from the analytics system 108, the asset 102 may store the model-workflow pair in data storage but then may rely on the analytics system 108 to centrally execute part or all of the model-workflow pair. In particular, the asset 102 may provide at least sensor and/or actuator data to the analytics system 108, which may then use such data to centrally execute a predictive model for the asset 102. Based on the output of the model, the analytics system 108 may then execute the corresponding workflow or the analytics system 108 may transmit to the asset 102 the output of the model or an instruction for the asset 102 to locally execute the workflow.

In other examples, the analytics system 108 may rely on the asset 102 to locally execute part or all of the model-workflow pair. Specifically, the asset 102 may locally execute part or all of the predictive model and transmit results to the analytics system 108, which may then cause the analytics system 108 to centrally execute the corresponding workflow. Or the asset 102 may also locally execute the corresponding workflow.

In yet other examples, the analytics system 108 and the asset 102 may share in the responsibilities of executing the model-workflow pair. For instance, the analytics system 108 may centrally execute portions of the model and/or workflow, while the asset 102 locally executes the other portions of the model and/or workflow. The asset 102 and analytics system 108 may transmit results from their respective executed responsibilities. Other examples are also possible.

At some point in time, the asset 102 and/or the analytics system 108 may determine that the execution of the model-workflow pair should be adjusted. That is, one or both may determine that the execution responsibilities should be modified. This operation may occur in a variety of manners.

Figure 11 is a flow diagram 1100 depicting one possible example of an adjustment phase that may be used for adjusting execution of a model-workflow pair. For purposes of illustration, the example adjustment phase is described as being carried out by the asset 102 and/or the analytics system 108, but this modification phase may be carried out by other systems as well.

One of ordinary skill in the art will appreciate that the flow diagram 1100 is provided for sake of clarity and explanation and that numerous other combinations of operations may be utilized to adjust the execution of a model-workflow pair.

At block 1102, the asset 102 and/or the analytics system 108 may detect an adjustment factor (or potentially multiple) that indicates conditions that require an adjustment to the execution of the model-workflow pair. Examples of such conditions include network conditions of the communication network 106 or processing conditions of the asset 102 and/or analytics system 108, among other examples. Example network conditions may include network latency, network bandwidth, signal strength of a link between the asset 102 and the communication network 106, or some other indication of network performance, among other examples. Example processing conditions may include processing capacity (e.g., available processing power), processing usage (e.g., amount of processing power being consumed) or some other indication of processing capabilities, among other examples.

In practice, detecting an adjustment factor may be performed in a variety of manners. For example, this operation may involve determining whether network (or processing) conditions reach one or more threshold values or whether conditions have changed in a certain manner. Other examples of detecting an adjustment factor are also possible.

In particular, in some cases, detecting an adjustment factor may involve the asset 102 and/or the analytics system 108 detecting an indication that a signal strength of a communication link between the asset 102 and the analytics system 108 is below a threshold signal strength or has been decreasing at a certain rate of change. In this example, the adjustment factor may indicate that the asset 102 is about to go “off-line.”

In another case, detecting an adjustment factor may additionally or alternatively involve the asset 102 and/or the analytics system 108 detecting an indication that network latency is above a threshold latency or has been increasing at a certain rate of change. Or the indication may be that a network bandwidth is below a threshold bandwidth or has been decreasing at a certain rate of change. In these examples, the adjustment factor may indicate that the communication network 106 is lagging.

In yet other cases, detecting an adjustment factor may additionally or alternatively involve the asset 102 and/or the analytics system 108 detecting an indication that processing capacity is below a particular threshold or has been decreasing at a certain rate of change and/or that processing usage is above a threshold value or increasing at a certain rate of change. In such examples, the adjustment factor may indicate that processing capabilities of the asset 102 (and/or the analytics system 108) are low. Other examples of detecting an adjustment factor are also possible.

At block 1104, based on the detected adjustment factor, the local execution responsibilities may be adjusted, which may occur in a number of manners. For example, the asset 102 may have detected the adjustment factor and then determined to locally execute the model-workflow pair or a portion thereof. In some cases, the asset 102 may then transmit to the analytics system 108 a notification that the asset 102 is locally executing the predictive model and/or workflow.

In another example, the analytics system 108 may have detected the adjustment factor and then transmitted an instruction to the asset 102 to cause the asset 102 to locally execute the model-workflow pair or a portion thereof. Based on the instruction, the asset 102 may then locally execute the model-workflow pair.

At block 1106, the central execution responsibilities may be adjusted, which may occur in a number of manners. For example, the central execution responsibilities may be adjusted based on the analytics system 108 detecting an indication that the asset 102 is locally executing the predictive model and/or the workflow. The analytics system 108 may detect such an indication in a variety of manners.

In some examples, the analytics system 108 may detect the indication by receiving from the asset 102 a notification that the asset 102 is locally executing the predictive model and/or workflow. The notification may take various forms, such as binary or textual, and may identify the particular predictive model and/or workflow that the asset is locally executing.

In other examples, the analytics system 108 may detect the indication based on received operating data for the asset 102. Specifically, detecting the indication may involve the analytics system 108 receiving operating data for the asset 102 and then detecting one or more characteristics of the received data. From the one or more detected characteristics of the received data, the analytics system 108 may infer that the asset 102 is locally executing the predictive model and/or workflow.

In practice, detecting the one or more characteristics of the received data may be performed in a variety of manners. For instance, the analytics system 108 may detect a type of the received data. In particular, the analytics system 108 may detect a source of the data, such as a particular sensor or actuator that generated sensor or actuator data. Based on the type of the received data, the analytics system 108 may infer that the asset 102 is locally executing the predictive model and/or workflow. For example, based on detecting a sensor-identifier of a particular sensor, the analytics system 108 may infer the that asset 102 is locally executing a predictive model and corresponding workflow that causes the asset 102 to acquire data from the particular sensor and transmit that data to the analytics system 108.

In another instance, the analytics system 108 may detect an amount of the received data. The analytics system 108 may compare that amount to a certain threshold amount of data. Based on the amount reaching the threshold amount, the analytics system 108 may infer that the asset 102 is locally executing a predictive model and/or workflow that causes the asset 102 to acquire an amount of data equivalent to or greater than the threshold amount. Other examples are also possible.

In example implementations, detecting the one or more characteristics of the received data may involve the analytics system 108 detecting a certain change in one or more characteristics of the received data, such as a change in the type of the received data, a change in the amount of data that is received, or change in the frequency at which data is received. In a particular example, a change in the type of the received data may involve the analytics system 108 detecting a change in the source of sensor data that it is receiving (e.g., a change in sensors and/or actuators that are generating the data provided to the analytics system 108).

In some cases, detecting a change in the received data may involve the analytics system 108 comparing recently received data to data received in the past (e.g., an hour, day, week, etc. before a present time). In any event, based on detecting the change in the one or more characteristics of the received data, the analytics system 108 may infer that the asset 102 is locally executing a predictive model and/or workflow that causes such a change to the data provided by the asset 102 to the analytics system 108.

Moreover, the analytics system 108 may detect an indication that the asset 102 is locally executing the predictive model and/or the workflow based on detecting the adjustment factor at block 1102. For example, in the event that the analytics system 108 detects the adjustment factor at block 1102, the analytics system 108 may then transmit to the asset 102 instructions that cause the asset 102 to adjust its local execution responsibilities and accordingly, the analytics system 108 may adjust its own central execution responsibilities. Other examples of detecting the indication are also possible.

In example implementations, the central execution responsibilities may be adjusted in accordance with the adjustment to the local execution responsibilities. For instance, if the asset 102 is now locally executing the predictive model, then the analytics system 108 may accordingly cease centrally executing the predictive model (and may or may not cease centrally executing the corresponding workflow). Further, if the asset 102 is locally executing the corresponding workflow, then the analytics system 108 may accordingly cease executing the workflow (and may or may not cease centrally executing the predictive model). Other examples are also possible.

In practice, the asset 102 and/or the analytics system 108 may continuously perform the operations of blocks 1102-1106. And at times, the local and central execution responsibilities may be adjusted to facilitate optimizing the execution of model-workflow pairs.

Moreover, in some implementations, the asset 102 and/or the analytics system 108 may perform other operations based on detecting an adjustment factor. For example, based on a condition of the communication network 106 (e.g., bandwidth, latency, signal strength, or another indication of network quality), the asset 102 may locally execute a particular workflow. The particular workflow may be provided by the analytics system 108 based on the analytics system 108 detecting the condition of the communication network, may be already stored on the asset 102, or may be a modified version of a workflow already stored on the asset 102 (e.g., the asset 102 may locally modify a workflow). In some cases, the particular workflow may include a data-acquisition scheme that increases or decreases a sampling rate and/or a data-transmission scheme that increases or decreases a transmission rate or amount of data transmitted to the analytics system 108, among other possible workflow operations.

In a particular example, the asset 102 may determine that one or more detected conditions of the communication network have reached respective thresholds (e.g., indicating poor network quality). Based on such a determination, the asset 102 may locally execute a workflow that includes transmitting data according to a data-transmission scheme that reduces the amount and/or frequency of data the asset 102 transmits to the analytics system 108. Other examples are also possible.

V. EXAMPLE METHODS

Turning now to Figure 12, a flow diagram is depicted illustrating an example method 1200 for defining and deploying an aggregate, predictive model and corresponding workflow that may be performed by the analytics system 108. For the method 1200 and the other methods discussed below, the operations illustrated by the blocks in the flow diagrams may be performed in line with the above discussion. Moreover, one or more operations discussed above may be added to a given flow diagram.

At block 1202, the method 1200 may involve the analytics system 108 receiving respective operating data for a plurality of assets (e.g., the assets 102 and 104). At block 1204, the method 1200 may involve the analytics system 108, based on the received operating data, defining a predictive model and a corresponding workflow (e.g., a failure model and corresponding workflow) that are related to the operation of the plurality of assets. At block 1206, the method 1200 may involve the analytics system 108 transmitting to at least one asset of the plurality of assets (e.g., the asset 102) the predictive model and the corresponding workflow for local execution by the at least one asset.

Figure 13 depicts a flow diagram of an example method 1300 for defining and deploying an individualized, predictive model and/or corresponding workflow that may be performed by the analytics system 108. At block 1302, the method 1300 may involve the analytics system 108 receiving operating data for a plurality of assets, where the plurality of assets includes at least a first asset (e.g., the asset 102). At block 1304, the method 1300 may involve the analytics system 108, based on the received operating data, defining an aggregate predictive model and an aggregate corresponding workflow that are related to the operation of the plurality of assets. At block 1306, the method 1300 may involve the analytics system 108 determining one or more characteristics of the first asset. At block 1308, the method 1300 may involve the analytics system 108, based on the one or more characteristics of the first asset and the aggregate predictive model and the aggregate corresponding workflow, defining at least one of an individualized predictive model or an individualized corresponding workflow that is related to the operation of the first asset. At block 1310, the method 1300 may involve the analytics system 108 transmitting to the first asset the defined at least one individualized predictive model or individualized corresponding workflow for local execution by the first asset.

Figure 14 depicts a flow diagram of an example method 1400 for dynamically modifying the execution of model-workflow pairs that may be performed by the analytics system 108. At block 1402, the method 1400 may involve the analytics system 108 transmitting to an asset (e.g., the asset 102) a predictive model and corresponding workflow that are related to the operation of the asset for local execution by the asset. At block 1404, the method 1400 may involve the analytics system 108 detecting an indication that the asset is locally executing at least one of the predictive model or the corresponding workflow. At block 1406, the method 1400 may involve the analytics system 108, based on the detected indication, modifying central execution by the computing system of at least one of the predictive model or the corresponding workflow.

Similar to method 1400, another method for dynamically modifying the execution of model-workflow pairs may be performed by an asset (e.g., the asset 102). For instance, such a method may involve the asset 102 receiving from a central computing system (e.g., the analytics system 108) a predictive model and corresponding workflow that are related to the operation of the asset 102. The method may also involve the asset 102 detecting an adjustment factor indicating one or more conditions associated with adjusting execution of the predictive model and the corresponding workflow. The method may involve, based on the detected adjustment factor, (i) modifying local execution by the asset 102 of at least one of the predictive model or the corresponding workflow and (ii) transmitting to the central computing system an indication that the asset 102 is locally executing the at least one of the predictive model or the corresponding workflow to facilitate causing the central computing system to modify central execution by the computing system of at least one of the predictive model or the corresponding workflow.

Figure 15 depicts a flow diagram of an example method 1500 for locally executing a model-workflow pair, for example, by the local analytics device of the asset 102. At block 1502, the method 1500 may involve the local analytics device receiving, via a network interface, a predictive model that is related to the operation of an asset (e.g. the asset 102) that is coupled to the local analytics device via an asset interface of the local analytics device, where the predictive model is defined by a computing system (e.g., the analytics system 108) located remote from the local analytics device based on operating data for a plurality of assets. At block 1504, the method 1500 may involve the local analytics device receiving, via the asset interface, operating data for the asset 102 (e.g., operating data that is generated by one or more sensors and/or actuators and may be received either indirectly via the asset's central processing unit or directly from the one or more sensors and/or actuators). At block 1506, the method 1500 may involve the local analytics device executing the predictive model based on at least a portion of the received operating data for the asset 102. At block 1508, the method 1500 may involve the local analytics device, based on executing the predictive model, executing a workflow corresponding to the predictive model, where executing the workflow includes causing the asset 102, via the asset interface, to perform an operation.

VI. CONCLUSION

Example embodiments of the disclosed innovations have been described above. Those skilled in the art will understand, however, that changes and modifications may be made to the embodiments described without departing from the true scope and sprit of the present invention, which will be defined by the claims.

Further, to the extent that examples described herein involve operations performed or initiated by actors, such as "humans", "operators", "users" or other entities, this is for purposes of example and explanation only. The claims should not be construed as requiring action by such actors unless explicitly recited in the claim language.

CLAIMS

1. A computing system comprising:

at least one processor;

a non-transitory computer-readable medium; and

program instructions stored on the non-transitory computer-readable medium that are executable by the at least one processor to cause the computing system to:

receive respective operating data for a plurality of assets;

based on the received operating data, define a predictive model and a corresponding workflow that are related to the operation of the plurality of assets; and

transmit to at least one asset of the plurality of assets the predictive model and the corresponding workflow for local execution by the at least one asset.

2. The computing system of claim 1, wherein the respective operating data comprises (i) abnormal-condition data associated with a failure that occurred at a given asset at a particular time and (ii) at least one of sensor or actuator data indicating at least one operating condition of the given asset at the particular time.

3. The computing system of claim 1, wherein the predictive model is defined to output a probability that a particular event will occur at a given asset within a period of time into the future.

4. The computing system of claim 3, wherein the corresponding workflow comprises one or more operations to be performed based on the determined probability.

5. The computing system of claim 1, wherein the corresponding workflow comprises a given asset controlling one or more actuators of the given asset to facilitate modifying an operating condition of the given asset.

6. The computing system of claim 1, wherein the corresponding workflow comprises one or more diagnostic tools to be executed locally by a given asset.

7. The computing system of claim 1, wherein the corresponding workflow comprises acquiring sensor data according to a data-acquisition scheme.

8. The computing system of claim 7, wherein the data-acquisition scheme indicates one or more sensors of a given asset from which data is acquired.

9. The computing system of claim 8, wherein the data-acquisition scheme further indicates an amount of data that the given asset will acquire from each of the one or more sensors.

10. The computing system of claim 1, wherein the corresponding workflow comprises transmitting data to the computing system according to a data-transmission scheme.

11. The computing system of claim 10, wherein the data-transmission scheme indicates a frequency at which a given asset transmits data to the computing system.

12. The computing system of claim 1, wherein the computing system is a first computing system, and wherein the corresponding workflow comprises a given asset transmitting instructions to a second computing system to facilitate causing the second computing system to carry out an operation related to the given asset.

13. The computing system of claim 1, wherein the at least one asset of the plurality of assets comprises a first asset and a second asset, and wherein transmitting the predictive model and the corresponding workflow comprises transmitting to the first asset and the second asset the predictive model and the corresponding workflow.

14. A non-transitory computer-readable medium having instructions stored thereon that are executable to cause a computing system to:

receive respective operating data for a plurality of assets;

based on the received operating data, define a predictive model and a corresponding workflow that are related to the operation of the plurality of assets; and

transmit to at least one asset of the plurality of assets the predictive model and the corresponding workflow for local execution by the at least one asset.

15. The non-transitory computer-readable medium of claim 14, wherein the predictive model is defined to output a probability that a particular event will occur at a given asset within a period of time into the future.

16. The non-transitory computer-readable medium of claim 14, wherein the corresponding workflow comprises a given asset controlling one or more actuators of the given asset to facilitate modifying an operating condition of the given asset.

17. The non-transitory computer-readable medium of claim 14, wherein the corresponding workflow comprises one or more diagnostic tools to be executed locally by a given asset.

18. The non-transitory computer-readable medium of claim 14, wherein the computing system is a first computing system, and wherein the corresponding workflow comprises a given asset transmitting instructions to a second computing system to facilitate causing the second computing system to carry out an operation related to the given asset.

19. A computer-implemented method comprising:

receiving respective operating data for a plurality of assets;

based on the received operating data, defining a predictive model and a corresponding workflow that are related to the operation of the plurality of assets; and

transmitting to at least one asset of the plurality of assets the predictive model and the corresponding workflow for local execution by the at least one asset.

20. The computer-implemented method of claim 19, wherein the corresponding workflow comprises acquiring sensor data according to a data-acquisition scheme, wherein the data-acquisition scheme indicates one or more sensors of a given asset from which data is acquired.

21. A computing system comprising:

at least one processor;

a non-transitory computer-readable medium; and

program instructions stored on the non-transitory computer-readable medium that are executable by the at least one processor to cause the computing system to:

receive operating data for a plurality of assets, wherein the plurality of assets comprises a first asset;

based on the received operating data, define an aggregate predictive model and an aggregate corresponding workflow that are related to the operation of the plurality of assets;

determine one or more characteristics of the first asset;

based on the one or more characteristics of the first asset and the aggregate predictive model and the aggregate corresponding workflow, define at least one of an individualized predictive model or an individualized corresponding workflow that is related to the operation of the first asset; and

transmit to the first asset the defined at least one individualized predictive model or individualized corresponding workflow for local execution by the first asset.

22. The computing system of claim 21, wherein the one or more characteristics of the first asset comprises at least one of an asset age or an asset health.

23. The computing system of claim 21, wherein determining the one or more characteristics of the first asset comprises determining the one or more characteristics of the first asset based on received operating data for the first asset.

24. The computing system of claim 21, wherein defining at least one of an individualized predictive model or an individualized corresponding workflow comprises defining the individualized predictive model and the individualized corresponding workflow, and wherein transmitting the at least one individualized predictive model or individualized corresponding workflow comprises transmitting the individualized predictive model and the individualized corresponding workflow.

25. The computing system of claim 21, wherein defining at least one of an individualized predictive model or an individualized corresponding workflow comprises defining the individualized corresponding workflow, and wherein transmitting the at least one individualized predictive model or individualized corresponding workflow comprises transmitting the aggregate predictive model and the individualized corresponding workflow.

26. The computing system of claim 25, wherein the aggregate corresponding workflow comprises a first operation, and wherein the individualized corresponding workflow comprises a second operation that differs from the first operation.

27. The computing system of claim 26, wherein the first operation comprises acquiring data according to a first acquisition scheme, and wherein the second operation comprises acquiring data according to a second acquisition scheme.

28. The computing system of claim 26, wherein the first operation comprises acquiring data according to an acquisition scheme, and wherein the second operation comprises executing one or more diagnostic tools.

29. The computing system of claim 21, wherein the plurality of assets further comprises a second asset, and wherein the program instructions further comprise instructions that are executable to cause the computing system to:

after transmitting the at least one individualized predictive model or individualized corresponding workflow, receive operating data for the second asset indicating an occurrence of an event at the second asset;

based on the received operating data for the second asset, modify the at least one individualized predictive model or individualized corresponding workflow; and

transmit to the first asset the modified at least one individualized predictive model or individualized corresponding workflow.

30. A non-transitory computer-readable medium having instructions stored thereon that are executable to cause a computing system to:

receive operating data for a plurality of assets, wherein the plurality of assets comprises a first asset;

based on the received operating data, define an aggregate predictive model and an aggregate corresponding workflow that are related to the operation of the plurality of assets;

determine one or more characteristics of the first asset;

based on the one or more characteristics of the first asset and the aggregate predictive model and the aggregate corresponding workflow, define at least one of an individualized predictive model or an individualized corresponding workflow that is related to the operation of the first asset; and

transmit to the first asset the defined at least one individualized predictive model or individualized corresponding workflow for local execution by the first asset.

31. The non-transitory computer-readable medium of claim 30, wherein defining at least one of an individualized predictive model or an individualized corresponding workflow comprises defining the individualized predictive model and the individualized corresponding workflow, and wherein transmitting the at least one individualized predictive model or

individualized corresponding workflow comprises transmitting the individualized predictive model and the individualized corresponding workflow.

32. The non-transitory computer-readable medium of claim 30, wherein defining at least one of an individualized predictive model or an individualized corresponding workflow comprises defining the individualized corresponding workflow, and wherein transmitting the at least one individualized predictive model or individualized corresponding workflow comprises transmitting the aggregate predictive model and the individualized corresponding workflow.

33. The non-transitory computer-readable medium of claim 32, wherein the aggregate corresponding workflow comprises a first operation, and wherein the individualized corresponding workflow comprises a second operation that differs from the first operation.

34. The non-transitory computer-readable medium of claim 33, wherein the first operation comprises acquiring data according to a first acquisition scheme, and wherein the second operation comprises acquiring data according to a second acquisition scheme.

35. The non-transitory computer-readable medium of claim 33, wherein the first operation comprises acquiring data according to an acquisition scheme, and wherein the second operation comprises executing one or more diagnostic tools.

36. The non-transitory computer-readable medium of claim 30, wherein the plurality of assets further comprises a second asset, and wherein the program instructions further comprise instructions that are executable to cause the computing system to:

after transmitting the at least one individualized predictive model or individualized corresponding workflow, receive operating data for the second asset indicating an occurrence of an event at the second asset;

based on the received operating data for the second asset, modify the at least one individualized predictive model or individualized corresponding workflow; and

transmit to the first asset the modified at least one individualized predictive model or individualized corresponding workflow.

37. A computer-implemented method comprising:

receiving operating data for a plurality of assets, wherein the plurality of assets comprises a first asset;

based on the received operating data, defining an aggregate predictive model and an aggregate corresponding workflow that are related to the operation of the plurality of assets;

determining one or more characteristics of the first asset;

based on the one or more characteristics of the first asset and the aggregate predictive model and the aggregate corresponding workflow, defining at least one of an individualized predictive model or an individualized corresponding workflow that is related to the operation of the first asset; and

transmitting to the first asset the defined at least one individualized predictive model or individualized corresponding workflow for local execution by the first asset.

38. The computer-implemented method of claim 37, wherein defining at least one of an individualized predictive model or an individualized corresponding workflow comprises defining the individualized corresponding workflow, and wherein transmitting the at least one individualized predictive model or individualized corresponding workflow comprises transmitting the aggregate predictive model and the individualized corresponding workflow.

39. The computer-implemented method of claim 38, wherein the aggregate corresponding workflow comprises a first operation, and wherein the individualized corresponding workflow comprises a second operation that differs from the first operation.

40. The computer-implemented method of claim 39, wherein one of the first operation or the second operation comprises executing one or more diagnostic tools.

41. A computing device comprising:

an asset interface configured to couple the computing device to an asset;

a network interface configured to facilitate communication between the computing device and a computing system located remote from the computing device;

at least one processor;

a non-transitory computer-readable medium; and

program instructions stored on the non-transitory computer-readable medium that are executable by the at least one processor to cause the computing device to:

receive, via the network interface, a predictive model that is related to the operation of the asset, wherein the predictive model is defined by the computing system based on operating data for a plurality of assets;

receive, via the asset interface, operating data for the asset;

execute the predictive model based on at least a portion of the received operating data for the asset; and

based on executing the predictive model, execute a workflow corresponding to the predictive model, wherein executing the workflow comprises causing the asset, via the asset interface, to perform an operation.

42. The computing device of claim 41, wherein the asset interface communicatively couples the computing device to an on-asset computer of the asset.

43. The computing device of claim 41, wherein the asset comprises an actuator, and wherein executing the workflow comprises causing the actuator to perform a mechanical operation.

44. The computing device of claim 41, wherein executing the workflow comprises causing the asset to execute a diagnostic tool.

45. The computing device of claim 41, wherein executing the workflow further comprises causing, via the network interface, execution of an operation remote from the asset.

46. The computing device of claim 45, wherein causing execution of an operation remote from the asset comprises instructing the computing system to execute an operation remote from the asset.

47. The computing device of claim 41, wherein the program instructions stored on the non-transitory computer-readable medium are further executable by the at least one processor to cause the computing device to:

before executing the predictive model, individualize the predictive model.

48. The computing device of claim 47, wherein individualizing the predictive model comprises modifying one or more parameters of the predictive model based at least on received operating data for the asset.

49. The computing device of claim 47, wherein the program instructions stored on the non-transitory computer-readable medium are further executable by the at least one processor to cause the computing device to:

after individualizing the predictive model, transmit to the computing system, via the network interface, an indication that the predictive model has been individualized.

50. The computing device of claim 41, wherein the predictive model is a first predictive model, and wherein the program instructions stored on the non-transitory computer-readable medium are further executable by the at least one processor to cause the computing device to:

before executing the first predictive model, transmit to the computing system, via the network interface, a given subset of the received operating data for the asset, wherein the given subset of received operating data comprises operating data generated by a given group of one or more sensors.

51. The computing device of claim 50, wherein the program instructions stored on the non-transitory computer-readable medium are further executable by the at least one processor to cause the computing device to:

after transmitting the given subset of the received operating data for the asset, receive a second predictive model that is related to the operation of the asset, wherein the second predictive model is defined by the computing system based on the given subset of the received operating data for the asset; and

execute the second predictive model instead of the first predictive model.

52. A non-transitory computer-readable medium having instructions stored thereon that are executable to cause a computing device coupled to an asset via an asset interface of the computing device to:

receive, via a network interface of the computing device configured to facilitate communication between the computing device and a computing system located remote from the computing device, a predictive model that is related to the operation of the asset, wherein the predictive model is defined by the computing system based on operating data for a plurality of assets;

receive, via the asset interface, operating data for the asset;

execute the predictive model based on at least a portion of the received operating data for the asset; and

based on executing the predictive model, execute a workflow corresponding to the predictive model, wherein executing the workflow comprises causing the asset, via the asset interface, to perform an operation.

53. The non-transitory computer-readable medium of claim 52, wherein the program instructions stored on the non-transitory computer-readable medium are further executable to cause the computing device to:

before executing the predictive model, individualize the predictive model.

54. The non-transitory computer-readable medium of claim 53, wherein individualizing the predictive model comprises modifying one or more parameters of the predictive model based at least on received operating data for the asset.

55. The non-transitory computer-readable medium of claim 52, wherein the predictive model is a first predictive model, and wherein the program instructions stored on the non-transitory computer-readable medium are further executable to cause the computing device to:

before executing the first predictive model, transmit to the computing system, via the network interface, a given subset of the received operating data for the asset, wherein the given subset of received operating data comprises operating data generated by a given group of one or more sensors.

56. The non-transitory computer-readable medium of claim 55, wherein the program instructions stored on the non-transitory computer-readable medium are further executable to cause the computing device to:

after transmitting the operating data from the particular group of the one or more sensors, receive a second predictive model that is related to the operation of the asset, wherein the second predictive model is defined by the computing system based on the given subset of the received operating data for the asset; and

execute the second predictive model instead of the first model.

57. A computer-implemented method, the method comprising:

receiving, via a network interface of a computing device that is coupled to an asset via an asset interface of the computing device, a predictive model that is related to the operation of the asset, wherein the predictive model is defined by a computing system located remote from the computing device based on operating data for a plurality of assets;

receiving, by the computing device via the asset interface, operating data for the asset;

executing, by the computing device, the predictive model based on at least a portion of the received operating data for the asset; and

based on executing the predictive model, executing, by the computing device, a workflow corresponding to the predictive model, wherein executing the workflow comprises causing the asset, via the asset interface, to perform an operation.

58. The computer-implemented method of claim 57, the method further comprising:

before executing the predictive model, individualizing, by the computing device, the predictive model.

59. The computer-implemented method of claim 58, wherein individualizing the predictive model comprises modifying one or more parameters of the predictive model based at least on received operating data for the asset.

60. The computer-implemented method of claim 57, wherein executing the workflow further comprises causing, via the network interface, execution of an operation remote from the asset.

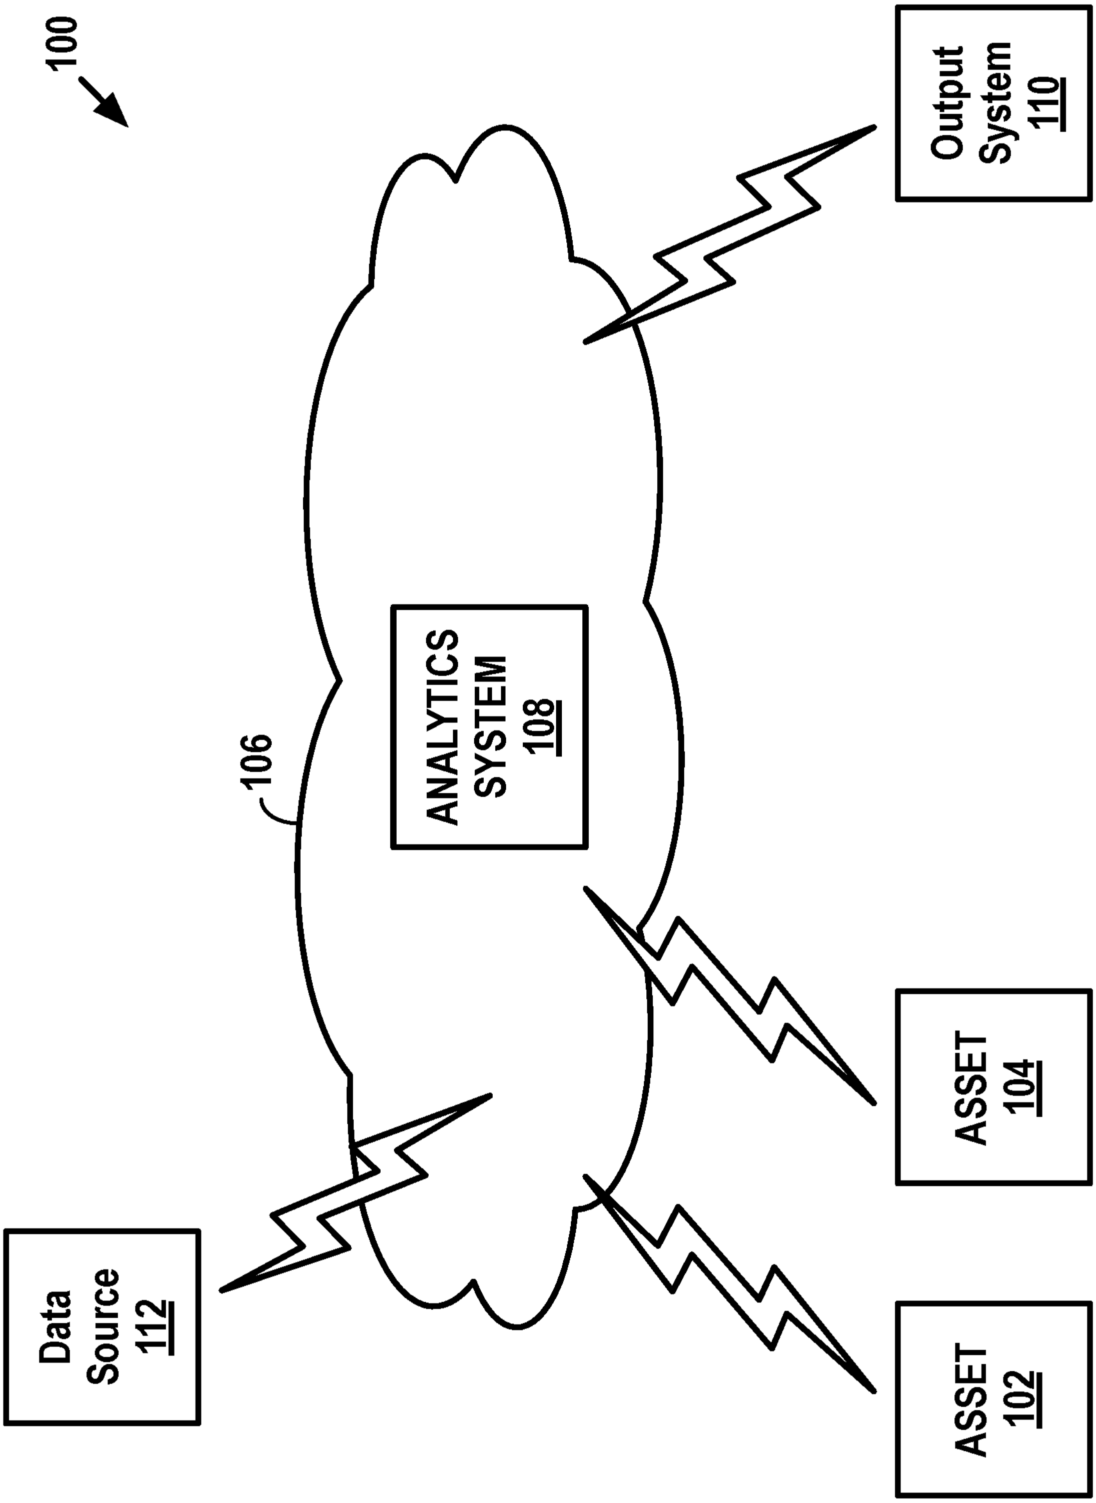
100
106
ANALYTICS SYSTEM 108
Output System 110
ASSET 104
ASSET 102
Data Source 112

FIGURE 1

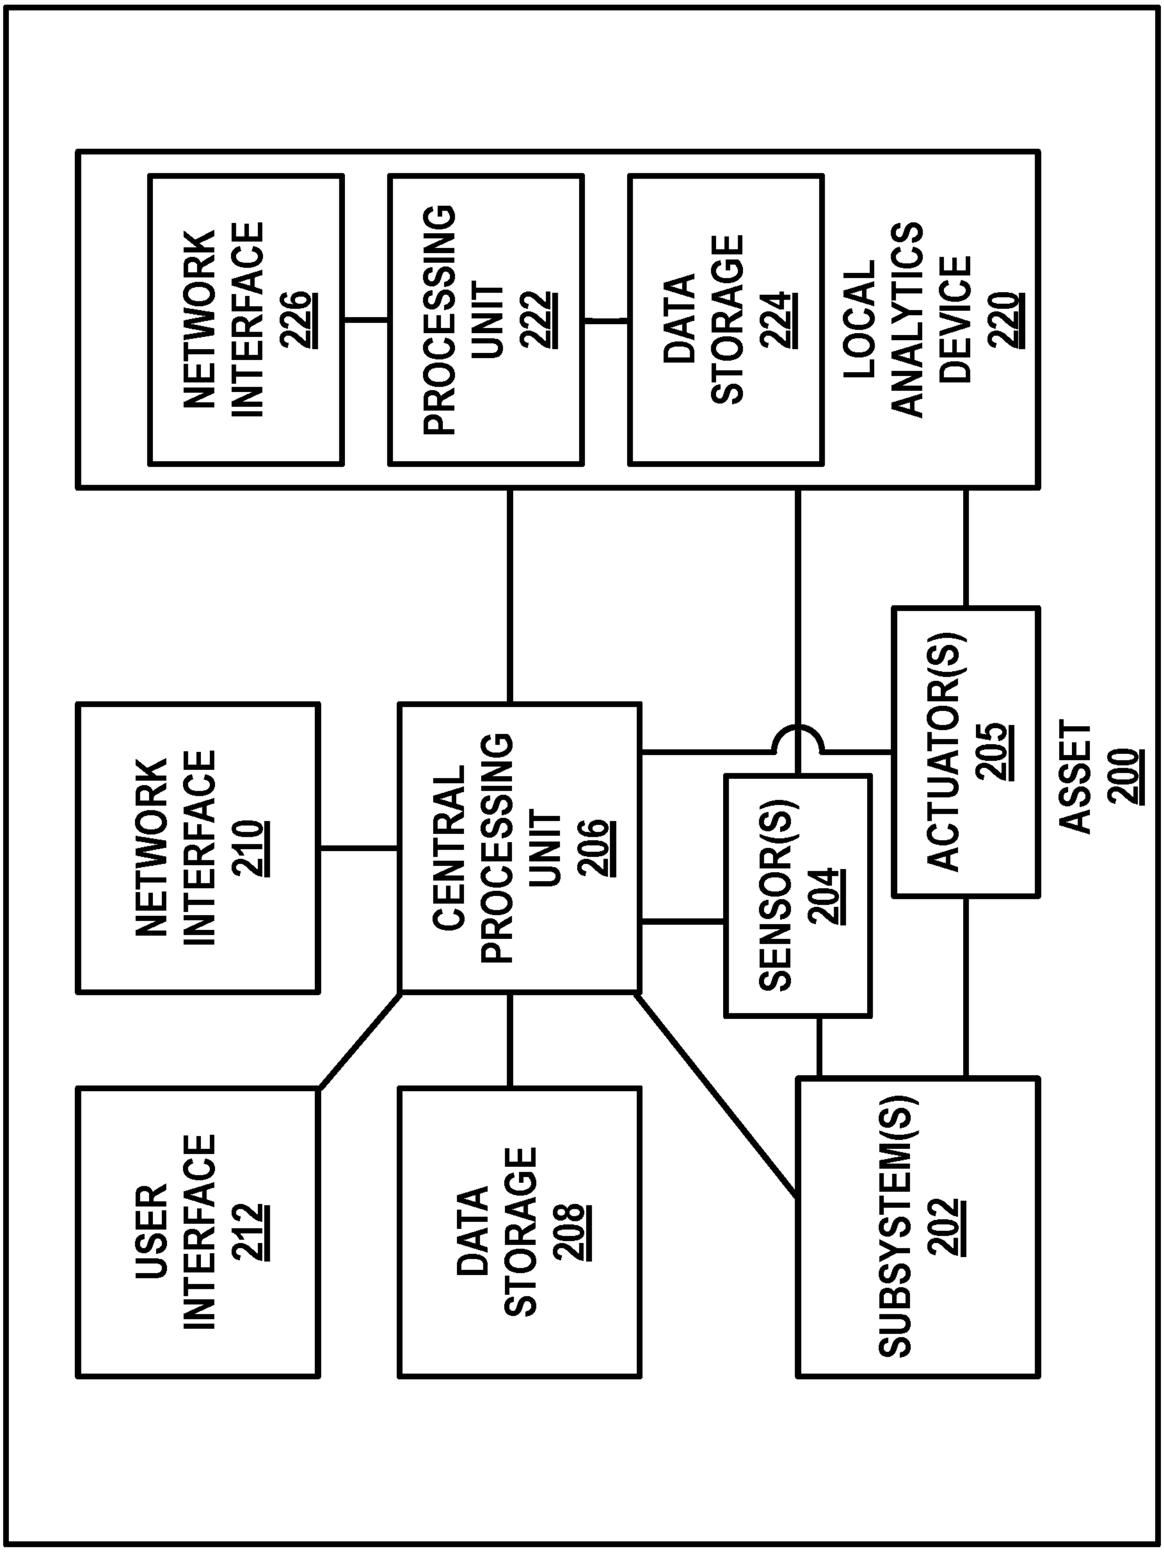
NETWORK INTERFACE 226
PROCESSING UNIT 222
DATA STORAGE 224
LOCAL ANALYTICS DEVICE 220
NETWORK INTERFACE 210
CENTRAL PROCESSING UNIT 206
SENSOR(S) 204
ACTUATOR(S) 205
ASSET 200
USER INTERFACE 212
DATA STORAGE 208
SUBSYSTEM(S) 202

FIGURE 2

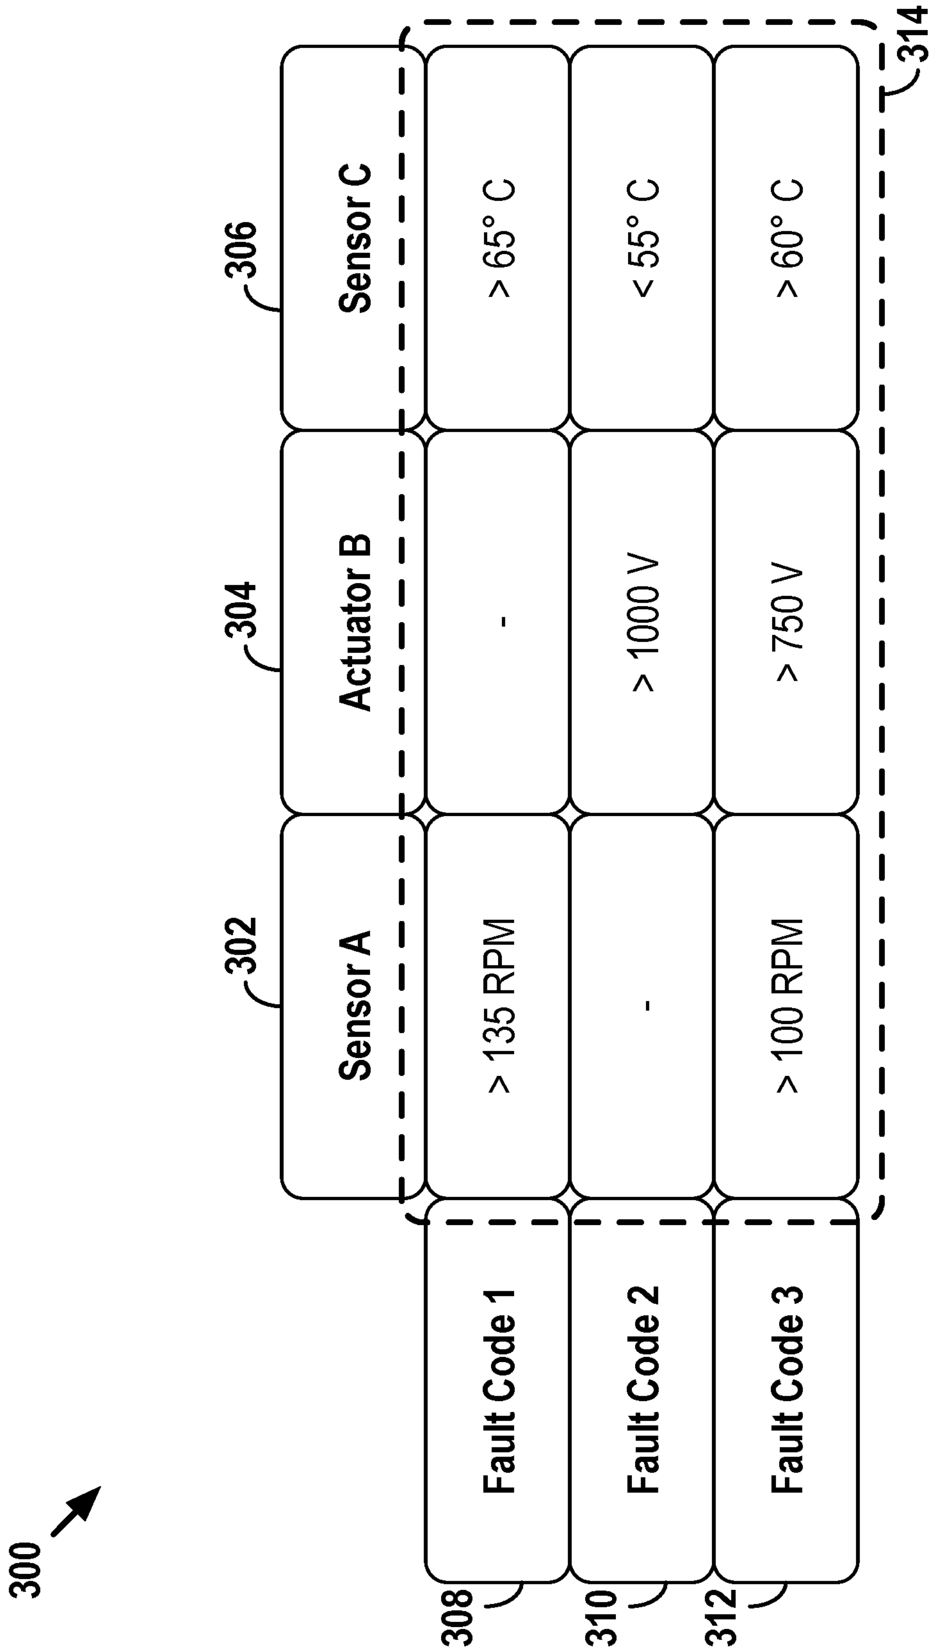
300
302
304
306
Sensor A
Actuator B
Sensor C
> 135 RPM
-
> 65° C
-
> 1000 V
< 55° C
> 100 RPM
> 750 V
> 60° C
Fault Code 1
Fault Code 2
Fault Code 3
308
310
312
314

FIGURE 3

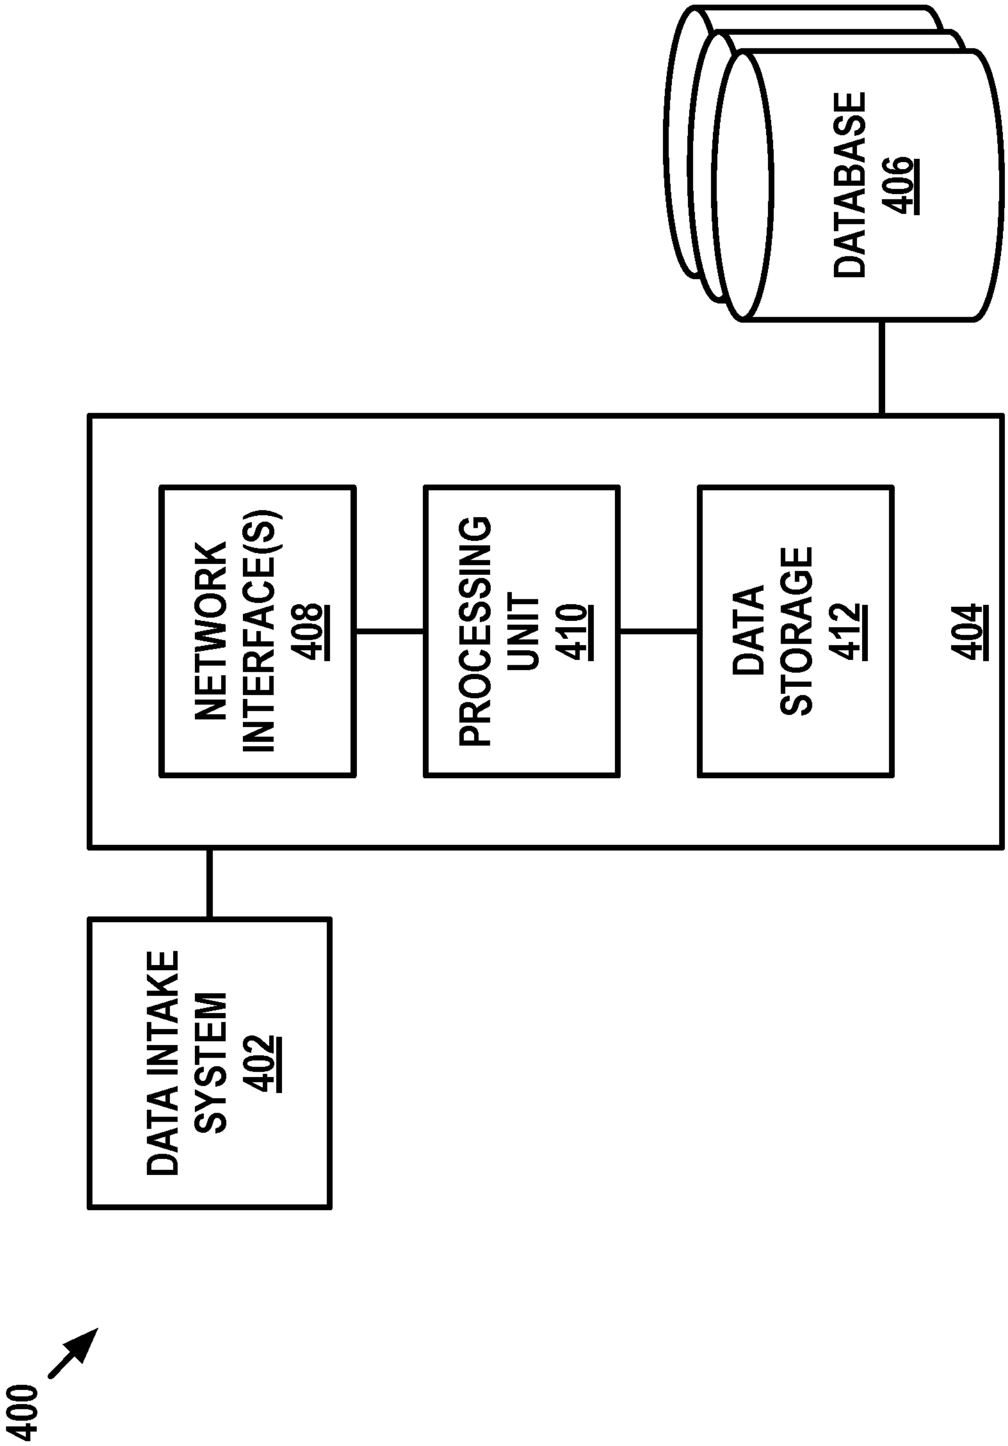
DATABASE
406
NETWORK INTERFACE(S)
408
PROCESSING UNIT
410
DATA STORAGE
412
404
DATA INTAKE SYSTEM
402
400

FIGURE 4

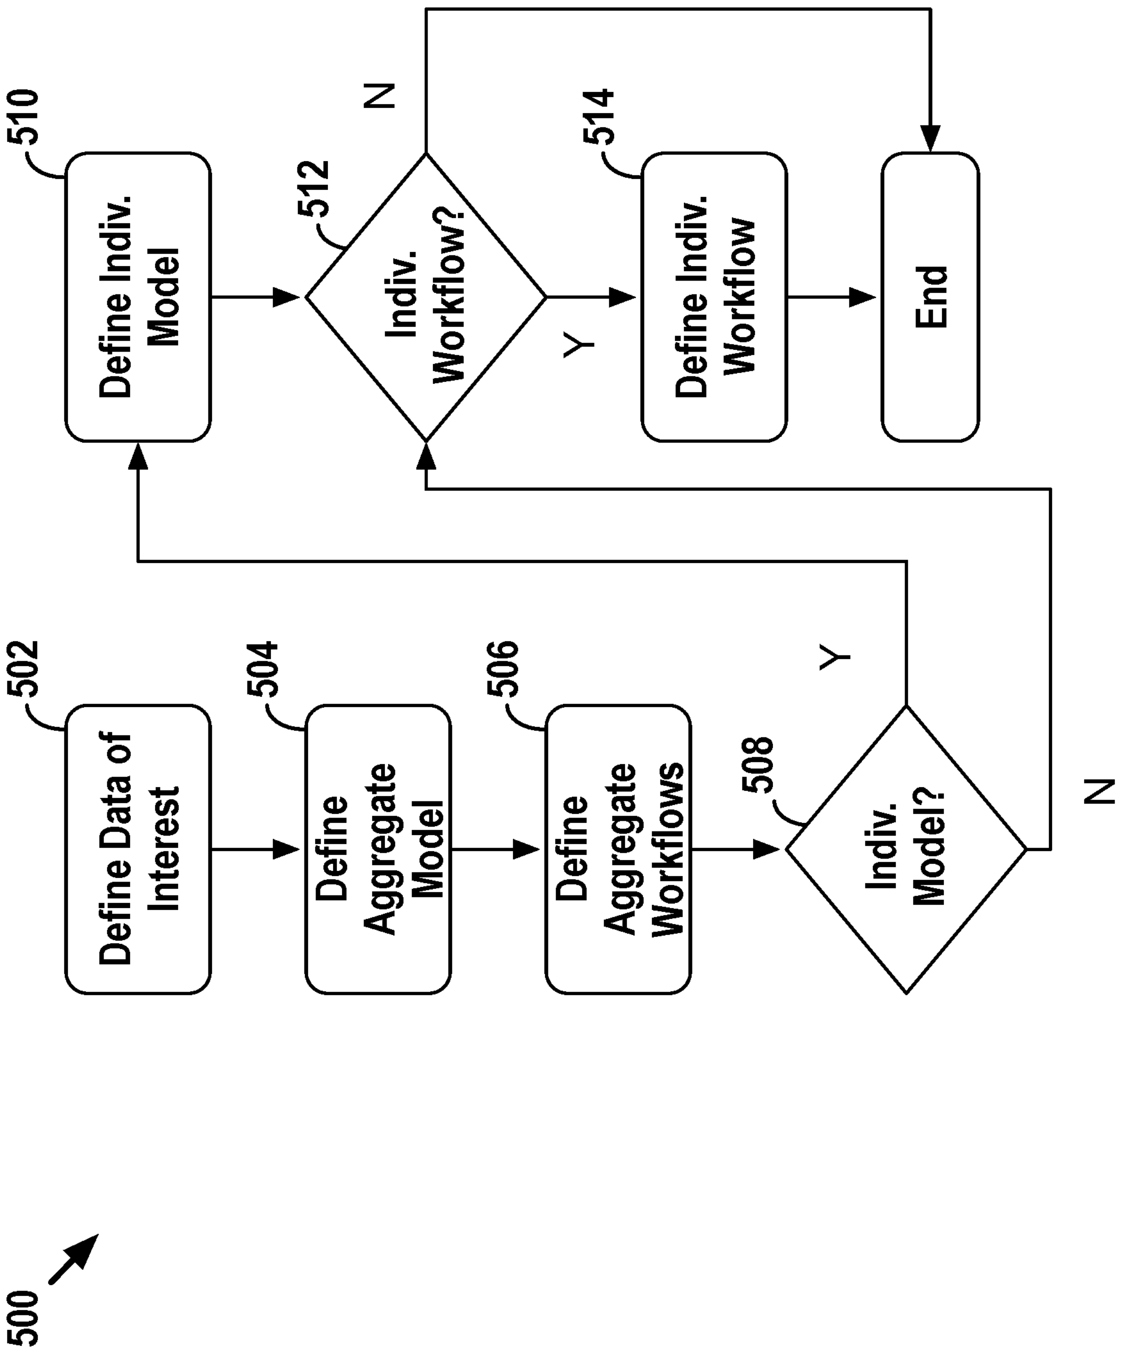
500
Define Data of Interest
502
Define Aggregate Model
504
Define Aggregate Workflows
506
Indiv. Model?
508
Y
N
Define Indiv. Model
510
Indiv. Workflow?
512
Y
N
Define Indiv. Workflow
514
End

FIGURE 5

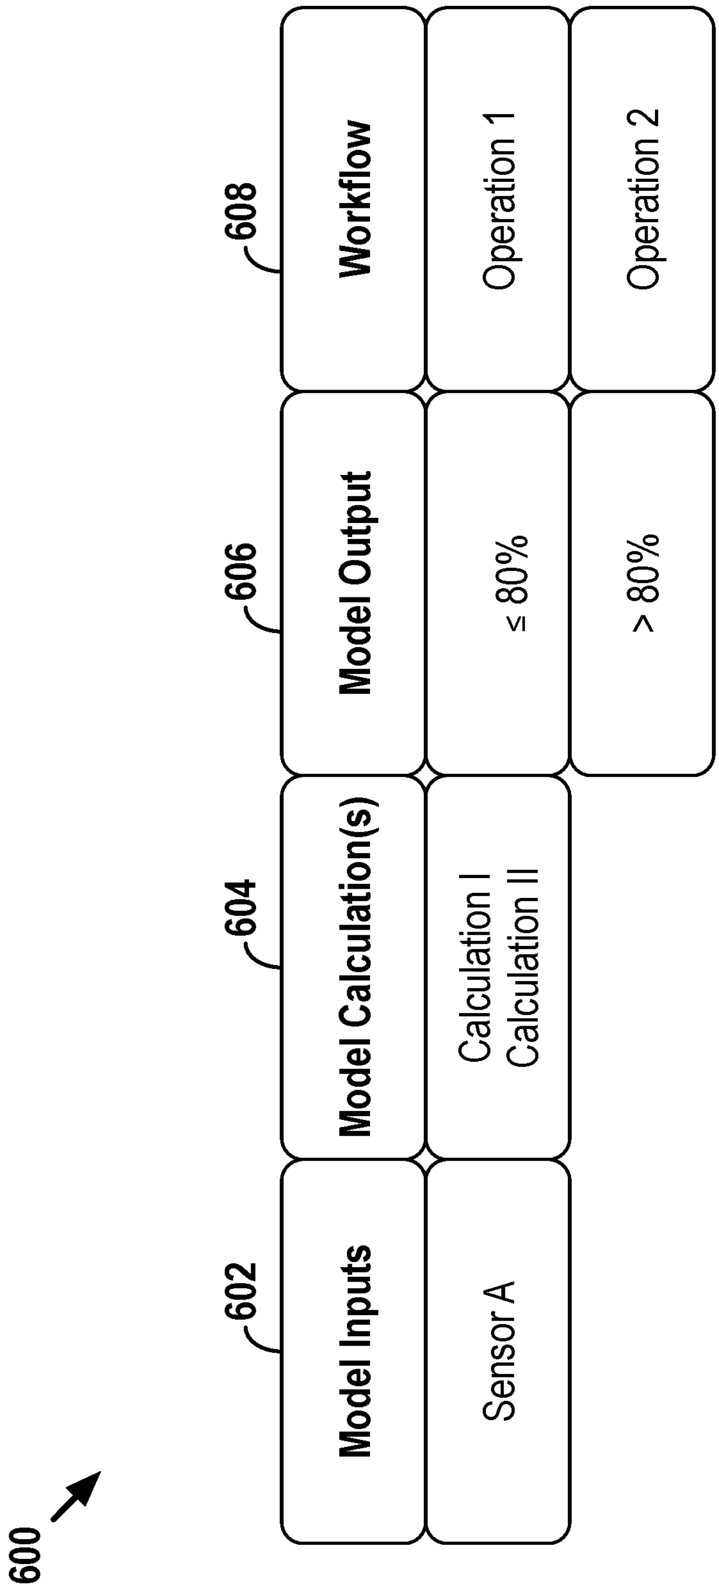
600
602
Model Inputs
Sensor A
604
Model Calculation(s)
Calculation I
Calculation II
606
Model Output
≤ 80%
> 80%
608
Workflow
Operation 1
Operation 2

FIGURE 6A

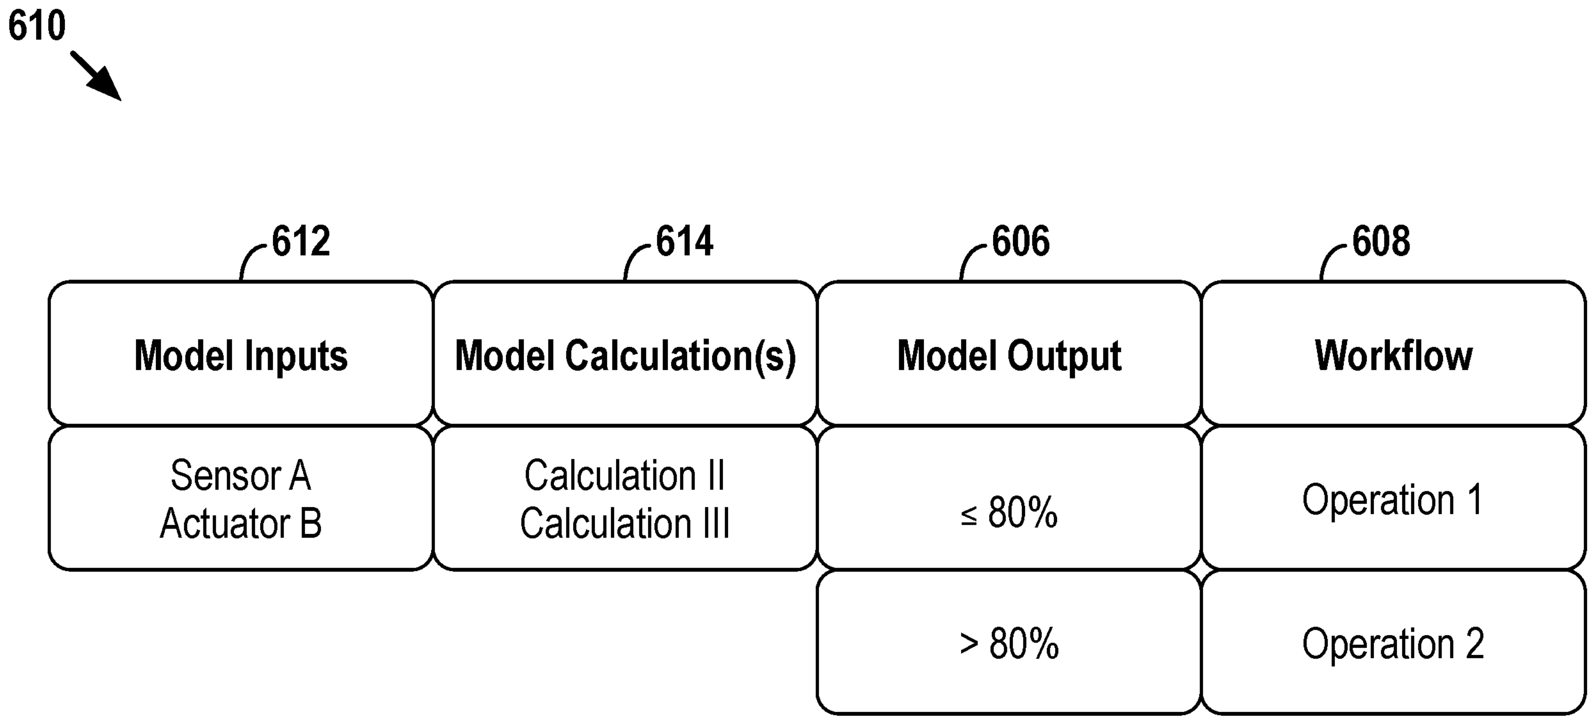
610
612
614
606
608
Model Inputs
Model Calculation(s)
Model Output
Workflow
Sensor A
Actuator B
Calculation II
Calculation III
≤ 80%
Operation 1
> 80%
Operation 2

FIGURE 6B

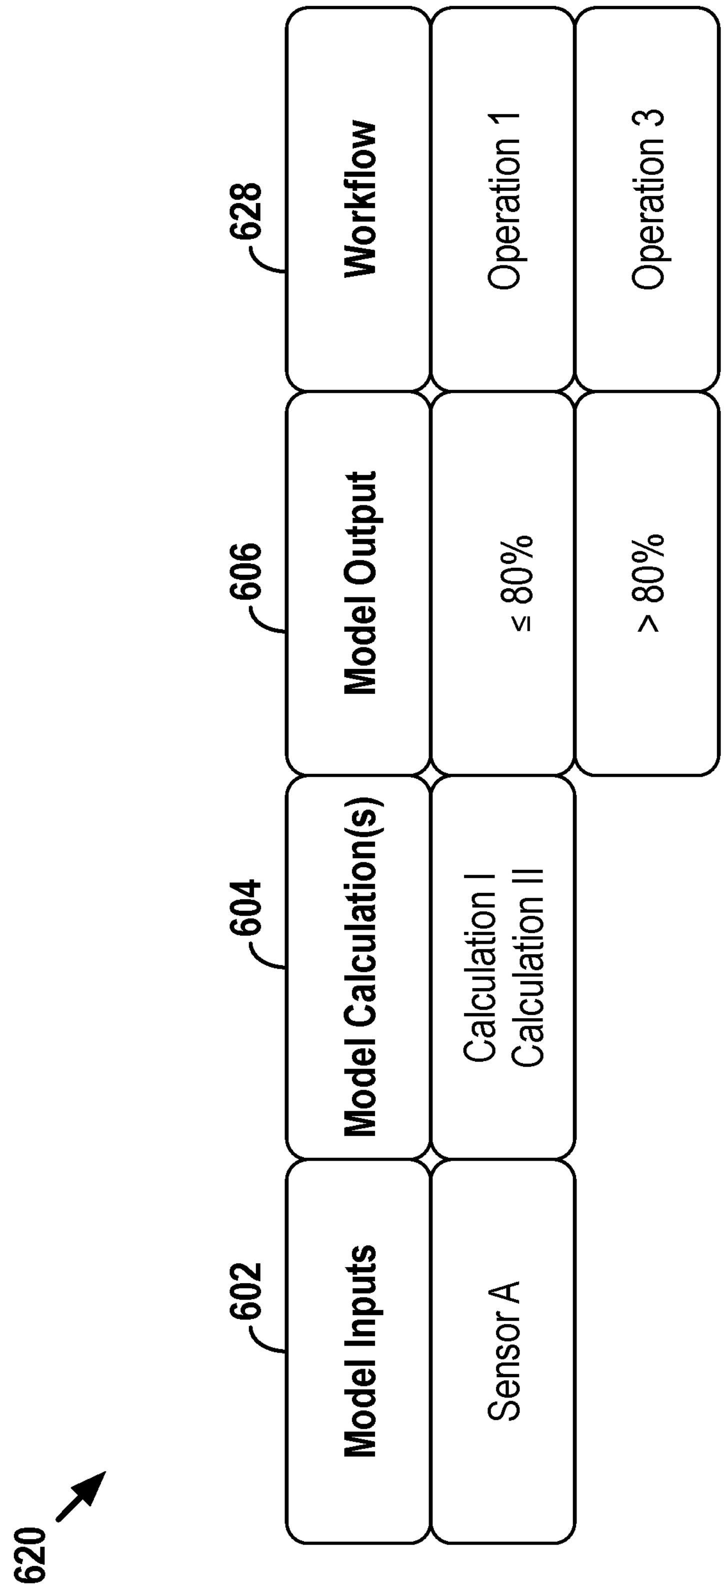
620
602
Model Inputs
Sensor A
604
Model Calculation(s)
Calculation I
Calculation II
606
Model Output
≤ 80%
> 80%
628
Workflow
Operation 1
Operation 3

FIGURE 6C

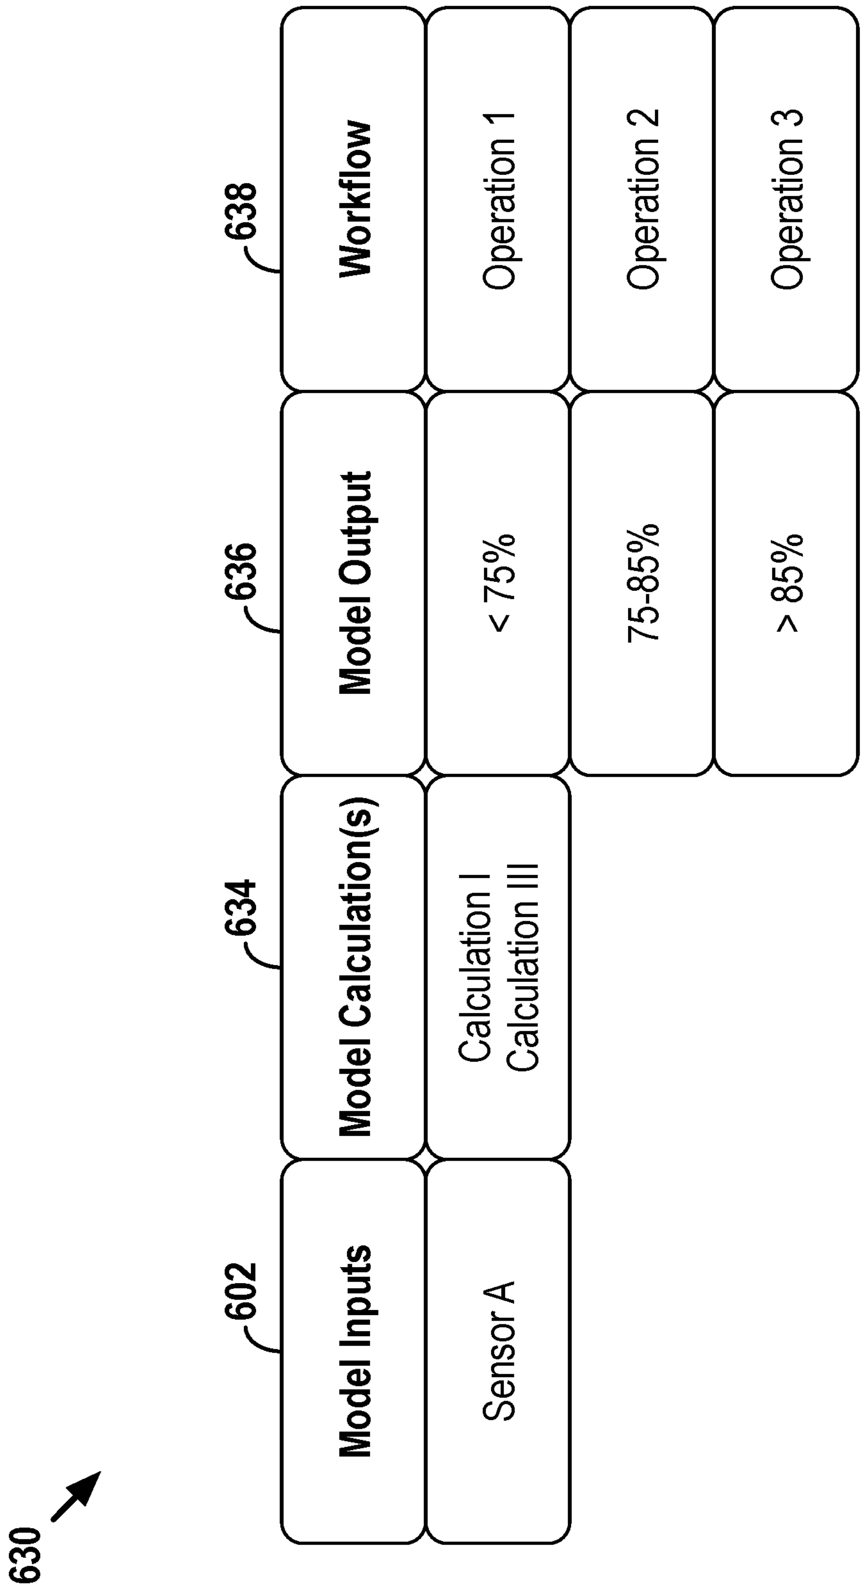
630
602
Model Inputs
Sensor A
634
Model Calculation(s)
Calculation I
Calculation III
636
Model Output
< 75%
75-85%
> 85%
638
Workflow
Operation 1
Operation 2
Operation 3

FIGURE 6D

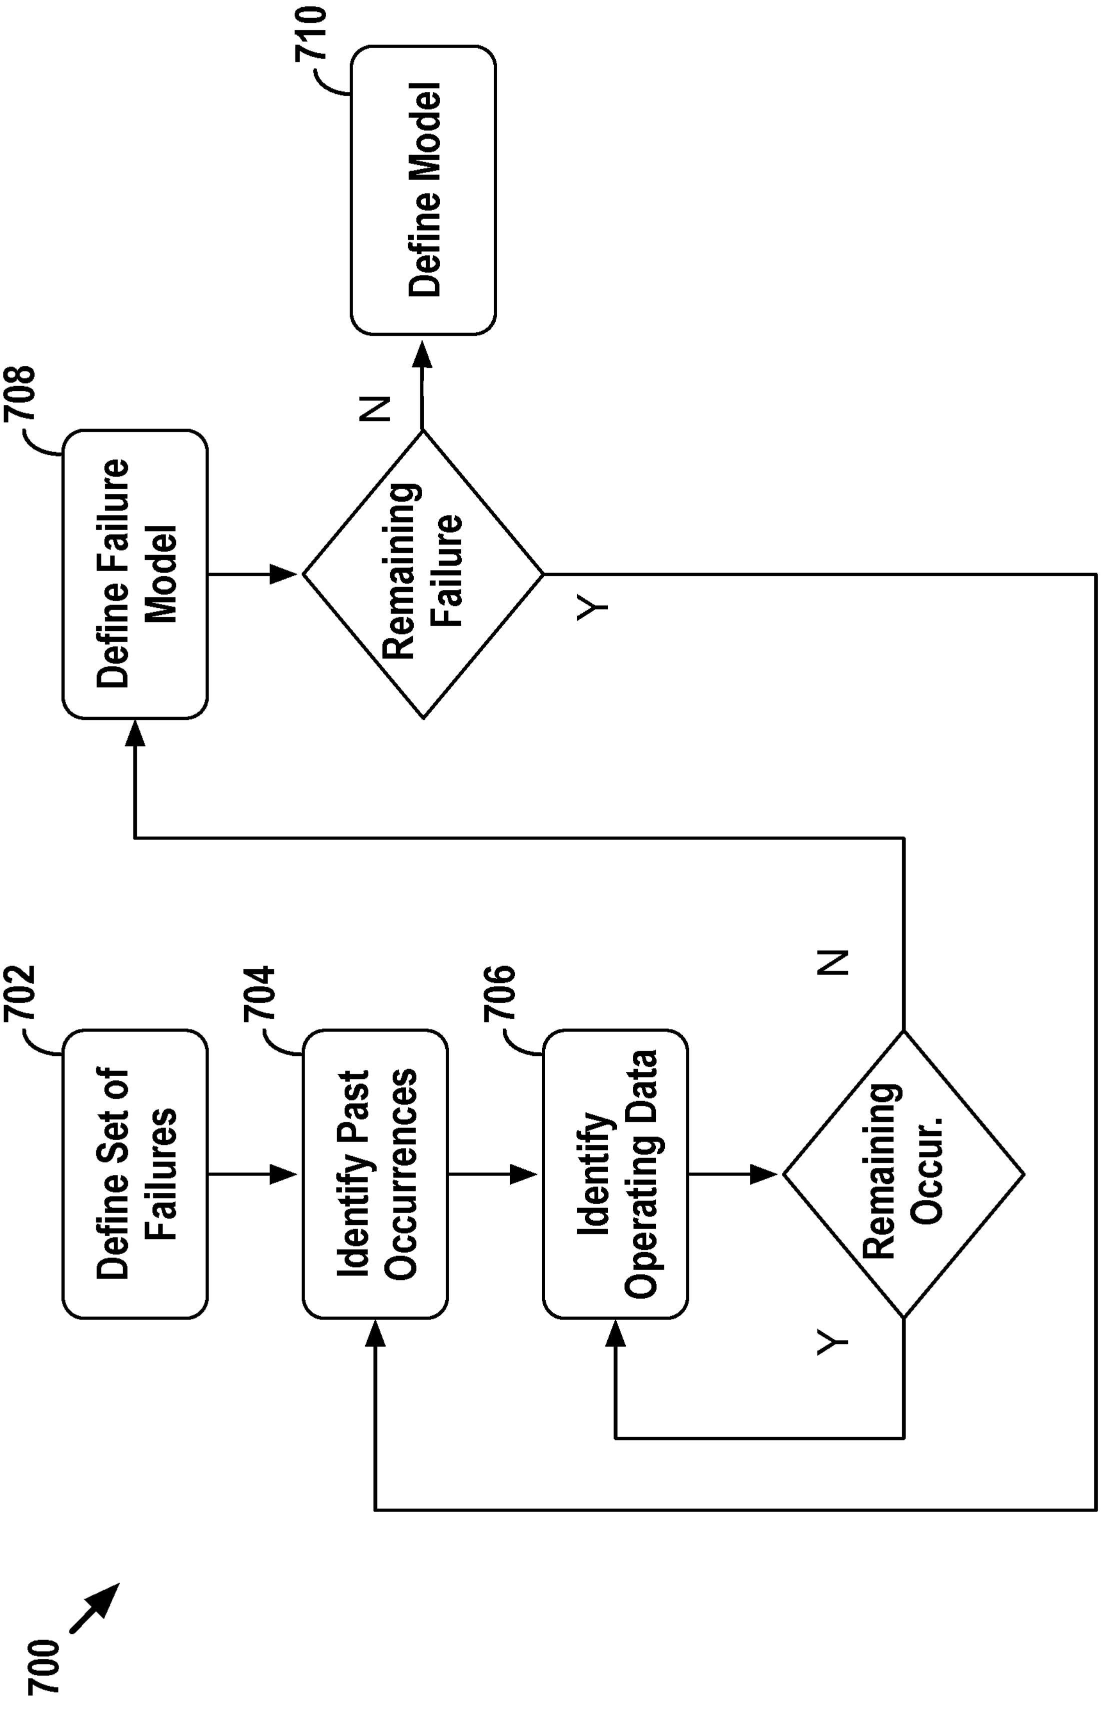
700
702
Define Set of Failures
704
Identify Past Occurrences
706
Identify Operating Data
Remaining Occur.
Y
N
708
Define Failure Model
Remaining Failure
Y
N
710
Define Model

FIGURE 7

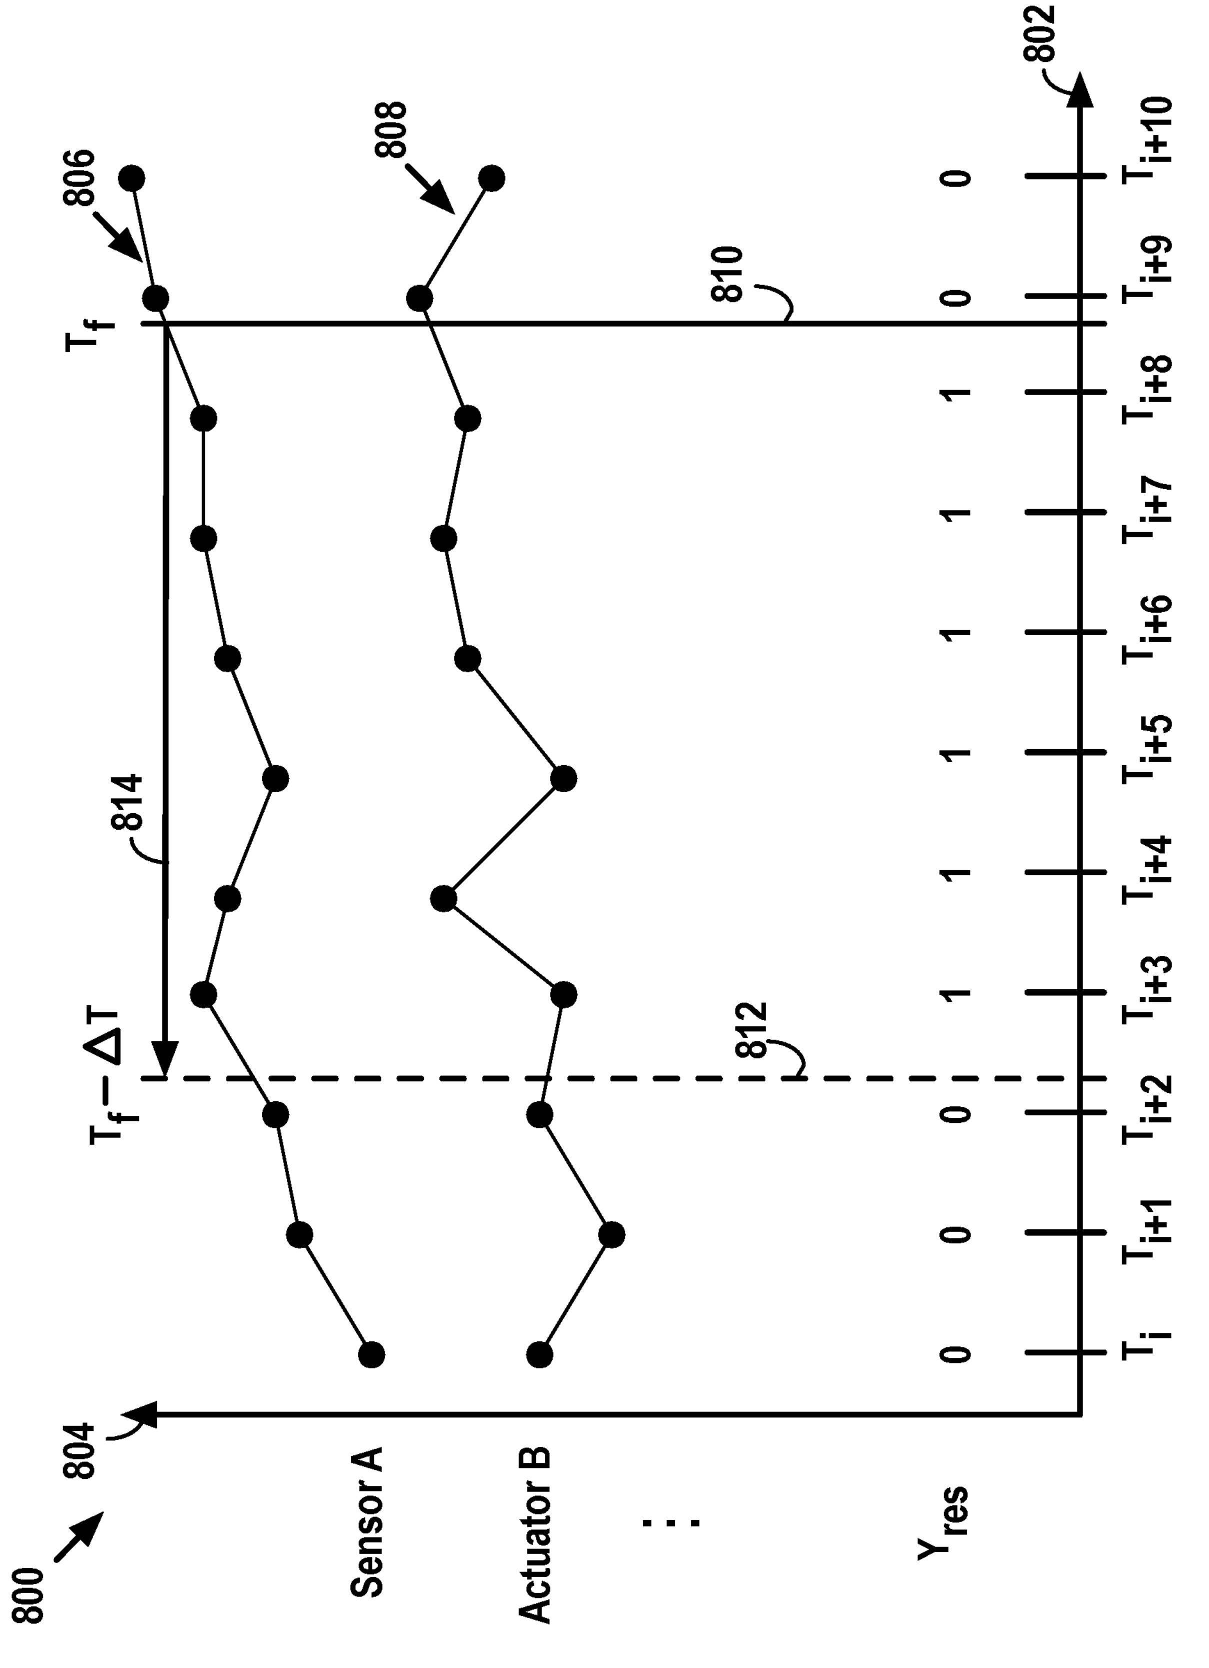
800
804
802
806
808
810
812
814
Sensor A
Actuator B
. . .
Y_{res}
T_f
$T_f - \Delta T$
0 0 0 1 1 1 1 1 1 0 0
T_i
T_{i+1}
T_{i+2}
T_{i+3}
T_{i+4}
T_{i+5}
T_{i+6}
T_{i+7}
T_{i+8}
T_{i+9}
T_{i+10}

FIGURE 8

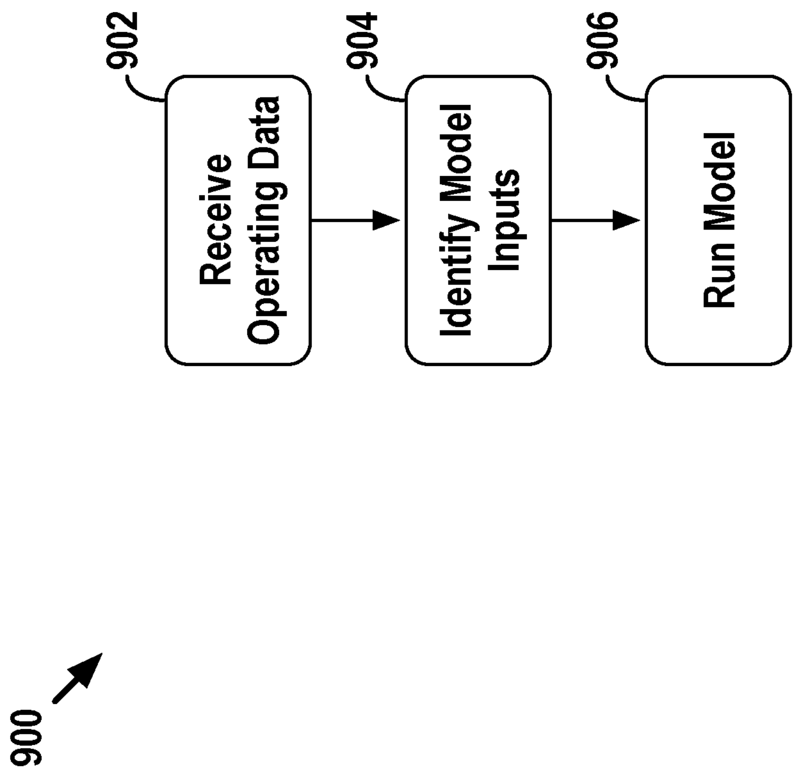
902
Receive Operating Data
904
Identify Model Inputs
906
Run Model
900

FIGURE 9

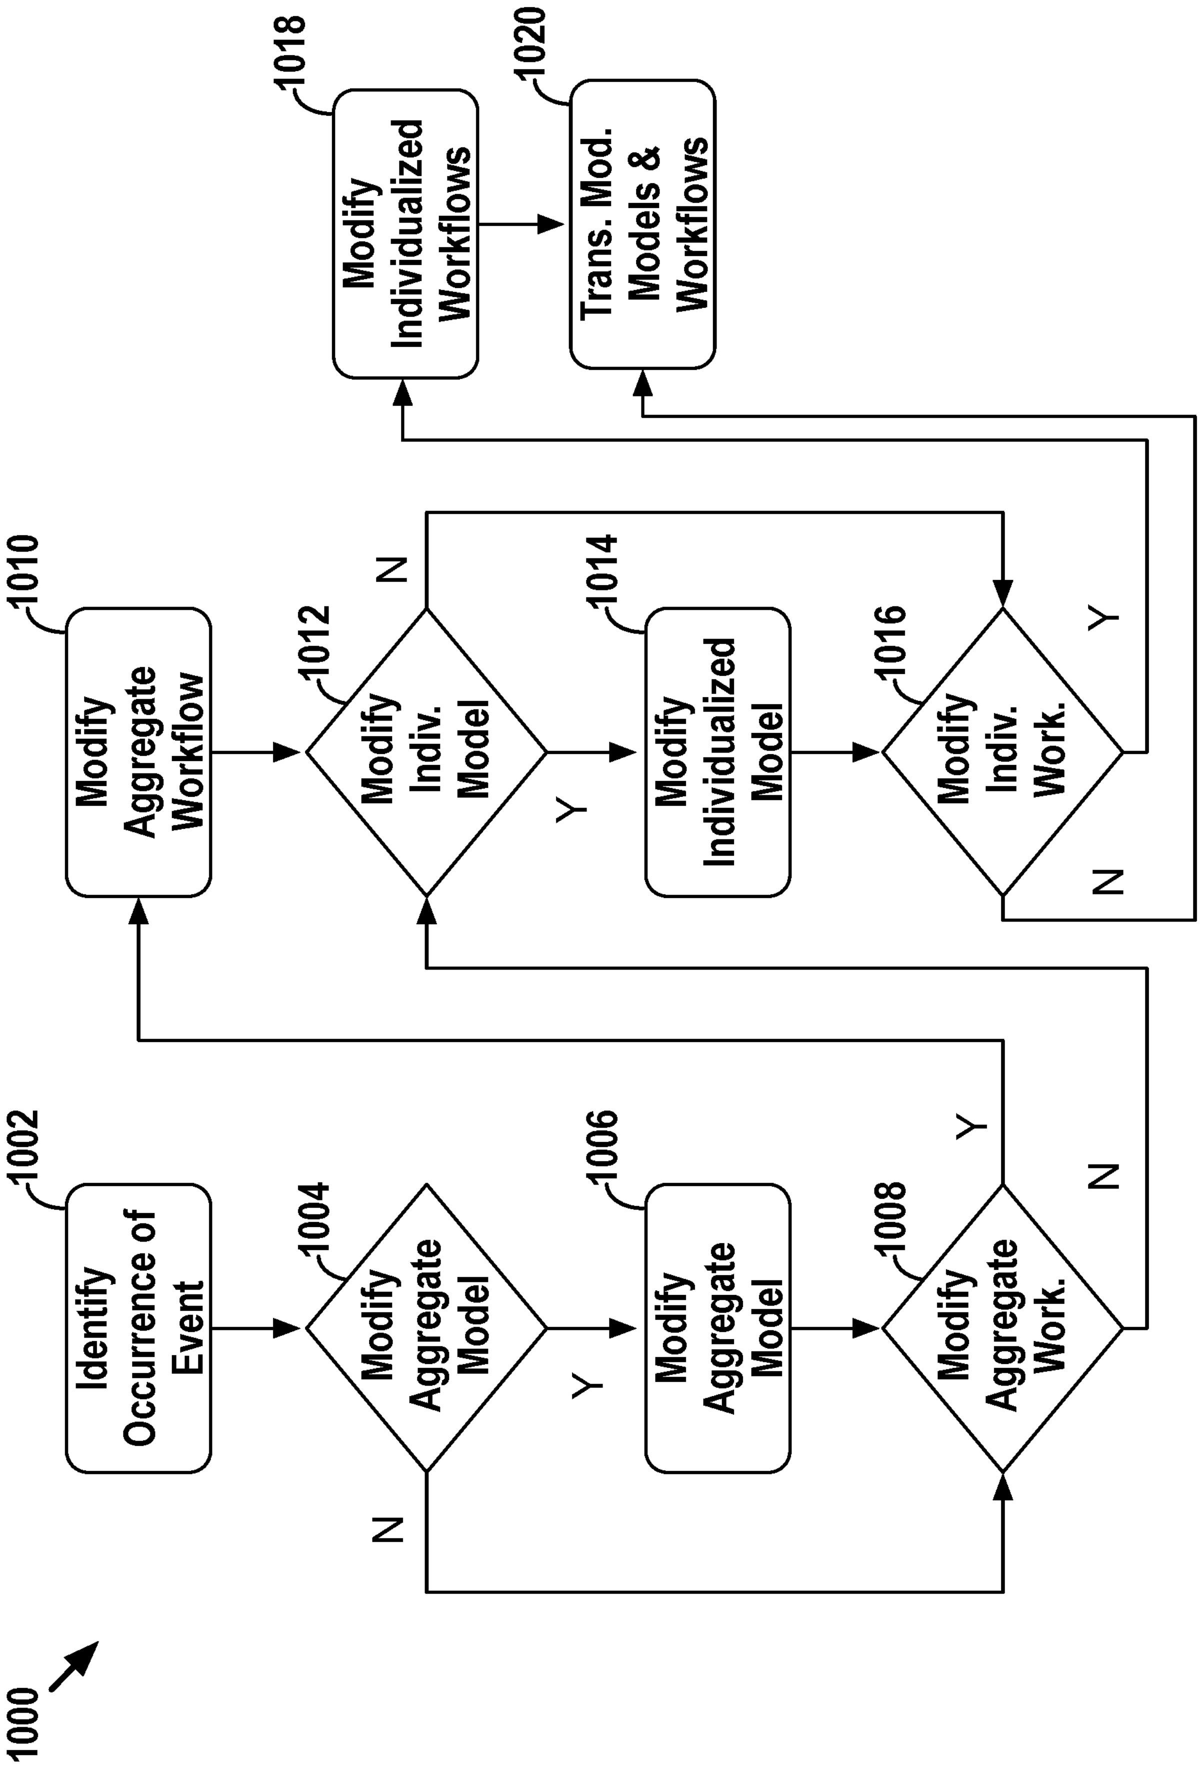
1000
1002
Identify Occurrence of Event
1004
Modify Aggregate Model
Y
N
1006
Modify Aggregate Model
1008
Modify Aggregate Work.
Y
N
1010
Modify Aggregate Workflow
1012
Modify Indiv. Model
Y
N
1014
Modify Individualized Model
1016
Modify Indiv. Work.
Y
N
1018
Modify Individualized Workflows
1020
Trans. Mod. Models & Workflows

FIGURE 10

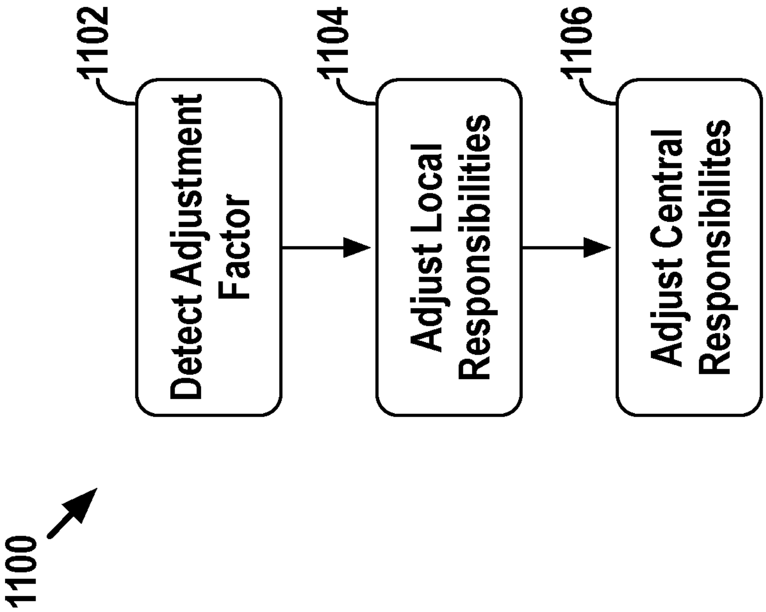
1100
1102
Detect Adjustment Factor
1104
Adjust Local Responsibilities
1106
Adjust Central Responsibilites

FIGURE 11

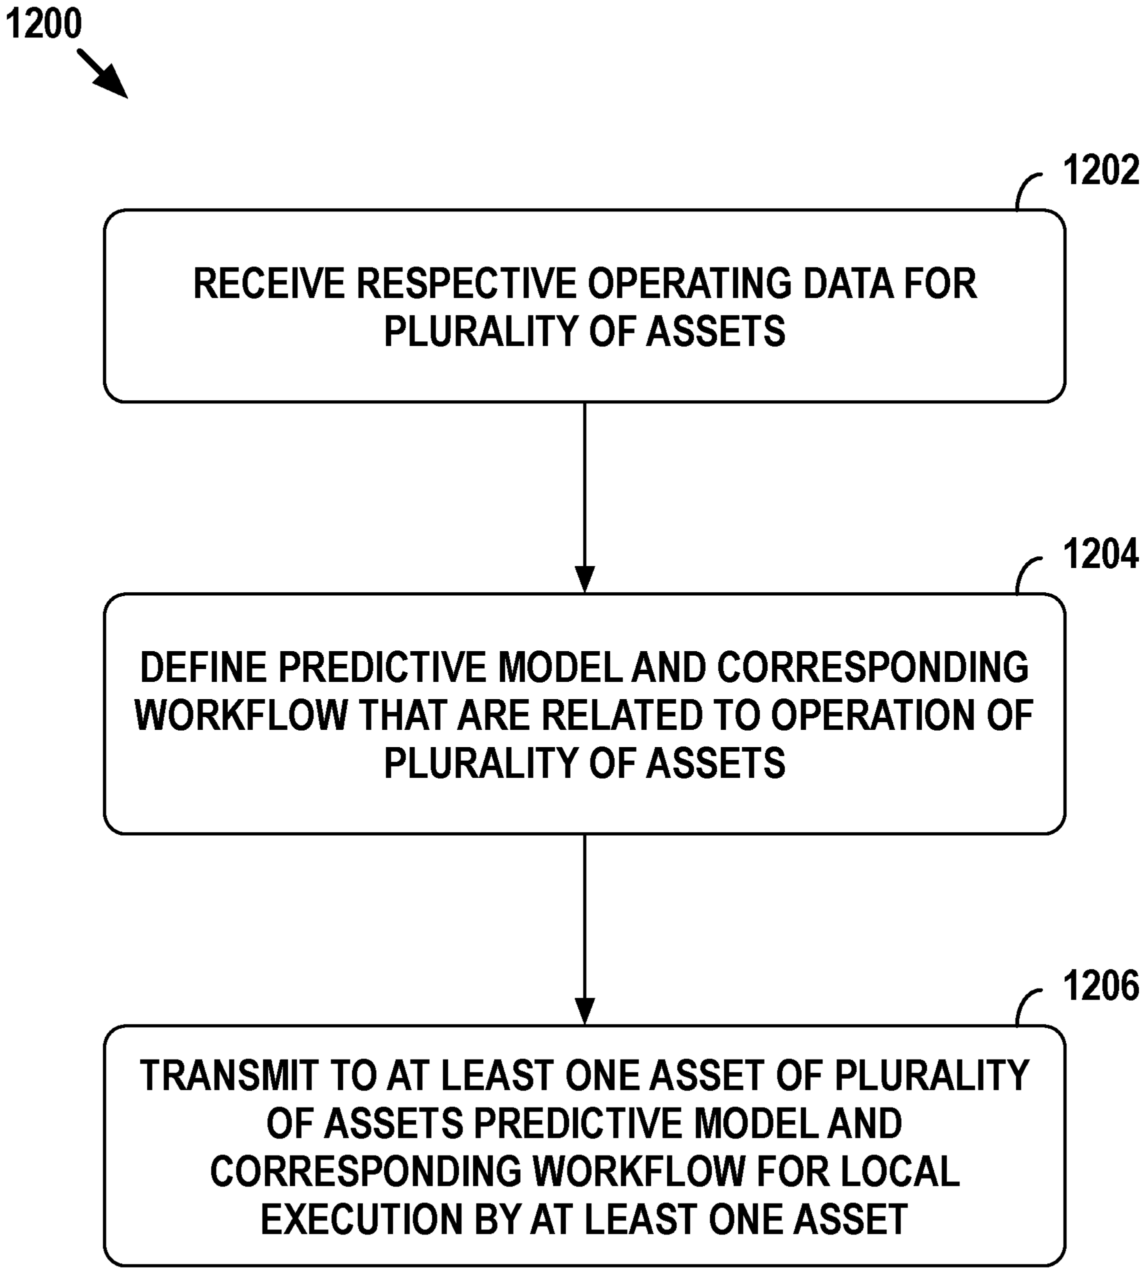
1200
1202
RECEIVE RESPECTIVE OPERATING DATA FOR PLURALITY OF ASSETS
1204
DEFINE PREDICTIVE MODEL AND CORRESPONDING WORKFLOW THAT ARE RELATED TO OPERATION OF PLURALITY OF ASSETS
1206
TRANSMIT TO AT LEAST ONE ASSET OF PLURALITY OF ASSETS PREDICTIVE MODEL AND CORRESPONDING WORKFLOW FOR LOCAL EXECUTION BY AT LEAST ONE ASSET

FIGURE 12

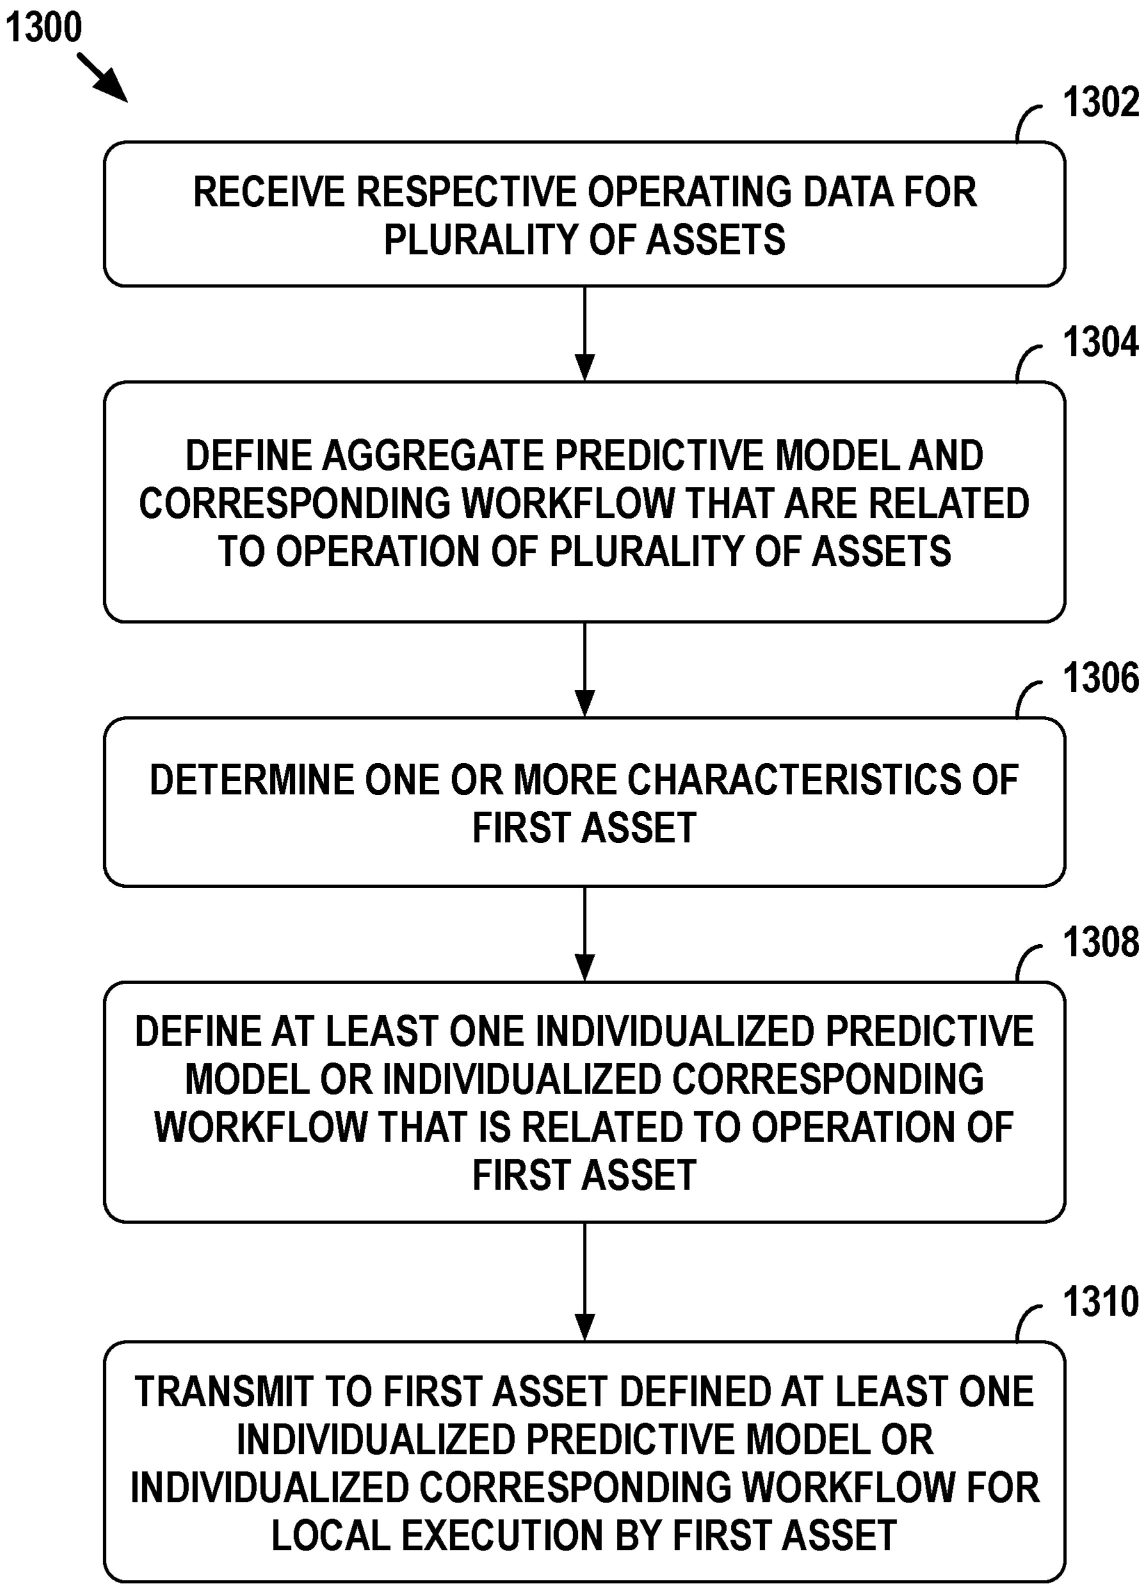
1300
1302
RECEIVE RESPECTIVE OPERATING DATA FOR PLURALITY OF ASSETS
1304
DEFINE AGGREGATE PREDICTIVE MODEL AND CORRESPONDING WORKFLOW THAT ARE RELATED TO OPERATION OF PLURALITY OF ASSETS
1306
DETERMINE ONE OR MORE CHARACTERISTICS OF FIRST ASSET
1308
DEFINE AT LEAST ONE INDIVIDUALIZED PREDICTIVE MODEL OR INDIVIDUALIZED CORRESPONDING WORKFLOW THAT IS RELATED TO OPERATION OF FIRST ASSET
1310
TRANSMIT TO FIRST ASSET DEFINED AT LEAST ONE INDIVIDUALIZED PREDICTIVE MODEL OR INDIVIDUALIZED CORRESPONDING WORKFLOW FOR LOCAL EXECUTION BY FIRST ASSET

FIGURE 13

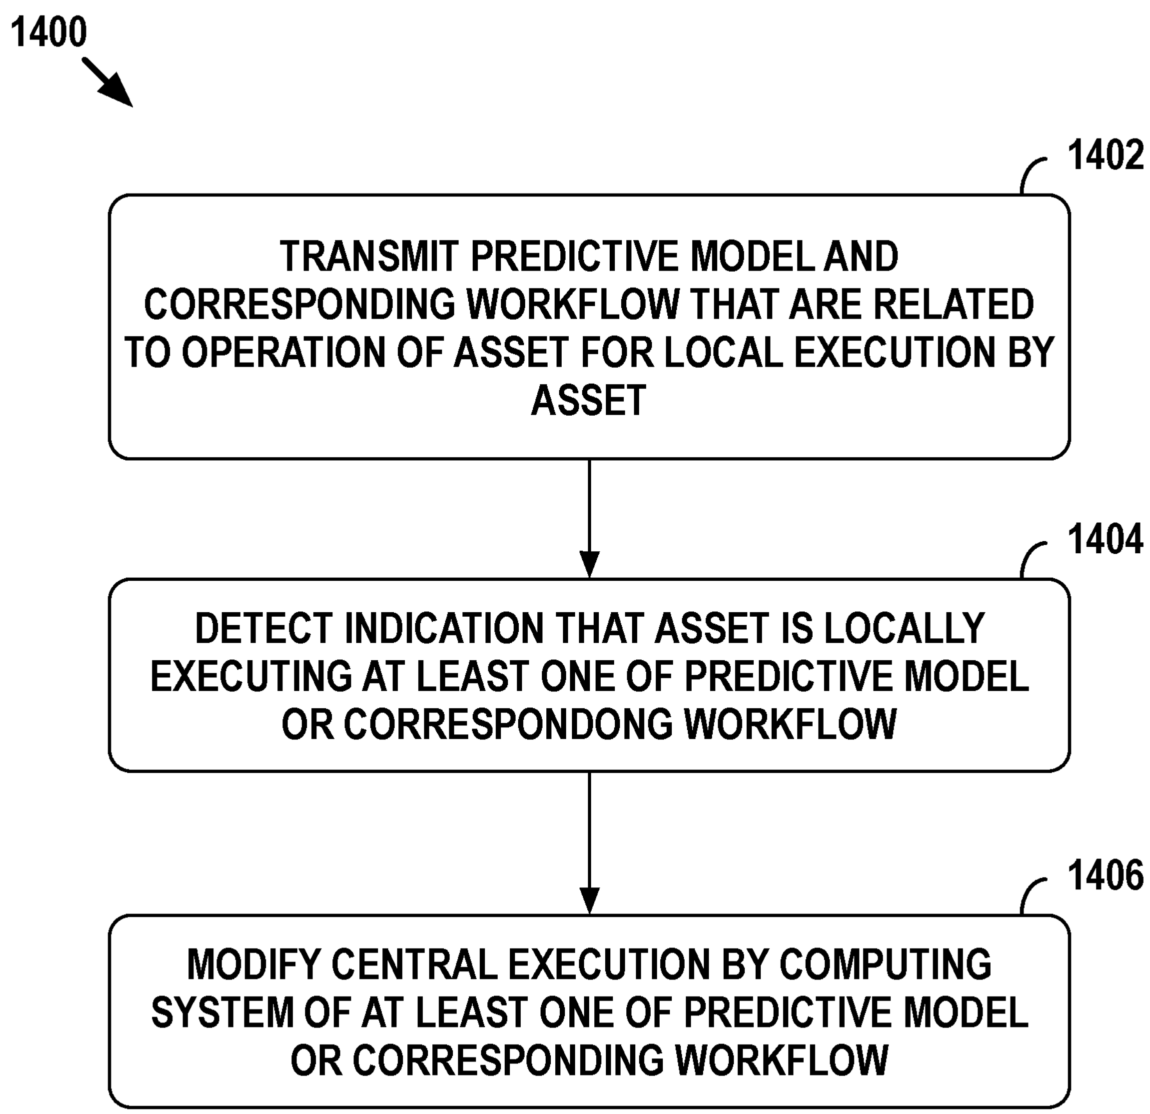
1400
1402
TRANSMIT PREDICTIVE MODEL AND CORRESPONDING WORKFLOW THAT ARE RELATED TO OPERATION OF ASSET FOR LOCAL EXECUTION BY ASSET
1404
DETECT INDICATION THAT ASSET IS LOCALLY EXECUTING AT LEAST ONE OF PREDICTIVE MODEL OR CORRESPONDONG WORKFLOW
1406
MODIFY CENTRAL EXECUTION BY COMPUTING SYSTEM OF AT LEAST ONE OF PREDICTIVE MODEL OR CORRESPONDING WORKFLOW

FIGURE 14

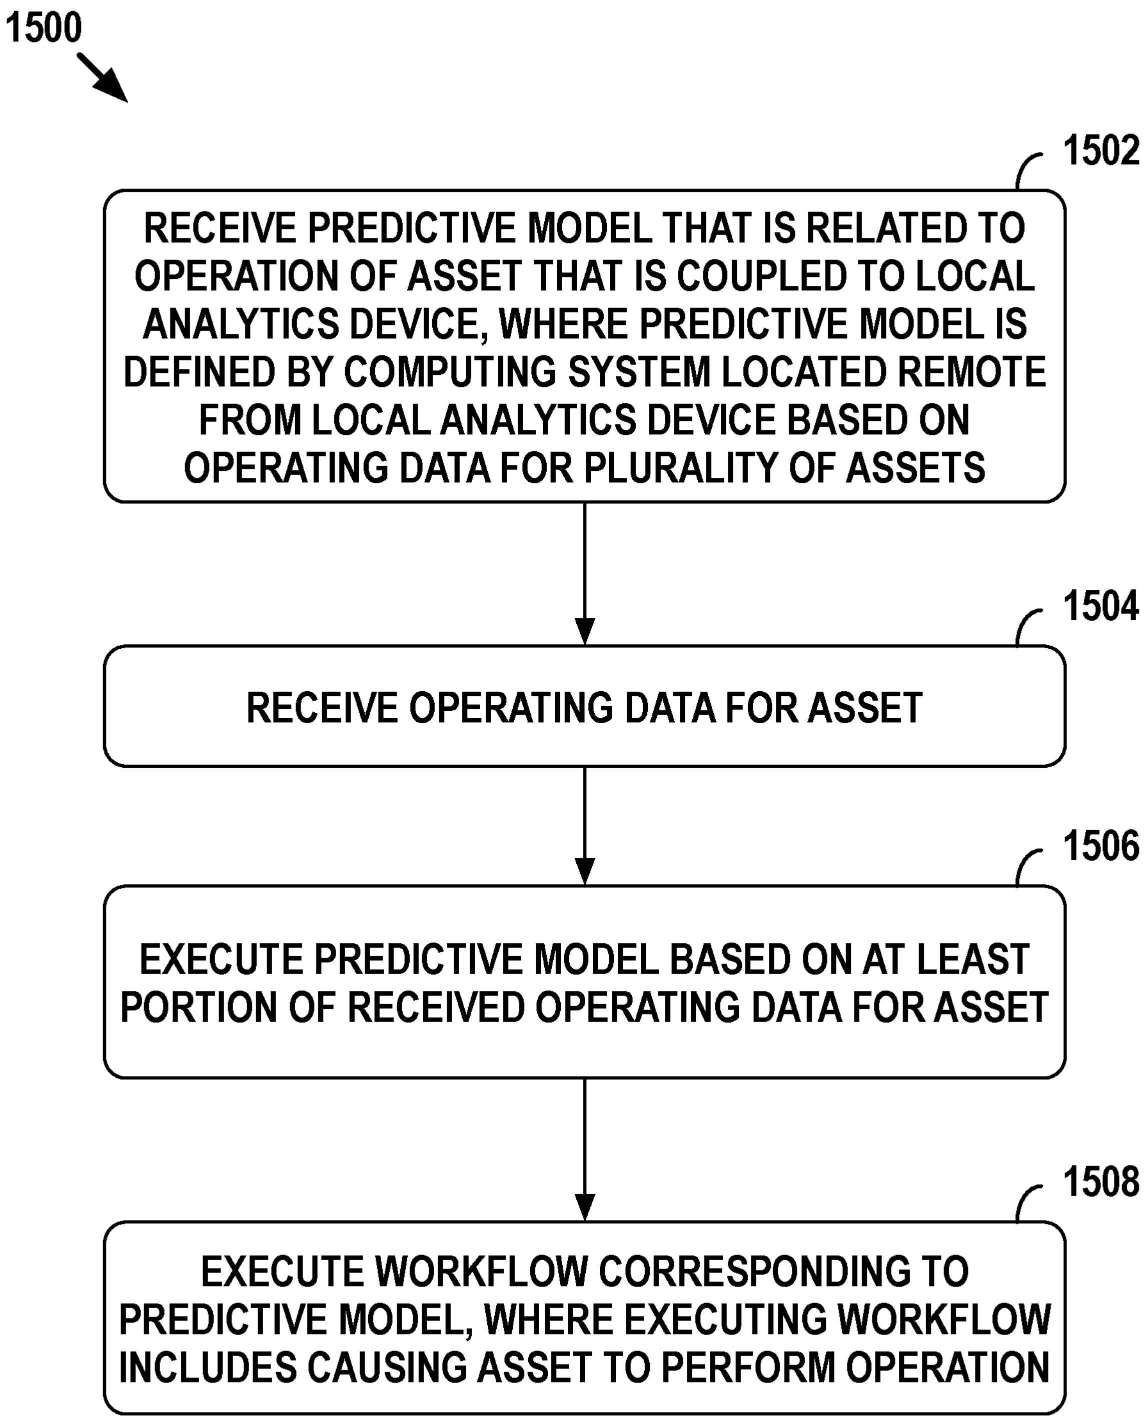
1500
1502
RECEIVE PREDICTIVE MODEL THAT IS RELATED TO OPERATION OF ASSET THAT IS COUPLED TO LOCAL ANALYTICS DEVICE, WHERE PREDICTIVE MODEL IS DEFINED BY COMPUTING SYSTEM LOCATED REMOTE FROM LOCAL ANALYTICS DEVICE BASED ON OPERATING DATA FOR PLURALITY OF ASSETS
1504
RECEIVE OPERATING DATA FOR ASSET
1506
EXECUTE PREDICTIVE MODEL BASED ON AT LEAST PORTION OF RECEIVED OPERATING DATA FOR ASSET
1508
EXECUTE WORKFLOW CORRESPONDING TO PREDICTIVE MODEL, WHERE EXECUTING WORKFLOW INCLUDES CAUSING ASSET TO PERFORM OPERATION

FIGURE 15

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2016/037247

A. CLASSIFICATION OF SUBJECT MATTER

G06Q 10/06(2012.01)I, G06Q 50/04(2012.01)I

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)
G06Q 10/06; G06F 11/07; G06F 19/00; H04N 1/00; E02F 9/20; E02F 9/24; G06F 11/22; G06F 11/00; G06Q 50/04

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
Korean utility models and applications for utility models
Japanese utility models and applications for utility models

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
eKOMPASS(KIPO internal) & Keywords: assets, operating, define, predictive, model, workflow, local, execution

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2011-0276828 A1 (KENJI TAMAKI et al.) 10 November 2011 See abstract, paragraphs [0033],[0042],[0089],[0093]-[0098],[0195],[0269], claims 1,16 and figures 1,26.	1-60
Y	JP 2009-206850 A (FUJI XEROX CO., LTD.) 10 September 2009 See abstract, paragraphs [0076],[0085],[0167]-[0170], claims 1,3 and figures 1,6,8.	1-60
A	US 2002-0013635 A1 (YASUHIRO GOTOU et al.) 31 January 2002 See abstract, claims 1-2,6-7 and figures 4-5,10.	1-60
A	US 7509537 B1 (DAVID W. JENSEN et al.) 24 March 2009 See abstract, claims 1-4 and figure 2.	1-60
A	EP 1403437 B1 (HITACHI CONSTRUCTION MACHINERY CO., LTD.) 11 December 2013 See claim 1 and figure 1.	1-60

☐ 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	Date of mailing of the international search report
19 September 2016 (19.09.2016)	**20 September 2016 (20.09.2016)**

Name and mailing address of the ISA/KR
International Application Division
Korean Intellectual Property Office
189 Cheongsa-ro, Seo-gu, Daejeon, 35208, Republic of Korea
Facsimile No. +82-42-481-8578

Authorized officer
KANG, Min Jeong
Telephone No. +82-42-481-8131

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2011-0276828 A1	10/11/2011	JP 05108116 B2	26/12/2012
		US 8515719 B2	20/08/2013
		WO 2010-082322 A1	22/07/2010
JP 2009-206850 A	10/09/2009	None	
US 2002-0013635 A1	31/01/2002	EP 1164550 A2	19/12/2001
		EP 1164550 A3	16/06/2004
		EP 1164550 B1	03/12/2008
		JP 03827928 B2	27/09/2006
		JP 04236022 B2	11/03/2009
		JP 2001-356808 A	26/12/2001
		JP 2002-071519 A	08/03/2002
		US 6983207 B2	03/01/2006
US 7509537 B1	24/03/2009	None	
EP 1403437 B1	11/12/2013	CN 1462328 A	17/12/2003
		CN 1462328 C	12/04/2006
		CN 1746438 A	15/03/2006
		CN 1746438 B	09/05/2012
		EP 1403437 A1	31/03/2004
		EP 1403437 A4	06/05/2009
		JP 04711321 B2	29/06/2011
		JP 04897152 B2	14/03/2012
		JP 2002-332664 A	22/11/2002
		JP 2003-030355 A	31/01/2003
		KR 10-0523228 B1	20/10/2005
		KR 10-2003-0014325 A	15/02/2003
		US 2004-0186687 A1	23/09/2004
		US 2006-0031042 A1	09/02/2006
		US 7079982 B2	18/07/2006
		US 7222051 B2	22/05/2007
		WO 02-090669 A1	14/11/2002

(19)中华人民共和国国家知识产权局

(12)发明专利申请

(10)申请公布号 CN 107851233 A
(43)申请公布日 2018.03.27

(21)申请号 201680043854.5

(22)申请日 2016.06.13

(30)优先权数据
14/744,352 2015.06.19 US
14/744,369 2015.06.19 US
14/963,207 2015.12.08 US

(85)PCT国际申请进入国家阶段日
2018.01.26

(86)PCT国际申请的申请数据
PCT/US2016/037247 2016.06.13

(87)PCT国际申请的公布数据
WO2016/205132 EN 2016.12.22

(71)申请人 阿普泰克科技公司
地址 美国伊利诺伊州

(72)发明人 B·尼古拉斯 J·科尔布

(74)专利代理机构 北京律盟知识产权代理有限责任公司 11287
代理人 李琳

(51)Int.Cl.
G06Q 10/06(2006.01)
G06Q 50/04(2006.01)

权利要求书7页 说明书35页 附图18页

(54)发明名称

资产处的本地分析

(57)摘要

本文中揭示了与资产相关的系统、装置及方法及与资产操作相关的预测模型及对应工作流。特定来说，实例涉及：定义及部署聚合预测模型及对应工作流；定义及部署个性化的预测模型及/或对应工作流；及动态地调整模型-工作流程对的执行。另外，实例涉及经配置以接收及本地执行预测模型、本地个性化预测模型及/或本地执行工作流或其部分的资产。

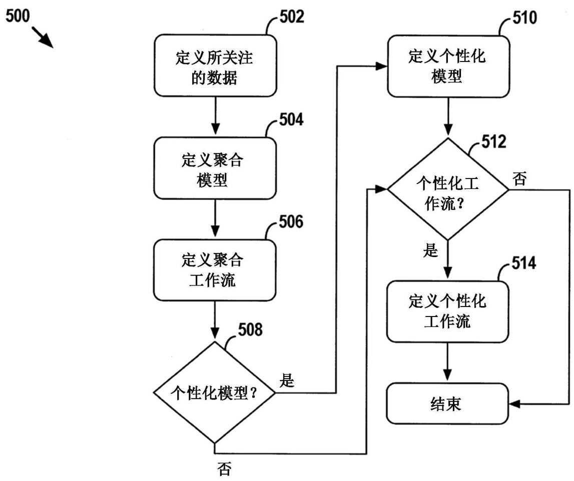

1.一种计算系统,其包括:

至少一个处理器;

非暂时性计算机可读媒体;及

存储在所述非暂时性计算机可读媒体上的程序指令,所述程序指令可由所述至少一个处理器执行以致使所述计算系统:

接收多个资产的相应操作数据;

基于所述所接收的操作数据,定义与所述多个资产的操作相关的预测模型及对应工作流;及

向所述多个资产中的至少一个资产发射所述预测模型及所述对应工作流以供所述至少一个资产本地执行。

2.根据权利要求1所述的计算系统,其中所述相应的操作数据包括(i)与在特定时间在给定资产处发生的故障相关联的异常状况数据,及(ii)指示所述给定资产在所述特定时间的至少一种操作状况的传感器或致动器数据中的至少一者。

3.根据权利要求1所述的计算系统,其中所述预测模型经定义以输出在未来时间段内将在给定资产处发生特定事件的概率。

4.根据权利要求3所述的计算系统,其中所述对应工作流包括基于所确定的概率来执行的一或多个操作。

5.根据权利要求1所述的计算系统,其中所述对应工作流包括给定资产控制所述给定资产的一或多个致动器以促进修改所述给定资产的操作状况。

6.根据权利要求1所述的计算系统,其中所述对应工作流包括要由给定资产本地执行的一或多个诊断工具。

7.根据权利要求1所述的计算系统,其中所述对应工作流包括根据数据采集方案来采集传感器数据。

8.根据权利要求7所述的计算系统,其中所述数据采集方案指示从其中采集数据的给定资产的一或多个传感器。

9.根据权利要求8所述的计算系统,其中所述数据采集方案进一步指示所述给定资产将从所述一或多个传感器中的每一者采集的数据量。

10.根据权利要求1所述的计算系统,其中所述对应工作流包括根据数据发射方案向所述计算系统发射数据。

11.根据权利要求10所述的计算系统,其中所述数据发射方案指示给定资产向所述计算系统发射数据的频率。

12.根据权利要求1所述的计算系统,其中所述计算系统是第一计算系统,且其中所述对应工作流包括给定资产向第二计算系统发射指令以促进致使所述第二计算系统实行与所述给定资产相关的操作。

13.根据权利要求1所述的计算系统,其中所述多个资产中的所述至少一个资产包括第一资产及第二资产,且其中发射所述预测模型及所述对应工作流包括向所述第一资产及所述第二资产发射所述预测模型及所述对应工作流。

14.一种上面存储有指令的非暂时性计算机可读媒体,所述指令可执行以致使计算系统执行以下操作:

接收多个资产的相应操作数据；

基于所述所接收的操作数据，定义与所述多个资产的操作相关的预测模型及对应工作流；及

向所述多个资产中的至少一个资产发射所述预测模型及对应工作流以供所述至少一个资产本地执行。

15.根据权利要求14所述的非暂时性计算机可读媒体，其中所述预测模型经定义以输出在未来时间段内将在给定资产处发生特定事件的概率。

16.根据权利要求14所述的非暂时性计算机可读媒体，其中所述对应工作流包括给定资产控制所述给定资产的一或多个致动器以促进修改所述给定资产的操作状况。

17.根据权利要求14所述的非暂时性计算机可读媒体，其中所述对应工作流包括要由给定资产本地执行的一或多个诊断工具。

18.根据权利要求14所述的非暂时性计算机可读媒体，其中所述计算系统是第一计算系统，且其中所述对应工作流包括给定资产向第二计算系统发射指令以促进致使所述第二计算系统执行与所述给定资产相关的操作。

19.一种计算机实施方法，其包括：

接收多个资产的相应操作数据；

基于所述所接收的操作数据，定义与所述多个资产的操作相关的预测模型及对应工作流；及

向所述多个资产中的至少一个资产发射所述预测模型及对应工作流以供所述至少一个资产本地执行。

20.根据权利要求19所述的计算机实施方法，其中所述对应工作流包括根据数据采集方案采集传感器数据，其中所述数据采集方案指示从其中采集数据的给定资产的一或多个传感器。

21.一种计算系统，其包括：

至少一个处理器；

非暂时性计算机可读媒体；及

存储在所述非暂时性计算机可读媒体上的程序指令，所述程序指令可由所述至少一个处理器执行以使所述计算系统执行以下操作：

接收多个资产的操作数据，其中所述多个资产包括第一资产；

基于所述所接收的操作数据，定义与所述多个资产的操作相关的聚合预测模型及聚合对应工作流；

确定所述第一资产的一或多个特性；

基于所述第一资产的所述一或多个特性以及所述聚合预测模型及所述聚合对应工作流，定义与所述第一资产的所述操作相关的个性化预测模型或个性化对应工作流中的至少一者；及

向所述第一资产发射所述所定义的至少一个个性化预测模型或个性化对应工作流以供所述第一资产本地执行。

22.根据权利要求21所述的计算系统，其中所述第一资产的所述一或多个特性包括资产年限或资产健康状况中的至少一者。

23.根据权利要求21所述的计算系统,其中确定所述第一资产的所述一或多个特性包括基于所述第一资产的所接收的操作数据来确定所述第一资产的所述一或多个特性。

24.根据权利要求21所述的计算系统,其中定义个性化预测模型或个性化对应工作流中的至少一者包括定义所述个性化预测模型及所述个性化对应工作流,且其中发射所述至少一个个性化预测模型或个性化对应工作流包括发射所述个性化预测模型及所述个性化对应工作流。

25.根据权利要求21所述的计算系统,其中定义个性化预测模型或个性化对应工作流中的至少一者包括定义所述个性化对应工作流,且其中发射所述至少一个个性化预测模型或个性化对应工作流包括发射所述聚合预测模型及所述个性化对应工作流。

26.根据权利要求25所述的计算系统,其中所述聚合对应工作流包括第一操作,且其中所述个性化对应工作流包括与所述第一操作不同的第二操作。

27.根据权利要求26所述的计算系统,其中所述第一操作包括根据第一采集方案采集数据,且其中所述第二操作包括根据第二采集方案采集数据。

28.根据权利要求26所述的计算系统,其中所述第一操作包括根据采集方案采集数据,且其中所述第二操作包括执行一或多个诊断工具。

29.根据权利要求21所述的计算系统,其中所述多个资产进一步包括第二资产,且其中所述程序指令进一步包括可执行以致使所述计算系统执行以下操作的指令:

在发射所述至少一个个性化预测模型或个性化对应工作流之后,接收指示所述第二资产处发生事件的所述第二资产的操作数据;

基于第二资产的所接收的操作数据,修改所述至少一个个性化预测模型或个性化对应工作流;及

向所述第一资产发射所述所修改的至少一个个性化预测模型或个性化对应工作流。

30.一种上面存储有指令的非暂时性计算机可读媒体,所述指令可执行以致使计算系统执行以下操作:

接收多个资产的操作数据,其中所述多个资产包括第一资产;

基于所述所接收的操作数据,定义与所述多个资产的操作相关的聚合预测模型及聚合对应工作流;

确定所述第一资产的一或多个特性;

基于所述第一资产的所述一或多个特性以及所述聚合预测模型及所述聚合工作流,定义与所述第一资产的所述操作相关的个性化预测模型或个性化对应工作流中的至少一者;及

向所述第一资产发射所述所定义的至少一个个性化预测模型或个性化对应工作流以供所述第一资产本地执行。

31.根据权利要求30所述的非暂时性计算机可读媒体,其中定义个性化预测模型或个性化对应工作流中的至少一者包括定义所述个性化预测模型及所述个性化对应工作流,且其中发射所述至少一个个性化预测模型或个性化对应工作流包括发射所述个性化预测模型及所述个性化对应工作流。

32.根据权利要求30所述的非暂时性计算机可读媒体,其中定义个性化预测模型或个性化对应工作流中的至少一者包括定义所述个性化对应工作流,且其中发射所述至少一个

个性化预测模型或个性化对应工作流包括发射所述聚合预测模型及所述个性化对应工作流。

33.根据权利要求32所述的非暂时性计算机可读媒体,其中所述聚合对应工作流包括第一操作,且其中所述个性化对应工作流包括与所述第一操作不同的第二操作。

34.根据权利要求33所述的非暂时性计算机可读媒体,其中所述第一操作包括根据第一采集方案采集数据,且其中所述第二操作包括根据第二采集方案采集数据。

35.根据权利要求33所述的非暂时性计算机可读媒体,其中所述第一操作包括根据采集方案采集数据,且其中所述第二操作包括执行一或多个诊断工具。

36.根据权利要求30所述的非暂时性计算机可读媒体,其中所述多个资产进一步包括第二资产,且其中所述程序指令进一步包括可执行以致使所述计算系统执行以下操作的指令:

在发射所述至少一个个性化预测模型或个性化对应工作流之后,接收指示在所述第二资产处发生事件的所述第二资产的操作数据;

基于第二资产的所接收的操作数据,修改所述至少一个个性化预测模型或个性化对应工作流;及

向所述第一资产发射所修改的至少一个个性化预测模型或个性化对应工作流。

37.一种计算机实施方法,其包括:

接收多个资产的操作数据,其中所述多个资产包括第一资产;

基于所述所接收的操作数据,定义与所述多个资产的操作相关的聚合预测模型及聚合对应工作流;

确定所述第一资产的一或多个特性;

基于所述第一资产的所述一或多个特性以及所述聚合预测模型及所述聚合对应工作流,定义与所述第一资产的所述操作相关的个性化预测模型或个性化对应工作流中的至少一者;及

向所述第一资产发射所述所定义的至少一个个性化预测模型或个性化对应工作流以供所述第一资产本地执行。

38.根据权利要求37所述的计算机实施方法,其中定义个性化预测模型或个性化对应工作流中的至少一者包括定义所述个性化对应工作流,且其中发射所述至少一个个性化预测模型或个性化对应工作流包括发射所述聚合预测模型及所述个性化对应工作流。

39.根据权利要求38所述的计算机实施方法,其中所述聚合对应工作流包括第一操作,且其中所述个性化对应工作流包括与所述第一操作不同的第二操作。

40.根据权利要求39所述的计算机实施方法,其中所述第一操作或所述第二操作中的一者包括执行一或多个诊断工具。

41.一种计算装置,其包括:

资产接口,其经配置以将所述计算装置耦合到资产;

网络接口,其经配置以促进所述计算装置与远离所述计算装置定位的计算系统之间的通信;

至少一个处理器;

非暂时性计算机可读媒体;及

存储在所述非暂时性计算机可读媒体上的程序指令，所述程序指令可由所述至少一个处理器执行以致使所述计算装置执行以下操作：

经由所述网络接口接收与所述资产的操作相关的预测模型，其中所述预测模型由所述计算系统基于多个资产的操作数据来定义；

经由所述资产接口接收所述资产的操作数据；

基于所述资产的所接收的操作数据的至少部分执行所述预测模型；及

基于执行所述预测模型，执行对应于所述预测模型的工作流，其中执行所述工作流包括经由所述资产接口致使所述资产执行操作。

42.根据权利要求41所述的计算装置，其中所述资产接口将所述计算装置通信地耦合到所述资产的资产上计算机。

43.根据权利要求41所述的计算装置，其中所述资产包括致动器，且其中执行所述工作流包括致使所述致动器执行机械操作。

44.根据权利要求41所述的计算装置，其中执行所述工作流包括致使所述资产执行诊断工具。

45.根据权利要求41所述的计算装置，其中执行所述工作流进一步包含经由所述网络接口致使执行远离所述资产的操作。

46.根据权利要求45所述的计算装置，其中致使执行远离所述资产的操作包括指示所述计算系统执行远离所述资产的操作。

47.根据权利要求41所述的计算装置，其中存储在所述非暂时性计算机可读媒体上的所述程序指令进一步可由所述至少一个处理器执行以致使所述计算装置：

在执行所述预测模型之前，个性化所述预测模型。

48.根据权利要求47所述的计算装置，其中个性化所述预测模型包括至少基于所述资产的所接收的操作数据来修改所述预测模型的一或多个参数。

49.根据权利要求47所述的计算装置，其中存储在所述非暂时性计算机可读媒体上的所述程序指令进一步可由所述至少一个处理器执行以致使所述计算装置执行以下操作：

在个性化所述预测模型之后，经由所述网络接口向所述计算系统发射所述预测模型已被个性化的指示。

50.根据权利要求41所述的计算装置，其中所述预测模型是第一预测模型，且其中存储在所述非暂时性计算机可读媒体上的所述程序指令进一步可由所述至少一个处理器执行以致使所述计算装置执行以下操作：

在执行所述第一预测模型之前，经由所述网络接口向所述计算系统发射所述资产的所述所接收的操作数据的给定子集，其中所接收的操作数据的所述给定子集包括由一或多个传感器的给定群组产生的操作数据。

51.根据权利要求50所述的计算装置，其中存储在所述非暂时性计算机可读媒体上的所述程序指令进一步可由所述至少一个处理器执行以致使所述计算装置执行以下操作：

在发射所述资产的所述所接收的操作数据的所述给定子集之后，接收与所述资产的所述操作相关的第二预测模型，其中所述第二预测模型由所述计算系统基于所述资产的所述接收的操作数据的所述给定子集资产定义；及

执行所述第二预测模型而不是所述第一预测模型。

52.一种上面存储有指令的非暂时性计算机可读媒体，所述指令可执行以致使经由计算装置的资产接口耦合到资产的所述计算装置执行以下操作：

经由所述计算装置的网络接口接收与所述资产的操作相关的预测模型，所述计算装置的所述网络接口经配置以促进所述计算装置与远离所述计算装置定位的计算系统之间的通信，其中所述预测模型由所述计算系统基于多个资产的操作数据来定义；

经由所述资产接口接收所述资产的操作数据；

基于所述资产的所接收的操作数据的至少部分执行所述预测模型；及

基于执行所述预测模型，执行对应于所述预测模型的工作流，其中执行所述工作流包括经由所述资产接口致使所述资产执行操作。

53.根据权利要求52所述的非暂时性计算机可读媒体，其中存储在所述非暂时性计算机可读媒体上的所述程序指令进一步可执行以致使所述计算装置：

在执行所述预测模型之前，个性化所述预测模型。

54.根据权利要求53所述的非暂时性计算机可读媒体，其中个性化所述预测模型包括至少基于所述资产的所接收的操作数据来修改所述预测模型的一或多个参数。

55.根据权利要求52所述的非暂时性计算机可读媒体，其中所述预测模型是第一预测模型，且其中存储在所述非暂时性计算机可读媒体上的所述程序指令进一步可执行以致使所述计算装置：

在执行所述第一预测模型之前，经由所述网络接口向所述计算系统发射所述资产的所述所接收的操作数据的给定子集，其中所接收的操作数据的所述给定子集包括由一或多个传感器的给定群组产生的操作数据。

56.根据权利要求55所述的非暂时性计算机可读媒体，其中存储在所述非暂时性计算机可读媒体上的所述程序指令进一步可执行以使所述计算装置执行以下操作：

在从所述一或多个传感器的特定群组发射所述操作数据之后，接收与所述资产的所述操作相关的第二预测模型，其中所述第二预测模型由所述计算系统基于所述资产的所述接收的操作数据的所述给定子集资产定义；及

执行所述第二预测模型而不是所述第一模型。

57.一种计算机实施方法，所述方法包括：

经由计算装置的网络接口接收与资产的操作相关的预测模型，所述计算装置的所述网络接口经由所述计算装置的资产接口耦合到所述资产，其中所述预测模型由远离所述计算装置定位的计算系统基于多个资产的操作数据来定义；

由所述计算装置经由所述资产接口接收所述资产的操作数据；

由所述计算装置基于所述资产的所接收的操作数据的至少部分执行所述预测模型；及

基于执行所述预测模型，由所述计算装置执行对应于所述预测模型的工作流，其中执行所述工作流包括经由所述资产接口致使所述资产执行操作。

58.根据权利要求57所述的计算机实施方法，所述方法进一步包括：

在执行所述预测模型之前，由所述计算装置个性化所述预测模型。

59.根据权利要求58所述的计算机实施方法，其中个性化所述预测模型包括至少基于所述资产的所接收的操作数据来修改所述预测模型的一或多个参数。

60.根据权利要求57所述的计算机实施方法，其中执行所述工作流进一步包含经由所

述网络接口致使执行远离所述资产的操作。

资产处的本地分析

[0001] 相关申请案的交叉参考

[0002] 本申请案主张以下项的优先权:(i)2015年6月19日申请且标题为用于本地执行的聚合预测模型及工作流(Aggregate Predictive Model&Workflow for Local Execution)的第14/744,352号美国非临时专利案;(ii)2015年6月19日申请且标题为用于资产的个性化预测模型及工作流(Individualized Predictive Model&Workflow for an Asset)的第14/744,369号美国非临时专利申请案;及(iii)2015年12月8日申请且标题为“资产处的本地分析(Local Analytics at an Asset)”的第14/963,207号美国非临时专利申请案,所述非临时专利申请案中的每一者的全部内容以引入方式全部并入本文中。本申请案还以引用方式并入了2015年6月5日申请且标题为“资产健康状况分数(Asset Health Score)”的第14/732,258号美国非临时专利申请案的全部内容。

背景技术

[0003] 今天,机器(在本文中也被称为“资产”)在许多行业中无处不在。从把货物运送到各国的机车到帮助护士及医生挽救生命的医疗设备,资产在日常生活中起着重要的作用。取决于资产所起的作用,其复杂性及成本可能会有所不同。例如,一些资产可包括多个子系统(例如,机车的发动机、传动装置等),其必须协调操作以使资产适当地运作。

[0004] 由于资产在日常生活中扮演着重要的角色,所以期望资产可在有限的停机时间内进行维修。因此,一些人已经开发了各机制来监测及检测资产内的异常状况,以促进可能在最小停机时间内维修资产。

发明内容

[0005] 用于监测资产的当前方法通常涉及资产上计算机,其从分布在整个资产中的各种传感器及/或致动器接收信号,所述传感器及/或致动器监测资产的操作状况。作为一个代表性实例,如果资产是机车,那么传感器及/或致动器可监测例如温度、电压及速度的参数以及其它参数。如果来自这些装置中的一或多者的传感器及/或致动器信号达到某些值,那么所述资产上计算机可产生异常状况指示符,例如“故障代码”,其是资产内已经发生异常状况的指示。

[0006] 通常,异常情况可为资产或其组件处的缺陷,其可能导致资产及/或组件的故障。因而,异常情况可与给定故障或可能多个故障相关联,因为异常情况是给定故障或多个故障的征兆。实际上,用户通常定义与每一异常状况指示符相关联的传感器及相应的传感器值。即,用户定义资产的“正常”操作状况(例如,不触发故障代码的操作状况)及“异常”操作状况(例如,触发故障代码的操作状况)。

[0007] 在资产上计算机产生异常状况指示符之后,指示符及/或传感器信号可被传递到远程位置,在所述远程位置中,用户可接收异常状况及/或传感器信号的一些指示且决定是否采取行动。用户可能采取的一个行动是指派机械师等来评估及可能地维修资产。一旦处于资产处,机械师可将计算装置连接到资产并操作计算装置以致使资产利用一或多个本地

诊断工具来促进诊断所产生的指示符的原因。

[0008] 虽然目前的资产监测系统在触发异常情况指示符方面通常是有效的，但是此类系统通常是保守的。即，到资产监测系统触发指示符的时候，资产内的故障可能已经发生(或即将发生)，这可能导致代价高昂的停机时间以及其它缺点。另外，由于此类资产监测系统中的资产上异常检测机制的简单本质，当前的资产监测方法倾向于涉及远程计算系统对资产执行监测计算，且如果检测到问题，那么向所述资产发射指令。当资产移动到通信网络的覆盖范围之外时，由于网络延迟及/或不可行，这可能是不利的。另外，由于存储在资产上的本地诊断工具的本质，当前的诊断程序倾向于又低效又麻烦，这是因为需要机械师来致使资产利用此类工具。

[0009] 本文中揭示的实例系统、装置及方法试图帮助解决这些问题中的一或多者。在实例实施方案中，网络配置可包含通信网络，其促进资产与远程计算系统之间的通信。在一些情况下，通信网络可促进资产与远程计算系统之间的安全通信(例如，经由加密或其它安全措施)。

[0010] 如上文所述，每一资产可包含分布在整个资产中的多个传感器及/或致动器，其促进监测资产的操作状况。多个资产可向远程计算系统提供指示每一资产的操作状况的相应数据，所述远程计算系统可经配置以基于所提供的数据执行一或多个操作。通常，传感器及/或致动器数据可用于一般的资产监测操作。然而，如本文中所述，远程计算系统及/或资产可利用此数据来促进执行更复杂的操作。

[0011] 在实例实施方案中，远程计算系统可经配置以定义及部署与资产的操作相关的预测模型及对应工作流(在本文中被称为“模型-工作流对”)。资产可经配置以接收模型-工作流对，并利用本地分析装置根据模型-工作流对来操作。

[0012] 通常，模型-工作流对可致使资产监测某些操作状况，且当存在某些状况时，修改可能帮助防止发生特定事件的行为。具体来说，预测模型可接收来自资产传感器及/或致动器的特定集合的数据作为输入，且输出一或多个特定事件可能在未来特定时间段内在资产处发生的可能性。工作流可涉及基于由模型输出的一或多个特定事件的可能性执行的一或多个操作。

[0013] 实际上，远程计算系统可定义聚合、预测模型及对应工作流、个性化预测模型及对应工作流，或其某个组合。“聚合”模型/工作流可指代对资产群组通用的模型/工作流，而“个性化”模型/工作流可指代针对来自资产群组的单个资产或资产子群组定制的模型/工作流。

[0014] 在实例实施方案中，远程计算系统可通过基于多个资产的历史数据定义聚合预测模型来开始。利用多个资产的数据可促进定义比利用单个资产的操作数据更准确的预测模型。

[0015] 形成聚合模型的基础的历史数据可包含指示给定资产的操作状况的至少操作数据。具体来说，操作数据可包含识别何时在资产处发生故障的情况的异常状况数据及/或指示在发生所述情况时在资产处测量的一或多个物理性质的数据。数据还可包含指示资产已在其中被操作的环境的环境数据及指示资产何时被利用的日期及时间的调度数据、用于定义聚合模型-工作流对的资产相关数据等等。

[0016] 基于历史数据，远程计算系统可定义预测特定事件发生的聚合模型。在特定的实

例实施方案中,聚合模型可输出将在未来特定时间段内在资产处发生故障的概率。此模型可在本文中被称为“故障模型”。除了其它实例预测模型外,其它聚合模型可预测资产将在未来特定时间段内完成任务的可能性。

[0017] 在定义聚合模型之后,远程计算系统可接着定义对应于所定义的聚合模型的聚合工作流。通常,工作流可包含资产可基于对应模型执行的一或多个操作。即,对应模型的输出可致使资产执行工作流操作。例如,聚合模型-工作流对可经定义使得当聚合模型输出特定范围内的概率时,资产将执行特定工作流操作(例如本地诊断工具)。

[0018] 在定义聚合模型-工作流对之后,远程计算系统可向一或多个资产发射所述对。一或多个资产接着可根据聚合模型-工作流对来操作。

[0019] 在实例实施方案中,远程计算系统可经配置以进一步定义一或多个资产的个性化预测模型及/或对应工作流。远程计算系统可基于每一给定资产的某些特性以及其它考虑来进行上述定义。在实例实施方案中,远程计算系统可以聚合模型-工作流对作为基准而开始,且基于资产的特性个性化给定资产的聚合模型及工作流中的一或两者。

[0020] 实际上,远程计算系统可经配置以确定与聚合模型-工作流对相关的资产特性(例如,所关注特性)。此类特性的实例可包含资产年限、资产使用状况、资产类别(例如,品牌及/或型号)、资产健康状况及资产的操作环境等其它特性。

[0021] 接着,远程计算系统可确定对应于所关注特性的给定资产的特性。至少基于给定资产特性中的一些,远程计算系统可经配置以个性化聚合模型及/或对应工作流。

[0022] 定义个性化模型及/或工作流可涉及远程计算系统对聚合模型及/或工作流进行某些修改。例如,除了其它实例外,个性化聚合模型可能涉及改变模型输入、改变模型计算,及/或改变计算的变量或输出的权重等等。除了其它实例外,个性化聚合工作流可能涉及改变工作流的一或多个操作及/或改变触发工作流的模型输出值或值范围。

[0023] 在定义给定资产的个性化模型及/或工作流之后,远程计算系统可接着向给定资产发射个性化模型及/或工作流。在其中模型或工作流中的仅一者被个性化的情况下,给定资产可利用模型或工作流的未被个性化的聚合版本。给定资产可接着根据其个性化模型-工作流对来操作。

[0024] 在实例实施方案中,给定资产可包含本地分析装置,其可经配置以致使给定资产根据由远程计算系统提供的模型-工作流对来操作。本地分析装置可经配置以利用来自资产传感器及/或致动器的操作数据(例如,通常用于其它资产相关目的的数据)来运行预测模型。当本地分析装置接收到某些操作数据时,其可执行所述模型,且取决于所述模型的输出,可执行对应工作流。

[0025] 执行对应工作流可帮助促进防止在给定资产处发生非期望事件。以此方式,给定资产可本地确定可能发生特定事件,且接着可执行特定工作流以帮助防止发生事件。如果给定资产与远程计算系统之间的通信受阻,那么这可能特别有用。例如,在某些情况下,故障可能发生在采取预防性动作的命令从远程计算系统到达给定资产之前。在此类情况下,本地分析装置可为有利的,因为其可在本地产生命令,由此避免任何网络延迟或由于给定资产“离线”而引起的任何问题。因而,本地分析装置执行模型-工作流对可促进致使资产适应其状况。

[0026] 在一些实例实施方案中,在首次执行模型-工作流对之前或之时,本地分析装置本

身可个性化其从远程计算系统接收到的模型-工作流对。通常,本地分析装置可通过评估在定义模型-工作流对时所进行的与给定资产相关的一些或全部预测、假设及/或一般化来个性化模型-工作流对。基于评估,本地分析装置可修改模型-工作流对,使得模型-工作流对的基础预测、假设及/或一般化更准确地反映给定资产的实际状态。本地分析装置接着可执行个性化模型-工作流对,而不是执行其最初从远程计算系统接收的模型-工作流对,这可导致对资产的更准确的监测。

[0027] 当给定资产根据模型-工作流对操作时,给定资产也可继续向远程计算系统提供操作数据。基于至少此数据,远程计算系统可修改聚合模型-工作流对及/或一或多个个性化模型-工作流对。远程计算系统可能由于多种原因而进行修改。

[0028] 在一个实例中,如果在资产处发生所述模型先前没有考虑到的新事件,那么远程计算系统可修改模型及/或工作流。例如,在故障模型中,新事件可能是在其数据用于定义聚合模型的资产中的任一者处尚未发生的新故障。

[0029] 在另一实例中,如果在资产处在通常不会导致事件发生的操作状况下发生所述事件,那么远程计算系统可修改模型及/或工作流。例如,再次返回到故障模型,如果在过去尚未导致故障发生的操作状况下发生所述故障,那么可修改所述故障模型或对应工作流。

[0030] 在又一实例中,如果所执行的工作流未能阻止事件的发生,那么远程计算系统可修改模型及/或工作流。具体来说,如果模型的输出致使资产执行旨在防止发生事件的工作流但是尽管如此所述事件还是发生在资产处,那么远程计算系统可修改模型及/或工作流。修改模型及/或工作流的原因的其它实例也是可能的。

[0031] 远程计算系统接着可将任何修改分配给其数据引起修改的资产及/或与远程计算系统通信的其它资产。以此方式,远程计算系统可动态地修改模型及/或工作流,且基于个别资产的操作状况将此类修改分配给整群资产。

[0032] 在一些实例实施方案中,资产及/或远程计算系统可经配置以动态调整执行预测模型及/或工作流。特定来说,资产及/或远程计算系统可经配置以检测触发关于资产及/或远程计算系统是否执行预测模型及/或工作流的责任的改变的某些事件。

[0033] 例如,在一些情况下,在资产从远程计算系统接收到模型-工作流对之后,资产可将模型-工作流对存储在数据存储装置中,而接着可依赖于远程计算系统集中地执行模型-工作流对中的部分或全部。另一方面,在其它情况下,远程计算系统可依赖于资产来本地执行模型-工作流对中的部分或全部。在又其它情况下,远程计算系统及资产可共享执行模型-工作流对的责任。

[0034] 无论如何,在某个时间点,可能发生触发资产及/或远程计算系统调整预测模型及/或工作流的执行的某些事件。例如,资产及/或远程计算系统可检测将资产耦合到远程计算系统的通信网络的某些特性。基于通信网络的特性,资产可调整其是否在本地执行预测模型及/或工作流,且远程计算系统可因此修改其是否集中执行模型及/或工作流。以此方式,资产及/或远程计算系统可适应资产的状况。

[0035] 在特定实例中,资产可检测如下指示:资产与远程计算系统之间的通信链路的信号强度相对较弱(例如,资产可确定即将“离线”)、网络延迟相对较高,及/或网络带宽相对较低。因此,资产可经编程以承担先前由远程计算系统处置的执行模型-工作流对的责任。进而,远程计算系统可停止集中执行模型-工作流对中的一些或全部。以此方式,资产可在

本地执行预测模型,且接着基于执行预测模型来执行对应工作流以潜在地帮助防止资产处发生故障。

[0036] 另外,在一些实施方案中,资产及/或远程计算系统可类似地基于各种其它考虑来调整执行(或可能修改)预测模型及/或工作流。例如,基于资产的处理能力,资产可在本地执行模型-工作流对,且远程计算系统也可因此进行调整。在另一实例中,基于将资产耦合到远程计算系统的通信网络的带宽,资产可执行所修改的工作流(例如,根据数据发射方案以减小的发射速率向远程计算系统发射数据)。其它实例也是可能的。

[0037] 如上文所讨论,本文中提供的实例涉及预测模型的部署及执行。在一个方面中,提供了一种计算系统。所述计算系统包括至少一个处理器、非暂时性计算机可读媒体及存储在所述非暂时性计算机可读媒体上的程序指令,所述程序指令可由所述至少一个处理器执行以致使所述计算系统:(a)接收多个资产的相应操作数据;(b)基于所述所接收的操作数据,定义与所述多个资产的操作相关的预测模型及对应工作流;及(c)向所述多个资产中的至少一个资产发射所述预测模型及对应工作流以供所述至少一个资产本地执行。

[0038] 在另一方面中,提供了一种上面存储有指令的非暂时性计算机可读媒体,所述指令可执行以致使计算系统:(a)接收多个资产的相应操作数据;(b)基于所述所接收的操作数据,定义与所述多个资产的操作相关的预测模型及对应工作流;及(c)向所述多个资产中的至少一个资产发射所述预测模型及对应工作流以供所述至少一个资产本地执行。

[0039] 在又一方面中,提供了一种计算机实施方法。所述方法包括:(a)接收多个资产的相应操作数据;(b)基于所述所接收的操作数据,定义与所述多个资产的操作相关的预测模型及对应工作流;及(c)向所述多个资产中的至少一个资产发射所述预测模型及对应工作流以供所述至少一个资产本地执行。

[0040] 如上文所讨论,本文中提供的实例涉及预测模型的部署及执行。在一个方面中,提供了一种计算系统。所述计算系统包括至少一个处理器、非暂时性计算机可读媒体及存储在所述非暂时性计算机可读媒体上的程序指令,所述程序指令可由所述至少一个处理器执行以致使所述计算系统:(a)接收多个资产的操作数据,其中所述多个资产包括第一资产;(b)基于所述所接收的操作数据,定义与所述多个资产的操作相关的聚合预测模型及聚合对应工作流;(c)确定所述第一资产的一或多个特性;(d)基于所述第一资产的所述一或多个特性以及所述聚合预测模型及所述聚合对应工作流,定义与所述第一资产的所述操作相关的个性化预测模型或个性化对应工作流中的至少一者;及(e)向所述第一资产发射所述所定义的至少一个个性化预测模型或个性化对应工作流以供所述第一资产本地执行。

[0041] 在另一方面中,提供了一种上面存储有指令的非暂时性计算机可读媒体,所述指令可执行以致使计算系统:(a)接收多个资产的操作数据,其中所述多个资产包括第一资产;(b)基于所述所接收的操作数据,定义与所述多个资产的操作相关的聚合预测模型及聚合对应工作流;(c)确定所述第一资产的一或多个特性;(d)基于所述第一资产的所述一或多个特性以及所述聚合预测模型及所述集合对应工作流,定义与所述第一资产的所述操作相关的个性化预测模型或个性化对应工作流中的至少一者;及(e)向所述第一资产发射所述所定义的至少一个个性化预测模型或个性化对应工作流以供所述第一资产本地执行。

[0042] 在又一方面中,提供了一种计算机实施方法。所述方法包括:(a)接收多个资产的操作数据,其中所述多个资产包括第一资产;(b)基于所述所接收的操作数据,定义与所述

多个资产的操作相关的聚合预测模型及聚合对应工作流；(c)确定所述第一资产的一或多个特性；(d)基于所述第一资产的所述一或多个特性以及所述聚合预测模型及所述聚合对应工作流，定义与所述第一资产的所述操作相关的个性化预测模型或个性化对应工作流中的至少一者；及(e)向所述第一资产发射所述所定义的至少一个个性化预测模型或个性化对应工作流以供所述第一资产本地执行。

[0043] 如上文所讨论，本文中提供的实例涉及在资产处接收并执行预测模型和/或工作流。在一个方面中，提供了一种计算装置。所述计算装置包括：(i)资产接口，其经配置以将所述计算装置耦合到资产；(ii)网络接口，其经配置以促进所述计算装置与远离所述计算装置定位的计算系统之间的通信；(iii)至少一个处理器；(iv)非暂时性计算机可读媒体；及(v)存储在所述非暂时性计算机可读媒体上的程序指令，所述程序指令可由所述至少一个处理器执行以致使所述计算装置：(a)经由所述网络接口接收与所述资产的操作相关的预测模型，其中所述预测模型由所述计算系统基于多个资产的操作数据来定义；(b)经由所述资产接口接收所述资产的操作数据；(c)基于所述资产的所接收的操作数据的至少部分执行所述预测模型；及(d)基于执行所述预测模型，执行对应于所述预测模型的工作流，其中执行所述工作流包括经由所述资产接口致使所述资产执行操作。

[0044] 在另一方面中，提供了一种上面存储有指令的非暂时性计算机可读媒体，所述指令可执行以致使经由计算装置的资产接口耦合到资产的所述计算装置：(a)经由所述计算装置的网络接口接收与所述资产的所述操作相关的预测模型，所述计算装置的所述网络接口经配置以促进所述计算装置与远离所述计算装置定位的计算系统之间的通信，其中所述预测模型由所述计算系统基于多个资产的操作数据来定义；(b)经由所述资产接口接收所述资产的操作数据；(c)基于所述资产的所接收的操作数据的至少部分执行所述预测模型；及(c)基于执行所述预测模型，执行对应于所述预测模型的工作流，其中执行所述工作流包括经由所述资产接口致使述资产执行操作。

[0045] 在又一方面中，提供了一种计算机实施方法。所述方法包括：(a)经由计算装置的网络接口接收与资产的所述操作相关的预测模型，所述计算装置的所述网络接口经由所述计算装置的资产接口耦合到所述资产，其中所述预测模型由远离所述计算装置定位的所述计算系统基于多个资产的操作数据来定义；(b)由所述计算装置经由所述资产接口接收所述资产的操作数据；(b)由所述计算装置基于所述资产的所接收的操作数据的至少部分执行所述预测模型；及(c)基于执行所述预测模型，由所述计算装置执行对应于所述预测模型的工作流，其中执行所述工作流包括经由所述资产接口致使所述资产执行操作。

[0046] 所属领域一般技术人员在阅读下面的揭示内容时将明白这些以及许多其它方面。

附图说明

[0047] 图1描绘了其中可实施实例实施例的实例网络配置。

[0048] 图2描绘了实例资产的简化框图。

[0049] 图3描绘了实例异常状况指示符及触发准则的概念说明。

[0050] 图4描绘了实例分析系统的简化框图。

[0051] 图5描绘了可用于定义模型-工作流对的定义阶段的实例流程图。

[0052] 图6A描绘了聚合模型-工作流对的概念说明。

[0053] 图6B描绘了个性化模型-工作流对的概念说明。

[0054] 图6C描绘了另一个性化模型-工作流对的概念说明。

[0055] 图6D描绘了所修改的模型-工作流对的概念说明。

[0056] 图7描绘了可用于定义输出健康状况指标的预测模型的建模阶段的实例流程图。

[0057] 图8描绘了用于定义模型的数据的概念说明。

[0058] 图9描绘了可用于本地执行预测模型的本地执行阶段的实例流程图。

[0059] 图10描绘了可用于修改模型-工作流对的修改阶段的实例流程图。

[0060] 图11描绘了可用于调整模型-工作流对的执行的调整阶段的实例流程图。

[0061] 图12描绘了用于定义及部署聚合预测模型及对应工作流的实例方法的流程图。

[0062] 图13描绘了用于定义及部署个性化预测模型及/或对应工作流的实例方法的流程图。

[0063] 图14描绘了用于动态地修改模型-工作流对的执行的实例方法的流程图。

[0064] 图15描绘了用于接收及本地执行模型-工作流对的实例方法的流程图。

具体实施方式

[0065] 以下揭示内容参考附图及几个示范性情况。所属领域的一般技术人员将会理解，此类参考仅用于解释的目的，且因此不意味着限制。所揭示的系统、装置及方法中的部分或全部可以各种方式重新布置、组合、添加及/或移除，其中的每一方式均在本文中予以考虑。

[0066] I.实例网络配置

[0067] 现在转向图式，图1描绘了其中可实施实例实施例的实例网络配置100。如所示，网络配置100包含资产102、资产104、通信网络106、可采取分析系统的形式的远程计算系统108、输出系统110及数据源112。

[0068] 通信网络106可通信地连接网络配置100中的组件中的每一者。例如，资产102及104可经由通信网络106与分析系统108通信。在一些情况下，资产102及104可与一或多个中间系统(例如资产网关(未描绘))通信，所述资产网关进而与分析系统108通信。类似地，分析系统108可经由通信网络106与输出系统110通信。在一些情况下，分析系统108可与一或多个中间系统(例如主机服务器(未描绘))通信，所述主机服务器进而与输出系统110通信。许多其它配置也是可能的。在实例情况下，通信网络106可促进网络组件之间的安全通信(例如，经由加密或其它安全措施)。

[0069] 通常，资产102及104可采取经配置以执行一或多个操作(其可基于领域来定义)的任何装置的形式，且还可包含经配置以发射指示给定资产的一或多个操作状况的数据的设备。在一些实例中，资产可包含经配置以执行一或多个相应操作的一或多个子系统。实际上，多个子系统可并行或按顺序操作以使资产操作。

[0070] 实例资产可包含运输机器(例如，机车、飞机、乘用车辆、半挂卡车、船舶等)、工业机器(例如，采矿设备、建筑设备、工厂自动化设备等)、医疗机器(例如，医学成像设备、外科手术设备、医疗监测系统、医学实验室设备等)及实用机器(例如，涡轮机、太阳能农场等)等等。所属领域的一般技术人员将明白，这些仅仅是资产的几个实例，且数个其它资产在本文中是可能的及予以考虑的。

[0071] 在实例实施方案中，资产102及104可各自为相同的类型(例如，一队机车或飞机、

风力涡轮机组或MRI机器集合等等)，且可能为相同类别(例如，相同品牌及/或型号)。在其它实例中，资产102及104在类型、品牌、型号等方面可能不同。下文参考图2进一步详细讨论资产。

[0072] 如所示，资产102及104及可能的数据源112可经由通信网络106与分析系统108通信。通常，通信网络106可包含经配置以促进在网络组件之间传送数据的一或多个计算系统及网络基础设施。通信网络106可为或可包含可为有线及/或无线的且支持安全通信的一或多个广域网(WAN)及/或局域网(LAN)。在一些实例中，通信网络106可包含一或多个蜂窝网络及/或因特网等网络。通信网络106可根据例如LTE、CDMA、GSM、LPWAN、WiFi、蓝牙、以太网、HTTP/S、TCP、CoAP/DTLS等的一或多个通信协议进行操作。虽然通信网络106被示为单个网络，但是应理解的是，通信网络106可包含本身通信地链接的多个不同的网络。通信网络106也可采取其它形式。

[0073] 如上文所述，分析系统108可经配置以从资产102及104及数据源112接收数据。一般来说，分析系统108可包含经配置以接收、处理、分析及输出数据的一或多个计算系统，例如服务器及数据库。分析系统108可根据给定数据流技术(例如TPL Dataflow或NiFi等等)来配置。在下文参考图3进一步详细讨论分析系统108。

[0074] 如所示，分析系统108可经配置以向资产102及104及/或输出系统110发射数据。所发射的特定数据可采取各种形式，并将在下文进一步详细描述。

[0075] 通常，输出系统110可采取经配置以接收数据并提供某种形式的输出的计算系统或装置的形式。输出系统110可采取各种形式。在一个实例中，输出系统110可为或包含输出装置，其经配置以接收数据并响应于所述数据而提供听觉、视觉及/或触觉输出。通常，输出装置可包含经配置以接收用户输入的一或多个输入接口，且输出装置可经配置以基于此用户输入而通过通信网络106发射数据。输出装置的实例包含平板计算机、智能电话、膝上型计算机、其它移动计算装置、台式计算机、智能电视等。

[0076] 输出系统110的另一实例可采取工作令系统的形式，所述工作令系统经配置以输出使机械师等维修资产的请求。输出系统110的又另一实例可采取经配置以对资产的零件下订单并输出其收据的零件订购系统的形式。许多其它输出系统也是可能的。

[0077] 数据源112可经配置以与分析系统108通信。通常，数据源112可为或包含一或多个计算系统，其经配置以收集、存储数据及/或向其它系统(例如分析系统108)提供数据，所述数据可与由分析系统108执行的功能相关。数据源112可经配置以独立于资产102及104产生及/或获得数据。因而，由数据源112提供的数据在本文中可被称为“外部数据”。数据源112可经配置以提供当前及/或历史数据。实际上，分析系统108可通过“订阅”由数据源提供的服务来从数据源112接收数据。然而，分析系统108也可以其它方式从数据源112接收数据。

[0078] 数据源112的实例包含环境数据源、资产管理数据源及其它数据源。通常，环境数据源提供指示资产的操作环境的某种特性的数据。环境数据源的实例包含提供关于给定区域的自然及人造特征的信息的气象数据服务器、全球导航卫星系统(GNSS)服务器、地图数据服务器及拓扑数据服务器等等。

[0079] 通常，资产管理数据源提供指示可影响资产的操作或维护的实体(例如，其它资产)的事件或状态的数据(例如，资产可在何时及何处操作或接收维护)。资产管理数据源的实例包含：交通数据服务器，其提供关于空中、水上及/或地面交通的信息；资产调度服务

器，其提供关于资产在特定日期及/或特定时间的预期路线及/或位置的信息；缺陷检测器系统(也称为“热箱”检测器)，其提供关于经过缺陷检测器系统附近的资产的一或多个操作状况的信息；零件供应商服务器，其提供关于特定供应商的库存中的零件及其价格的信息；及维修车间服务器，其提供关于维修车间产能等的信息；等等。

[0080] 其它数据源的实例包含提供关于电力消耗的信息的电网服务器及存储资产的历史操作数据的外部数据库等等。所属领域的一般技术人员将明白，这些仅仅是数据源的几个实例，且数个其它实例是可能的。

[0081] 应理解的是，网络配置100是其中可实施本文中描述的实施例的网络的一个实例。数个其它布置是可能的且在本文中予以考虑。例如，其它网络配置可包含未描绘的额外组件及/或更多或更少的所描绘的组件。

[0082] II.实例资产

[0083] 转向图2，描绘了实例资产200的简化框图。来自图1的资产102及104中的任一者或两者可如同资产200一样进行配置。如所示，资产200可包含一或多个子系统202、一或多个传感器204、一或多个致动器205、中央处理单元206、数据存储装置208、网络接口210、用户接口212及本地分析装置220，其全部可通过系统总线、网络或其它连接机构(直接地或间接地)通信地链接。所属领域的一般技术人员将明白，资产200可包含未展示的额外部件及/或更多或更少的所描绘组件。

[0084] 一般来说，资产200可包含经配置以执行一或多个操作的一或多个电气、机械及/或机电组件。在一些情况下，一或多个组件可被分组到给定子系统202中。

[0085] 通常，子系统202可包含作为资产200的部分的相关组件群组。单个子系统202可独立地执行一或多个操作，或单个子系统202可与一或多个其它子系统一起操作以执行一或多个操作。通常，不同类型的资产且甚至不同类别的相同类型的资产可包含不同子系统。

[0086] 例如，在运输资产的背景下，子系统202的实例可包含发动机、变速装置、传动系、燃料系统、电池系统、排气系统、制动系统、电气系统、信号处理系统、发电机、齿轮箱、转子及液压系统，以及数个其它子系统。在医疗机器的背景下，子系统202的实例可包含扫描系统、电动机、线圈及/或磁体系统、信号处理系统、转子及电气系统，以及数个其它子系统。

[0087] 如上文所指示，资产200可配备有：各种传感器204，其经配置以监测资产200的操作状况；及各种致动器205，其经配置以与资产200或其组件交互并监测资产200的操作状况。在一些情况下，传感器204及/或致动器205中的一些可基于特定子系统202被分组。以此方式，传感器204及/或致动器205的群组可经配置以监测特定子系统202的操作状况，且来自所述群组的致动器可经配置以按照某种方式与特定子系统202交互，所述方式可基于那些操作状况而改变子系统的行为。

[0088] 通常，传感器204可经配置以检测可指示资产200的一或多个操作状况的物理性质，且提供所检测的物理性质的指示，例如电信号。在操作中，传感器204可经配置以连续地、周期性地(例如，基于采样频率)及/或响应于某个触发事件来获得测量值。在一些实例中，传感器204可预先配置有用于执行测量的操作参数及/或可根据由中央处理单元206提供的操作参数(例如，指示传感器204获得测量值的采样信号)来执行测量。在实例中，不同传感器204可具有不同操作参数(例如，一些传感器可基于第一频率进行采样，而其它传感器基于第二不同频率进行采样)。无论如何，传感器204可经配置以向中央处理单元206发射

指示所测量的物理性质的电信号。传感器204可连续地或周期性地将此类信号提供给中央处理单元206。

[0089] 例如，传感器204可经配置以测量资产200的物理性质，例如资产200的位置及/或移动，在所述情况下，传感器可采取GNSS传感器、基于航位推测的传感器、加速度计、陀螺仪、计步器、磁力计等形式。

[0090] 另外，各种传感器204可经配置以测量资产200的其它操作状况，其实例可包含温度、压力、速度、加速率或减速率、摩擦、功率使用量、燃料使用量、液面、运行时间、电压及电流、磁场、电场、物体的存在或不存在、组件的位置及发电等等。所属领域的一般技术人员将明白，这些仅仅是传感器可经配置以测量的一些实例操作状况。取决于行业应用或特定资产，可使用更多或更少的传感器。

[0091] 如上文所指示，致动器205的配置在某些方面可类似于传感器204。具体来说，致动器205可经配置以检测指示资产200的操作状况的物理性质，且以与传感器204类似的方式提供其指示。

[0092] 另外，致动器205可经配置以与资产200、一或多个子系统202及/或其一些组件交互。因而，致动器205可包含经配置以执行机械操作(例如，移动)或以其它方式控制组件、子系统或系统的电动机等。在特定实例中，致动器可经配置以测量燃料流量并改变燃料流量(例如限制燃料流量)，或致动器可经配置以测量液压压力并改变液压压力(例如，增大或减小液压压力)。致动器的数个其它实例交互也是可能的且在本文中予以考虑。

[0093] 通常，中央处理单元206可包含一或多个处理器及/或控制器，其可采取通用或专用处理器或控制器的形式。具体来说，在实例实施方案中，中央处理单元206可为或包含微处理器、微控制器、专用集成电路、数字信号处理器等。进而，数据存储装置208可为或包含一或多个非暂时性计算机可读存储媒体，例如光学、磁性、有机或快闪存储器等等。

[0094] 中央处理单元206可经配置以存储、存取及执行存储在数据存储装置208中的计算机可读程序指令，以执行本文中描述的资产的操作。例如，如上文所指示，中央处理单元206可经配置以从传感器204及/或致动器205接收相应的传感器信号。中央处理单元206可经配置以将传感器及/或致动器数据存储在数据存储装置208中且随后从数据存储装置208存取所述数据。

[0095] 中央处理单元206还可经配置以确定所接收的传感器及/或致动器信号是否触发任何异常状况指示符，例如故障代码。例如，中央处理单元206可经配置以在数据存储装置208中存储异常状况规则，其中的每一规则包含表示特定异常状况的给定异常状况指示符及触发异常状况指示符的相应触发准则。即，每一异常状况指示符对应于在异常状态指示符被触发之前必须得到满足的一或多个传感器及/或致动器测量值。实际上，资产200可预编程有异常状况规则及/或可从计算系统(例如分析系统108)接收新的异常状况规则或对现有规则的更新。

[0096] 无论如何，中央处理单元206可经配置以确定所接收的传感器及/或致动器信号是否触发任何异常状况指示符。即，中央处理单元206可确定所接收的传感器及/或致动器信号是否满足任何触发准则。当此确定为肯定时，中央处理单元206可产生异常状态数据，且还可致使资产的用户接口212输出异常情况的指示，例如视觉及/或声讯警报。另外，中央处理单元206可可能以时间戳记在数据存储装置208中记录所触发的异常状况指示符的发生。

[0097] 图3描绘了资产的实例异常状况指示符及相应触发准则的概念说明。特定来说，图3描绘了实例故障代码的概念说明。如所示，表300包含分别对应于传感器A、致动器B及传感器C的列302、304及306及分别对应于故障代码1、2及3的行308、310及312。条目314接着指定与给定故障代码对应的传感器准则(例如，传感器值阈值)。

[0098] 例如，当传感器A检测到大于135转每分钟(RPM)的旋转测量值且传感器C检测到大于65摄氏度(C)的温度测量值时，将触发故障代码1。当致动器B检测到大于1000伏(V)的电压测量值且传感器C检测到小于55℃的温度测量值时，将触发故障代码2。当传感器A检测到大于100RPM的旋转测量值、致动器B检测到大于750V的电压测量值且传感器C检测到大于60℃的温度测量值时，将触发故障代码3。所属领域的一般技术人员将明白，图3的提供仅仅是出于实例及解释目的，且许多其它故障代码及/或触发准则在本文中是可能的且予以考虑。

[0099] 返回参考图2，中央处理单元206还可经配置以实行用于管理及/或控制资产200的操作的各种额外功能。例如，中央处理单元206可经配置以向子系统202及/或致动器205提供指令信号，所述指令信号致使子系统202及/或致动器205执行某个操作(例如修改节气门位置)。另外，中央处理单元206可经配置以修改其处理来自传感器204及/或致动器205的数据的速率，或中央处理单元206可经配置以向传感器204及/或致动器205提供指令信号，所述指令信号致使传感器204及/或致动器205例如修改采样速率。另外，中央处理单元206可经配置以从子系统202、传感器204、致动器205、网络接口210及/或用户接口212接收信号，且基于此类信号致使操作发生。另外，中央处理单元206可经配置以从例如诊断装置等的计算装置接收信号，所述信号致使中央处理单元206根据存储在数据存储装置208中的诊断规则执行一或多个诊断工具。下文讨论中央处理单元206的其它功能性。

[0100] 网络接口210可经配置以提供资产200与连接到通信网络106的各种网络组件之间的通信。例如，网络接口210可经配置以促进去向及来自通信网络106的无线通信，且因此可采取用于发射及接收各种无线信号(over-the-air signal)的天线结构及相关联装置的形式。其它实例也是可能的。实际上，网络接口210可根据通信协议(例如但不限于上述通信协议中的任一者)来配置。

[0101] 用户接口212可经配置以促进用户与资产200的交互，且还可经配置以促进致使资产200响应于用户交互来执行操作。用户接口212的实例包含触敏接口、机械接口(例如，杠杆、按钮、滚轮、拨号盘、键盘等)以及其它输入接口(例如麦克风)等等。在一些情况下，用户接口212可包含或提供到输出组件(例如显示屏、扬声器、耳机插孔等)的连接性。

[0102] 本地分析装置220通常可经配置以接收及分析与资产200相关的数据，且基于此分析可致使在资产200处发生一或多个操作。例如，本地分析装置220可接收资产200的操作数据(例如，由传感器204及/或致动器205产生的数据)且基于此数据可向中央处理单元206、传感器204及/或致动器205提供致使资产200执行操作的指令。

[0103] 为了促进此操作，本地分析装置220可包含一或多个资产接口，其经配置以将本地分析装置220耦合到资产的机载系统中的一或多者。例如，如图2中所示，本地分析装置220可具有到资产的中央处理单元206的接口，其可使得本地分析装置220能够从中央处理单元206接收操作数据(例如，由传感器204及/或致动器205产生并发送到中央处理单元206的操作数据)，且接着向中央处理单元206提供指令。以此方式，本地分析装置220可经由中央处理单元206间接地与资产200的其它机载系统(例如，传感器204及/或致动器205)介接并从

其中接收数据。另外或替代地，如图2中所示，本地分析装置220可具有到一或多个传感器204及/或致动器205的接口，其可使得本地分析装置220能够与传感器204及/或致动器205直接通信。本地分析装置220也可以其它方式与资产200的机载系统介接，包含图2中所说明的接口由未展示的一或多个中间系统促进的可能性。

[0104] 实际上，本地分析装置220可使得资产200能够在本地执行高级分析及相关联的操作，例如执行预测模型及对应工作流，所述操作可能不能够使用其它资产上组件执行。因而，本地分析装置220可帮助向资产200提供额外的处理能力及/或智能。

[0105] 应当理解的是，本地分析装置220还可经配置以致使资产200执行与预测模型无关的操作。例如，本地分析装置220可从远程源(例如分析系统108或输出系统110)接收数据，且基于所接收的数据致使资产200执行一或多个操作。一个特定实例可涉及本地分析装置220从远程源接收资产200的固件更新，且接着致使资产200更新其固件。另一特定实例可涉及本地分析装置220从远程源接收诊断指令，且接着根据所接收的指令致使资产200执行本地诊断工具。数个其它实例也是可能的。

[0106] 如所示，除了上文讨论的一或多个资产接口之外，本地分析装置220还可包含处理单元222、数据存储装置224及网络接口226，其全部均可通过系统总线、网络或其它连接机制通信地链接。处理单元222可包含上文关于中央处理单元206讨论的组件中的任一者。进而，数据存储装置224可为或可包含一或多个非暂时性计算机可读存储媒体，其可采用上文讨论的计算机可读存储媒体的形式中的任一者。

[0107] 处理单元222可经配置以存储、存取及执行存储在数据存储装置224中的计算机可读程序指令，以执行本文中描述的本地分析装置的操作。例如，处理单元222可经配置以接收由传感器204及/或致动器205产生的相应的传感器及/或致动器信号，且可基于此类信号执行预测模型-工作流对。其它功能在下文予以描述。

[0108] 网络接口226可与上述网络接口相同或类似。实际上，网络接口226可促进本地分析装置220与分析系统108之间的通信。

[0109] 在一些实例实施方案中，本地分析装置220可包含可与用户接口212类似的用户接口及/或与所述用户接口通信。实际上，用户接口可位于远离本地分析装置220(及资产200)的位置。其它实例也是可能的。

[0110] 虽然图2展示了本地分析装置220经由一或多个资产接口物理地且通信地耦合到其相关联的资产(例如，资产200)，但是也应当理解的是，情况可能并非总是如此。例如，在一些实施方案中，本地分析装置220可不物理地耦合到其相关联的资产，而是可位于远离资产220的位置。在此实施方案的实例中，本地分析装置220可无线地通信地耦合到资产200。其它布置及配置也是可能的。

[0111] 所属领域的一般技术人员将明白，图2中所示的资产200仅仅是资产的简化表示的一个实例，且许多其它实例也是可能的。例如，其它资产可包含未描绘的额外组件及/或更多或更少的所描绘的组件。另外，给定资产可包含一致操作以执行给定资产的操作的多个个别资产。其它实例也是可能的。

[0112] III.实例分析系统

[0113] 现在转向图4，描绘了实例分析系统400的简化框图。如上文所指示，分析系统400可包含通信地链接且经布置以实行本文中描述的各种操作的一或多个计算系统。具体来

说，如所示，分析系统400可包含数据摄入系统402、数据科学系统404及一或多个数据库406。这些系统组件可经由一或多个无线及/或有线连接通信地耦合，所述连接可经配置以促进安全通信。

[0114] 数据摄入系统402通常可用于接收及处理数据并将数据输出到数据科学系统404。因而，数据摄入系统402可包含一或多个网络接口，其经配置以从网络配置100的各种网络组件(例如资产102及104、输出系统110及/或数据源112)接收数据。具体来说，数据摄入系统402可经配置以接收模拟信号、数据流及/或网络分组等等。因而，网络接口可包含一或多个有线网络接口(例如端口等)及/或无线网络接口，类似于上文描述的无线网络接口。在一些实例中，数据摄入系统402可为或包含根据给定数据流技术配置的组件，例如NiFi接收器等。

[0115] 数据摄入系统402可包含经配置以执行一或多个操作的一或多个处理组件。实例操作可包含压缩及/或解压缩、加密及/或解密、模/数转换及/或数/模转换、筛选及放大等其它操作。另外，数据摄入系统402可经配置以基于数据的数据类型及/或数据特性来解析、分类、组织及/或路由数据。在一些实例中，数据摄入系统402可经配置以基于数据科学系统404的一或多个特性或操作参数来格式化、封装及/或路由数据。

[0116] 通常，由数据摄入系统402接收的数据可采取各种形式。例如，数据的有效载荷可包含单个传感器或致动器测量值、多个传感器及/或致动器测量值及/或一或多个异常状况数据。其它实例也是可能的。

[0117] 另外，所接收的数据可包含某些特性，例如源标识符及时间戳记(例如，获得信息的日期及/或时间)。例如，可将唯一标识符(例如，计算机产生的字母、数字、字母数字或类似标识符)指派给每一资产，且可能指派给每一传感器及致动器。此类标识符可操作以识别数据所来源于的资产、传感器或致动器。在一些情况下，另一特性可包含获得信息的位置(例如，GPS坐标)。数据特性可以信号签名或元数据等等形式出现。

[0118] 数据科学系统404通常可用于(例如，从数据摄入系统402)接收数据并分析数据，并基于此分析致使一或多个操作发生。因而，数据科学系统404可包含一或多个网络接口408、处理单元410及数据存储装置412，其全部均可通过系统总线、网络或其它连接机制通信地链接。在一些情况下，数据科学系统404可经配置以存储及/或存取促进实行本文中揭示的功能性中的一些的一或多个应用程序接口(API)。

[0119] 网络接口408可与上述任何网络接口相同或类似。实际上，网络接口408可促进数据科学系统404与各种其它实体(例如数据摄入系统402、数据库406、资产102、输出系统110等)之间的通信(例如，具有某种安全级别)。

[0120] 处理单元410可包含一或多个处理器，其可采用上述处理器形式中的任一者。进而，数据存储装置412可为或可包含一或多个非暂时性计算机可读存储媒体，其可采用上文讨论的计算机可读存储媒体的形式中的任一者。处理单元410可经配置以存储、存取及执行存储在数据存储装置412中的计算机可读程序指令，以执行本文中描述的分析系统的操作。

[0121] 通常，处理单元410可经配置以对从数据摄入系统402接收的数据执行分析。为此，处理单元410可经配置以执行一或多个模块，其各自可采取存储在数据存储装置412中的一或多个程序指令集的形式。模块可经配置以促进致使结果基于相应程序指令的执行而出现。来自给定模块的实例结果可包含将数据输出到另一模块、更新给定模块及/或另一模块

的程序指令，且将数据输出到网络接口408以发射到资产及/或输出系统110等等。

[0122] 数据库406通常可用于(例如，从数据科学系统404)接收数据并存储数据。因而，每一数据库406可包含一或多个非瞬时计算机可读存储媒体，例如上文提供的实例中的任一者。实际上，数据库406可与数据存储装置412分离或与数据存储装置412集成。

[0123] 数据库406可经配置以存储数种类型的数据，其中的一些在下文讨论。实际上，存储在数据库406中的数据中的一些可包含时间戳记，其指示数据被产生或添加到数据库的日期及时间。另外，数据可以数种方式存储在数据库406中。除其它实例外，例如，数据可按照时间顺序、表格方式存储，及/或基于数据源类型(例如，基于资产、资产类型、传感器、传感器类型、致动器或致动器类型)或异常状况指示符来组织。

[0124] IV.实例操作

[0125] 现在将在下文进一步详细讨论图1中描绘的实例网络配置100的操作。为了帮助描述此类操作中的一些，可参考流程图来描述可执行的操作的组合。在一些情况下，每一框可表示程序代码的模块或部分，其包含可由处理器执行以实施过程中的特定逻辑功能或步骤的指令。程序代码可存储在任何类型的计算机可读媒体上，例如非暂时性计算机可读媒体上。在其它情况下，每一框可表示经布线以执行过程中的特定逻辑功能或步骤的电路。另外，流程图中所示的框可重新布置为不同次序，组合成更少的框，分成额外的框，及/或基于特定实施例而移除。

[0126] 以下描述可参考其中例如资产102等的单个数据源向接着执行一或多个功能的分析系统108提供数据的实例。应当理解的是，这仅仅是为了清楚及解释而进行的，且并不意味着限制。实际上，分析系统108通常同时从多个源接收数据，且基于此聚合的接收数据来执行操作。

[0127] A.操作数据的收集

[0128] 如上文所提及，典型资产102可采取各种形式且可经配置以执行多个操作。在非限制性实例中，资产102可采取可操作以在美国各地转运货物的机车的形式。在运送时，资产102的传感器及/或致动器可获得反映资产102的一或多个操作状况的数据。传感器及/或致动器可向资产102的处理单元发射数据。

[0129] 处理单元可经配置以从传感器及/或致动器接收数据。实际上，处理单元可同时或按顺序接收来自多个传感器的传感器数据及/或来自多个致动器的致动器数据。如上文所讨论，在接收到此数据时，处理单元还可经配置以确定数据是否满足触发任何异常状况指示符(例如故障代码)的触发准则。在处理单元确定触发了一或多个异常状态指示符的情况下，处理单元可经配置以执行一或多个本地操作，例如经由用户接口输出被触发指示符的指示。

[0130] 资产102接着可经由资产102的网络接口及通信网络106向分析系统108发射操作数据。在操作中，资产102可连续地、周期性地及/或响应于触发事件(例如，异常状况)向操作分析系统108发射操作数据。具体来说，资产102可基于特定频率(例如，每天、每小时、每十五分钟、每分钟一次、每秒一次等)周期性地发射操作数据，或资产102可经配置以发射操作数据的连续的实时馈送。另外或替代地，例如当传感器及/或致动器测量值满足用于任何异常状况指示符的触发准则时，资产102可经配置以基于某些触发来发射操作数据。资产102也可以其它方式发射操作数据。

[0131] 实际上，资产102的操作数据可包含传感器数据、致动器数据及/或异常状况数据。在一些实施方案中，资产102可经配置以在单个数据流中提供操作数据，而在其它实施方案中，资产102可经配置以在多个不同的数据流中提供操作数据。例如，资产102可向分析系统108提供传感器及/或致动器数据的第一数据流及异常状况数据的第二数据流。其它可能性也存在。

[0132] 传感器及致动器数据可能采取各种形式。例如，有时，传感器数据(或致动器数据)可包含由资产102的传感器(或致动器)中的每一者获得的测量值。而在其它时间，传感器数据(或致动器数据)可包含由资产102的传感器(或致动器)的子集获得的测量值。

[0133] 具体来说，传感器及/或致动器数据可包含由与给定被触发的异常状况指示符相关联的传感器及/或致动器获得的测量值。例如，如果被触发的故障代码是来自图3的故障代码1，那么传感器数据可包含由传感器A及C获得的原始测量值。另外或替代地，数据可包含由与被触发的故障代码不直接相关联的一或多个传感器或致动器获得的测量值。继续最后的实例，数据可另外包含由致动器B及/或其它传感器或致动器获得的测量值。在一些实例中，资产102可基于由分析系统108提供的故障代码规则或指令在操作数据中包含特定的传感器数据，其例如可确定在致动器B正进行测量的测量值与起初引起故障代码1被触发的测量值之间存在相关性。其它实例也是可能的。

[0134] 另外，数据可包含基于特定所关注时间的来自每一所关注的感测器及/或致动器的一或多个传感器及/或致动器测量值，所述时间可基于数个因素来选择。在一些实例中，所关注的特定时间可基于采样速率。在其它实例中，所关注的特定时间可基于触发异常状况指示符的时间。

[0135] 特定来说，基于触发异常状况指示符的时间，数据可包含来自每一所关注的传感器及/或致动器(例如，与被触发的指示符直接及间接相关联的传感器及/或致动器)的一或多个相应的传感器及/致动器测量值。一或多个测量值可基于在触发的异常状况指示符的时间周围的测量的特定次数或特定持续时间。

[0136] 例如，如果触发的故障代码是来自图3的故障代码2，那么所关注的传感器及致动器可包含致动器B及传感器C。一或多个测量值可包含在触发故障代码(例如，触发测量)之前由致动器B及传感器C获得的最近的相应测量值或在触发测量之前、之后或附近的相应测量值集合。例如，一组五次测量可包含触发测量之前或之后的五次测量(例如，不包含触发测量)、触发测量之前或之后的四次测量及触发测量，或触发测量之前的两次测量及触发测量之后的两次测量以及触发测量以及其它可能性。

[0137] 与传感器及致动器数据类似，异常状况数据可采取各种形式。通常，异常情况数据可包含或采取指示符的形式，所述指示符可操作以从在资产102处可能发生的所有其它异常状况中唯一地识别在资产102处发生的特定异常状况。异常状况指示符可采取字母、数字或字母数字标识符等等的形式。另外，异常状况指示符可采取描述异常状况的字串的形式，例如“发动机过热”或“燃料不足”等等。

[0138] 分析系统108，且特别是分析系统108的数据摄入系统可经配置以从一或多个资产及/或数据源接收操作数据。数据摄入系统可经配置以对所接收的数据执行一或多个操作，且接着将数据中继到分析系统108的数据科学系统。进而，数据科学系统可分析所接收的数据并基于此分析执行一或多个操作。

[0139] B.定义预测模型及工作流

[0140] 作为一个实例，分析系统108可经配置以基于一或多个资产的所接收的操作数据及/或与一或多个资产相关的所接收的外部数据来定义预测模型及对应工作流。分析系统108也可基于各种其它数据定义模型-工作流对。

[0141] 通常，模型-工作流对可包含程序指令集，其致使资产监测某些操作状况并实行某些操作，所述操作帮助促进防止发生监测操作状况所指示的特定事件。具体来说，预测模型可包含一或多个算法，所述算法的输入是来自资产的一或多个传感器及/或致动器的传感器及/或致动器数据，且其输出用于确定在未来特定时间段内在资产处可能发生特定事件的概率。进而，工作流可包含一或多个触发(例如，模型输出值)及资产基于触发实行的对应操作。

[0142] 如上文所指示，分析系统108可经配置以定义聚合及/或个性化预测模型及/或工作流。“聚合”模型/工作流可指代如下模型/工作流：对资产群组通用且在不考虑部署有模型/工作流的资产的特定特性的情况下进行定义的模型/工作流。另一方面，“个性化”模型/工作流可指代如下模型/工作流：针对来自资产群组的单个资产或资产子群组具体定制且基于部署有模型/工作流的单个资产或资产子群组的特定特性进行定义的模型/工作流。下文进一步详细讨论这些不同类型的模型/工作流及由分析系统108执行以定义模型/工作流的操作。

[0143] 1.聚合模型及工作流

[0144] 在实例实施方案中，分析系统108可经配置以基于多个资产的聚合数据来定义聚合模型-工作流对。定义聚合模型-工作流对可以各种方式执行。

[0145] 图5是描绘可用于定义模型-工作流对的定义阶段的一个可能实例的流程图500。出于说明目的，实例定义阶段被描述为由分析系统108实行，但是此定义阶段也可由其它系统实行。所属领域的一般技术人员将明白，出于清楚及解释目的提供流程图500，且可利用操作的数个其它组合来定义模型-工作流对。

[0146] 如图5中所示，在框502处，分析系统108可开始于定义形成给定预测模型的基础的数据集合(例如，所关注的数据)。所关注的数据可源自数个来源，例如资产102及104及数据源112，且可被存储在分析系统108的数据库中。

[0147] 所关注的数据可包含来自资产群组中的特定资产集合或来自资产群组中的所有资产(例如，所关注的资产)的历史数据。另外，所关注的数据可包含来自所关注资产中的每一者的特定传感器及/或致动器集合或来自所关注的资产中的每一者的所有传感器及/或致动器的测量值。另外，所关注的数据可包含来自过去特定时间段的数据，例如两周的历史数据。

[0148] 所关注的数据可包含各种类型的数据，这可能取决于给定的预测模型。在一些情况下，所关注的数据可至少包含指示资产的操作状况的操作数据，其中操作数据如上文在操作数据的收集段落中讨论。另外，所关注的数据可包含指示资产通常在其中被操作的环境的环境数据及/或指示在其期间资产将实行某些任务的计划日期及时间的调度数据。其它类型的数据也可被包含在所关注的数据中。

[0149] 实际上，所关注的数据可以多种方式来定义。在一个实例中，所关注的数据可为用户定义的。特定来说，用户可操作输出系统110，其接收指示所关注的某些数据的选择的用

户输入,且输出系统110可向分析系统108提供指示此类选择的数据。基于所接收的数据,分析系统108接着可定义所关注的数据。

[0150] 在另一实例中,所关注的数据可为机器定义的。特定来说,分析系统108可执行各种操作,例如模拟,以确定产生最准确的预测模型的所关注的数据。其它实例也是可能的。

[0151] 返回到图5,在框504处,分析系统108可经配置以基于所关注的数据定义与资产操作相关的聚合预测模型。通常,聚合、预测模型可定义资产的操作状况与在资产处发生事件的可能性之间的关系。具体来说,聚合、预测模型可接收来自资产的传感器的传感器数据及/或来自资产的致动器的致动器数据作为输入,且输出在未来某个时间量内将在资产处发生事件的概率。

[0152] 预测模型预测的事件可取决于特定实施情况而有所不同。例如,事件可能是故障,且因此,预测模型可能是预测在未来某个时间段内是否将会发生故障的故障模型(下文在健康状况分数模型及工作流段落中详细讨论故障模型)。在另一实例中,事件可为资产完成任务,且因此,预测模型可预测资产将按时完成任务的可能性。在其它实例中,事件可为流体或组件更换,且因此,预测模型可预测特定资产流体或组件需要被更换之前的时间量。在又其它实例中,事件可为资产生产率的变化,且因此,预测模型可预测未来特定时间段内资产的生产率。在另一实例中,事件可为“领先指示符”事件的发生,所述事件可指示与预期资产行为不同的资产行为,且因此,预测模型可预测在未来发生一或多个领先指示符事件的可能性。预测模型的其它实例也是可能的。

[0153] 无论如何,分析系统108可以各种方式定义聚合预测模型。通常,此操作可涉及利用一种或多种建模技术来产生返回介于0与1之间的概率的模型,例如随机森林技术、逻辑回归技术或其它回归技术以及其它建模技术。在特定实例实施方案中,分析系统108可根据下文参考图7的讨论来定义聚合预测模型。分析系统108也可以其它方式定义聚合模型。

[0154] 在框506处,分析系统108可经配置以定义对应于来自框504的定义模型的聚合工作流。通常,工作流可采取基于预测模型的特定输出实行的动作的形式。在实例实施方案中,工作流可包含资产基于所定义的预测模型的输出执行的一或多个操作。可为工作流的部分的操作的实例包含资产根据特定数据采集方案采集数据,根据特定数据发射方案向分析系统108发射数据,执行本地诊断工具及/或修改资产的操作状况等工作流操作。

[0155] 特定的数据采集方案可指示资产如何采集数据。特定来说,数据采集方案可指示资产从其中获得数据的某些传感器及/或致动器,例如资产的多个传感器及致动器(例如,所关注的传感器/致动器)中的传感器及/或致动器的子集。另外,数据采集方案可指示资产从所关注的传感器/致动器获得的数据的量及/或资产采集此类数据的采样频率。数据采集方案也可包含各种其它属性。在特定的实例实施方案中,特定数据采集方案可对应于资产健康状况的预测模型,且可基于资产健康状况降低而调整以(例如,从特定传感器)采集更多数据及/或特定数据。或者特定的数据采集方案可对应于领先指示符预测模型,且可基于发生领先指示符事件的可能性增加而被调整为由资产传感器及/或致动器采集的修改数据,所述领先指示符事件可表明子系统可能发生故障。

[0156] 特定的数据发射方案可指示资产如何向分析系统108发射数据。具体来说,数据发射方案可指示资产应当发射的数据类型(且还可指示数据的格式及/或结构),例如来自某些传感器或致动器的数据、资产应当发射的多个数据样本、发射频率,及/或资产应在其数

据发射中包含的数据的优先级方案。在一些情况下，特定的数据采集方案可包含数据发射方案或数据采集方案可与数据发射方案配对。在一些实例实施方案中，特定的数据发射方案可对应于资产健康状况的预测模型，且可基于高于阈值的资产健康状况而调整为不太频繁地发射数据。其它实例也是可能的。

[0157] 如上文所指示，本地诊断工具可为本地存储在资产处的程序集合等。本地诊断工具通常可促进诊断资产处的错误或故障的原因。在一些情况下，当执行时，本地诊断工具可将测试输入传递到资产的子系统或其部分中以获得测试结果，这可促进诊断错误或故障的原因。这些本地诊断工具通常在资产上休眠，且除非资产接收到特定的诊断指令，否则将不会被执行。其它本地诊断工具也是可能的。在一个实例实施方案中，特定的本地诊断工具可对应于资产的子系统的健康状况的预测模型，且可基于等于或低于阈值的子系统健康状况来执行。

[0158] 最后，工作流可涉及修改资产的操作状况。例如，可控制资产的一或多个致动器以促进修改资产的操作状况。可修改各种操作状况，例如速度、温度、压力、液面、电流消耗及功率分布等等。在特定的实例实施方案中，操作状况修改工作流可对应于用于预测资产是否将按时完成任务的预测模型，且可基于低于阈值的预测完成百分比来致使资产提高其行进速度。

[0159] 无论如何，总体工作流可以各种方式来定义。在一个实例中，聚合工作流可为用户定义的。具体来说，用户可操作接收指示某些工作流操作的选择的用户输入的计算装置，且计算装置可向分析系统108提供指示此类选择的数据。基于此数据，分析系统108可接着定义聚合工作流。

[0160] 在另一实例中，聚合工作流可为机器定义的。特定来说，分析系统108可执行各种操作(例如模拟)来确定可促进确定由预测模型输出的概率的原因及/或防止由模型预测的事件的发生的工作流。定义聚合工作流的其它实例也是可能的。

[0161] 在定义对应于预测模型的工作流时，分析系统108可定义工作流的触发。在实例实施方案中，工作流触发可为由预测模型输出的概率的值或由预测模型输出的值的范围。在一些情况下，工作流可具有多个触发，其中的每一触发可致使发生不同的操作或多个操作。

[0162] 为了说明，图6A是聚合模型-工作流对600的概念说明。如所示，聚合模型-工作流对说明600包含用于模型输入602、模型计算604、模型输出范围606及对应工作流操作608的列。在此实例中，预测模型具有来自传感器A的单个输入数据，且具有两个计算值：计算值I及II。此预测模型的输出影响所执行的工作流操作。如果输出概率小于或等于80%，那么执行工作流操作1。否则，执行工作流操作2。其它实例模型-工作流对在本文中是可能的且予以考虑。

[0163] 2.个性化模型及工作流

[0164] 在另一方面中，分析系统108可经配置以定义资产的个性化的预测模型及/或工作流，这可涉及利用聚合模型-工作流对作为基准。个性化可基于资产的某些特性。以此方式，分析系统108可对给定资产提供与聚合模型-工作流对相比更准确及稳健的模型-工作流对。

[0165] 特定来说，返回到图5，在框508处，分析系统108可经配置以决定是否个性化在框504处针对给定资产(例如，资产102)定义的聚合模型。分析系统108可以多种方式实行此决

定。

[0166] 在一些情况下,分析系统108可经配置以默认地定义个性化的预测模型。在其它情况下,分析系统108可经配置以基于资产102的某些特性来决定是否定义个性化预测模型。例如,在一些情况下,只有某些类型或类别或在特定环境中操作的或具有某些健康状况评分的资产可能会收到个性化的预测模型。在又其它情况下,用户可定义是否针对资产102定义个性化模型。其它实例也是可能的。

[0167] 无论如何,如果分析系统108决定定义资产102的个性化预测模型,那么分析系统108可在框510处这样做。否则,分析系统108可前进到框512。

[0168] 在框510处,分析系统108可经配置以按照多种方式定义个性化预测模型。在实例实施方案中,分析系统108可至少部分地基于资产102的一或多个特性来定义个性化预测模型。

[0169] 在定义资产102的个性化预测模型之前,分析系统108可能已经确定了形成个性化模型的基础的一或多个所关注的资产特性。实际上,不同的预测模型可能具有不同的对应的所关注的特性。

[0170] 通常,所关注的特性可为与聚合模型-工作流对相关的特性。例如,所关注的特性可为分析系统108已经确定影响聚合模型-工作流对的准确性的特性。此类特性的实例可包含资产年限、资产使用状况、资产能力、资产负荷、资产健康状况(可能由下文讨论的资产健康状况指标指示)、资产类别(例如,品牌及/或型号)及操作资产的环境等特性。

[0171] 分析系统108可能已经以多种方式确定了所关注的特性。在一个实例中,分析系统108可通过执行促进识别所关注的特性的一或多个建模模拟来完成。在另一实例中,所关注的特性可能已经预定义并存储在分析系统108的数据存储装置中。在又另一实例中,所关注的特性可能已经由用户定义且经由输出系统110被提供给分析系统108。其它实例也是可能的。

[0172] 无论如何,在确定所关注的特性之后,分析系统108可确定与所确定的所关注的特性对应的资产102的特性。即,分析系统108可确定与所关注的特性对应的资产102的特性的类型、价值、其存在或缺乏等。分析系统108可以多种方式执行此操作。

[0173] 例如,分析系统108可经配置以基于源自资产102及/或数据源112的数据来执行此操作。特定来说,分析系统108可利用资产102的操作数据及/或来自数据源112的外部数据来确定资产102的一或多个特性。其它实例也是可能的。

[0174] 基于所确定的资产102的一或多个特性,分析系统108可通过修改聚合模型来定义个性化预测模型。聚合模型可以多种方式进行修改。例如,可通过改变(例如,增加、移除、重新排序等)一或多个模型输入、改变与资产操作极限对应的一或多个传感器及/或致动器测量范围(例如,改变与“领先指示符”事件对应的操作极限)、改变一或多个模型计算值、对计算的变量或输出加权(或改变其权重)、利用与用于定义聚合模型的建模技术不同的建模技术及/或利用与用于定义聚合模型的响应变量不同的响应变量等等来修改聚合模型。

[0175] 为了说明,图6B是个性化模型-工作流对610的概念说明。-具体来说,个性化模型-工作流对说明610是来自图6A的聚合模型-工作流对的修改版本。如所示,个性化模型-工作流对说明610包含用于模型输入612及模型计算614的修改列,且包含来自图6A的模型输出范围606及工作流操作608的原始列。在此实例中,个性化模型有两个输入,来自传感器A及

致动器B的数据，且具有两个计算值：计算值II及III。输出范围及对应工作流操作与图6A的输出范围及对应工作流操作相同。分析系统108可能已经基于确定资产102例如相对较老且健康状况相对较差等原因来以此方式定义个性化模型。

[0176] 实际上，个性化聚合模型可取决于给定资产的一或多个特性。特定来说，某些特性可能会以与其它特性不同的方式影响聚合模型的修改。另外，特性的类型、价值、存在等也可影响修改。例如，资产年限可能会影响聚合模型的第一部分，而资产类别可能会影响聚合模型的第二不同部分。且在第一年限范围内的资产年限可能以第一方式影响聚合模型的第一部分，而在与第一范围不同的第二年限范围内的资产年限可能以第二方式影响聚合模型的第一部分。其它实例也是可能的。

[0177] 在一些实施方案中，个性化聚合模型可取决于作为资产特性的补充或替代的考虑。例如，当已知资产处于相对良好的操作状态（例如，如机械师等所定义的）时，可基于资产的传感器及/或致动器读数对聚合模型进行个性化。更特定来说，在领先指示符预测模型的实例中，分析系统108可经配置以接收资产处于良好操作状态的指示（例如，来自机械师操作的计算装置）以及操作来自资产的数据。至少基于操作数据，分析系统108接着可通过修改与“领先指示符”事件的对应的相应操作极限来个性化资产的领先指示符预测模型。其它实例也是可能的。

[0178] 返回到图5，在框512处，分析系统108还可经配置以决定是否个性化资产102的工作流。分析系统108可以多种方式实行此决定。在一些实施方案中，分析系统108可根据框508来执行此操作。在其它实施方案中，分析系统108可基于个性化预测模型来决定是否定义个性化工作流。在又另一实施方案中，如果定义了个性化预测模型，那么分析系统108可决定定义个性化工作流。其它实例也是可能的。

[0179] 无论如何，如果分析系统108决定定义资产102的个性化工作流，那么分析系统108可在框514处这样做。否则，分析系统108可结束定义阶段。

[0180] 在框514处，分析系统108可经配置以按照多种方式定义个性化工作流。在实例实施方案中，分析系统108可至少部分地基于资产102的一或多个特性来定义个性化工作流。

[0181] 在定义资产102的个性化工作流之前，类似于定义个性化预测模型，分析系统108可能已经确定了形成个性化工作流的基础的一或多个所关注的资产特性，其可能已经根据框510的讨论进行了确定。通常，这些所关注的特性可为影响聚合工作流的效力的特性。此类特性可包含上文讨论的实例特性中的任一者。其它特性也是可能的。

[0182] 再次类似于框510，分析系统108可确定与个性化工作流的所确定的所关注的特性对应的资产102的特性。在实例实施方案中，分析法系统108可以与参考框510所讨论的特性确定类似的方式来确定资产102的特性，且实际上可利用所述确定中的一些或全部。

[0183] 无论如何，基于资产102的所确定的一或多个特性，分析系统108可通过修改聚合工作流来个性化资产102的工作流。聚合工作流可以多种方式进行修改。例如，可通过改变（例如，增加、移除、重新排序、替换等）一或多个工作流操作（例如，从第一数据采集方案改变为第二方案或从特定数据采集方案改变为特定的本地诊断工具）及/或改变（例如，增加、减少、增加、移除等）触发特定工作流操作的对应模型输出值或值范围等等来修改聚合工作流。实际上，对聚合工作流的修改可以类似于对聚合模型的修改的方式取决于资产102的一或多个特性。

[0184]　为了说明，图6C是个性化模型-工作流对620的概念说明。具体来说，个性化模型-工作流对说明620是来自图6A的聚合模型-工作流对的修改版本。如所示，个性化模型-工作流对说明620包含来自图6A的模型输入602、模型计算604及模型输出范围606的原始列，但是包含用于工作流操作628的修改列。在此实例中，个性化模型-工作流对与图6A中的聚合模型-工作流对类似，除了当聚合模型的输出大于80％时，工作流操作3被触发而不是操作1之外。除其它原因外，分析系统108可基于确定资产102例如在历史上增加资产故障的发生的环境中操作来定义此个别工作流。

[0185]　在定义个性化工作流之后，分析系统108可结束定义阶段。那时，分析系统108可接着具有用于资产102的个性化模型-工作流对。

[0186]　在一些实例实施方案中，分析系统108可经配置以定义用于给定资产的个性化预测模型及/或对应工作流，而不首先定义聚合预测模型及/或对应工作流。其它实例也是可能的。

[0187]　虽然上文讨论了分析系统108个性化预测模型及/或工作流，但是其它装置及/或系统可执行个性化。例如，资产102的本地分析装置可个性化预测模型及/或工作流，或可与分析系统108一起工作以执行此类操作。下文进一步详细讨论本地分析装置执行此类操作。

[0188]　3.健康状况分数模型及工作流

[0189]　在特定实施方案中，如上所提及，分析系统108可经配置以定义与资产的健康状况相关联的预测模型及对应工作流。在实例实施方案中，用于监测资产的健康状况的一或多个预测模型可用于输出资产的健康状况指标(例如，“健康状况分数”)，所述健康状况指标是单个聚合指标，其指示在未来的给定时间范围内(例如，接下来的两周)是否将在给定资产处发生故障。特定来说，健康状况指标可指示在未来的给定时间范围内将不会在资产处发生故障群组中的任何故障的可能性，或健康状况指标可指示在未来的给定时间范围内将在资产处发生故障群组中的至少一个故障的可能性。

[0190]　实际上，用于输出健康状况指标的预测模型及对应工作流可根据上述讨论被定义为聚合或个性化模型及/或工作流。

[0191]　另外，取决于健康状况指标的期望粒度，分析系统108可经配置以定义输出不同水平的健康状况指标的不同的预测模型且定义不同的对应工作流。例如，分析系统108可定义输出整个资产的健康状况指标(即，资产水平健康状况指标)的预测模型。作为另一实例，分析系统108可定义输出资产的一或多个子系统的相应健康状况指标(即，子系统级健康状况指标)的相应预测模型。在一些情况下，每一子系统级预测模型的输出可经组合以产生资产水平健康状况指标。其它实例也是可能的。

[0192]　通常，定义输出健康状况指标的预测模型可以各种方式执行。图7是描绘可用于定义输出健康状况指标的模型的建模阶段的一个可能实例的流程图700。出于说明目的，实例建模阶段被描述为由分析系统108实行，但是此建模阶段也可由其它系统实行。所属领域的一般技术人员将明白，出于清楚及解释目的提供流程图700，且可利用操作的数个其它组合来确定健康状况指标。

[0193]　如图7中所示，在框702处，分析系统108可通过定义形成健康状况指标的基础的一或多个故障(即，所关注的故障)的集合而开始。实际上，一或多个故障可为在发生的情况下可能使资产(或其子系统)不可操作的故障。基于所定义的故障集合，分析系统108可采取步

骤来定义用于预测在未来的给定时间范围内(例如,接下来的两周)发生故障中的任一者的可能性的模型。

[0194] 特定来说,在框704处,分析系统108可分析一或多个资产群组的历史操作数据,以从所述故障集合中识别给定故障的过去发生。在框706处,分析系统108可识别与给定故障的每一识别的过去发生相关联的相应操作数据集合(例如,在发生给定故障之前,来自给定时间范围的传感器及/或致动器数据)。在框708处,分析系统108可分析与给定故障的过去发生相关联的所识别的操作数据集合,以定义(1)用于给定的操作指标集合的值与(2)在未来的给定时间范围内(例如,未来两周)发生故障的可能性之间的关系(例如,故障模型)。最后,在框710处,将所定义集合中的每一故障的所定义关系(例如,个别故障模型)接着可组合成用于预测故障发生的总体可能性的模型。

[0195] 当分析系统108继续接收一或多个资产群组的更新的操作数据时,分析系统108还可通过对更新的操作数据重复步骤704到710来继续针对一或多个故障的所定义集合改进预测模型。

[0196] 现在将进一步详细地描述图7中所说明的实例建模阶段的功能。从框702开始,如上所述,分析系统108可通过定义形成健康状况指标的基础的一或多个故障的集合开始。分析系统108可以各种方式执行此功能。

[0197] 在一个实例中,一或多个故障的集合可基于一或多个用户输入。具体来说,分析系统108可从由用户操作的计算系统(例如输出系统110)接收指示一或多个故障的用户选择的输入数据。因而,一或多个故障的集合可为用户定义的。

[0198] 在其它实例中,一或多个故障的集合可基于由分析系统108做出的确定(例如,机器定义的)。特定来说,分析系统108可经配置以定义可能以多种方式发生的一或多个故障的集合。

[0199] 例如,分析系统108可经配置以基于资产102的一或多个特性来定义故障集合。即,某些故障可能对应于资产的某些特性,例如资产类型、类别等。例如,每一类型及/或类别的资产可能具有相应的所关注的故障。

[0200] 在另一实例中,分析系统108可经配置以基于存储在分析系统108的数据库中的历史数据及/或由数据源112提供的外部数据来定义故障集合。例如,分析系统108可利用此数据来确定哪些故障导致了最长的维修时间及/或历史上哪些故障之后接着出现额外故障等等。

[0201] 在又其它实例中,可基于用户输入及由分析系统108做出的确定的组合来定义一或多个故障的集合。其它实例也是可能的。

[0202] 在框704处,对于故障集合中的故障中的每一者,分析系统108可分析一或多个资产群组的历史操作数据(例如,异常行为数据)以识别给定故障的过去发生。一或多个资产群组可包含单个资产(例如资产102),或包含相同或类似类型的多个资产,例如包含资产102及104的资产群的队伍。分析系统108可分析特定数量的历史操作数据,例如一定量的时间价值的数据(例如,一个月的数据价值)或一定数量的数据点(例如,最近的一千个数据点)等等。

[0203] 实际上,识别给定故障的过去发生可涉及分析系统108识别指示给定故障的操作数据(例如异常状况数据)的类型。通常,给定故障可与一或多个异常状况指示符(例如故障

代码)相关联。即,当给定故障发生时,可能触发一或多个异常状态指示符。因而,异常状况指示符可反映给定故障的潜在征兆。

[0204] 在识别指示给定故障的操作数据的类型之后,分析系统108可以多种方式识别给定故障的过去发生。例如,分析系统108可根据存储在分析系统108的数据库中的历史操作数据来定位与关联于给定故障的异常状况指示符对应的异常状况数据。每一经定位的异常状况数据将指示发生给定的故障。基于此经定位的异常状况数据,分析系统108可识别发生过去故障的时间。

[0205] 在框706处,分析系统108可识别与每一识别的给定故障的过去发生相关联的操作数据的相应集合。特定来说,分析系统108可识别来自在给定故障的给定发生的时间周围的某个时间范围内的传感器及/或致动器数据的集合。例如,所述数据集合可来自在给定的故障发生之前、之后或周围的特定时间范围(例如,两周)。在其它情况下,可从在给定的故障发生之前、之后或周围的一定数量的数据点来识别所述数据集合。

[0206] 在实例实施方案中,所述操作数据集合可包含来自资产102的一些或全部传感器及致动器的传感器及/或致动器数据。例如,所述操作数据集合可包含来自与对应于给定故障的异常状况指示符相关联的传感器及/或致动器的数据。

[0207] 为了说明,图8描绘了分析系统108可分析以促进定义模型的历史操作数据的概念说明。曲线图800可对应于源自资产102的传感器及致动器中的一些(例如,传感器A及致动器B)或所有者的历史数据片段。如所示,曲线图800包含x轴802上的时间、y轴804上的测量值及与传感器A对应的传感器数据806及与致动器B对应的致动器数据808,所述数据中的每一者包含表示特定时间点T_i处的测量值的各种数据点。另外,曲线图800包含在过去时间T_f(例如,“故障时间”)发生的故障810的发生的指示及发生故障之前的时间量812的指示ΔT,从所述指示识别操作数据集合。因而,$T_f-\Delta T$定义了所关注数据点的时间范围814。

[0208] 返回到图7,在分析系统108识别给定故障的给定发生(例如,T_f处的发生)的操作数据集合之后,分析系统108可确定是否存在应针对其识别操作数据集合的任何剩余发生。在存在剩余发生的情况下,将对每一剩余发生重复框706。

[0209] 此后,在框708处,分析系统108可分析与给定故障的过去发生相关联的所识别的操作数据集合,以定义(1)操作指标的给定集合(例如,传感器及/或致动器测量值的给定集合)与(2)在未来的给定时间范围内(例如,未来两周)发生故障的可能性之间的关系(例如,故障模型)。即,给定故障模型可将来自一或多个传感器及/或致动器的传感器及/或致动器测量值作为输入,并输出将在未来的给定时间范围内发生给定故障的概率。

[0210] 通常,故障模型可定义资产102的操作状况与发生故障的可能性之间的关系。在一些实施方案中,除了来自资产102的传感器及/或致动器的原始数据信号之外,故障模型可接收从传感器及/或致动器信号导出的多个其它数据输入,也被称为特征。此类特征可包含发生故障时在历史上测量的值的平均值或范围、发生故障之前在历史上测量的值梯度的平均值或范围(例如,测量值的变化率)、特征之间的持续时间(例如,在第一次发生故障与第二次发生故障之间的时间量或数据点数目)及/或指示发生故障周围的传感器及/或致动器测量值趋势的一或多个故障模式。所属领域的一般技术人员将明白,这些仅仅是可从传感器及/或致动器信号导出的几个实例特征,且许多其它特征是可能的。

[0211] 实际上,故障模型可以多种方式来定义。在实例实施方案中,分析系统108可通过

利用一或多种建模技术来定义故障模型，所述建模技术返回介于0与1之间的概率，所述建模技术可采用上述任何建模技术的形式。

[0212] 在特定实例中，定义故障模型可涉及分析系统108基于在框706识别的历史操作数据来产生响应变量。具体来说，分析系统108可确定在特定时间点所接收的传感器及/或致动器测量值的每一集合的相关联响应变量。因而，响应变量可采取与故障模型相关联的数据集合的形式。

[0213] 响应变量可指示给定测量值集合是否在框706处确定的时间范围中的任一者内。即，响应变量可反映给定数据集合是否来自发生故障周围的所关注时间。响应变量可为二进制值响应变量，使得如果给定测量值集合在所确定的时间范围中的任一者内，那么相关联响应变量被指派值1，否则相关联响应变量被指派值0。

[0214] 返回到图8，在曲线图800上展示了响应变量矢量Y_{res}的概念说明。如所示，与时间范围814内的测量值集合相关联的响应变量具有值1(例如，在时间T_{i+3}到T_{i+8}处的Y_{res})，而与时间范围814之外的测量值集合相关联的响应变量具有值0(例如，在时间T_i到T_{i+2}及T_{i+9}到T_{i+10}处的Y_{res})。其它响应变量也是可能的。

[0215] 在基于响应变量定义故障模型的特定实例中继续，分析系统108可利用在框706识别的历史操作数据及所产生的响应变量来训练故障模型。基于此训练过程，分析系统108接着可定义故障模型，其接收各种传感器及/或致动器数据作为输入，且输出在与用于产生响应变量等效的时间段内将发生故障的介于0与1之间的概率。

[0216] 在一些情况下，利用在框706处识别的历史操作数据及所产生的响应变量的训练可导致用于每一传感器及/或致动器的可变重要性统计量。给定可变重要性统计量可指示传感器或致动器对将在未来的时间段内发生给定故障的概率的相对影响。

[0217] 另外地或替代地，分析系统108可经配置以基于一或多种生存分析技术(例如Cox比例风险技术)来定义故障模型。分析系统108可以在某些方面类似于上文讨论的建模技术的方式来利用生存分析技术，但是分析系统108可确定生存时间响应变量，其指示从上次故障到下一个预期事件的时间量。下一个预期事件可能是接收传感器及/或致动器的测量值或故障的发生，以先发生者为准。此响应变量可包含与接收测量值的特定时间点中的每一者相关联的一对值。接着，可利用响应变量来确定在未来的给定时间范围内将发生故障的概率。

[0218] 在一些实例实施方案中，可部分地基于例如天气数据及“热箱”数据等外部数据以及其它数据定义故障模型。例如，基于此数据，故障模型可增大或减小输出故障概率。

[0219] 实际上，可在与资产传感器及/或致动器获得测量值的时间不一致的时间点处观察外部数据。例如，收集“热箱”数据的时间(例如，机车沿装有热箱传感器的铁路轨道段行进的时间)可与传感器及/或致动器测量时间不一致。在此类情况下，分析系统108可经配置以执行一或多个操作以确定原本将在与传感器测量时间对应的时间处观察到的外部数据观察值。

[0220] 具体来说，分析系统108可利用外部数据观察值的时间及测量值的时间来内插外部数据观察值以产生与测量时间对应的时间的外部数据值。外部数据的内插可允许将外部数据观察值或从其导出的特征作为输入包含在故障模型中。实际上，除了其它实例外，还可使用各种技术来利用传感器及/或致动器数据来内插外部数据，例如最近邻内插、线性内

插、多项式内插及样条内插。

[0221] 返回到图7，在分析系统108确定来自在框702处定义的故障集合中的给定故障的故障模型之后，分析系统108可确定是否存在应针对其确定故障模型的任何剩余故障。在仍然存在应针对其确定故障模型的故障情况下，分析系统108可重复框704到708的回路。在一些实施方案中，分析系统108可确定包含在框702处定义的故障的所有者的单个故障模型。在其它实施方案中，分析系统108可确定资产102的每一子系统的故障模型，接着可利用所述故障模型来确定资产水平故障模型。其它实例也是可能的。

[0222] 最后，在框710处，接着可将所定义集合中的每一故障的所定义关系(例如，个别故障模型)组合成用于预测在未来的给定时间范围(例如，接下来的两周)发生故障的总体可能性的模型(例如，健康状况指标模型)。即，所述模型可接收来自一或多个传感器及/或致动器的传感器及/或致动器测量值作为输入，并输出将在未来的给定时间范围内发生故障集合中的至少一个故障的单个概率。

[0223] 分析系统108可以多种方式定义健康状况指标模型，这可取决于健康状况指标的期望粒度。即，在存在多个故障模型的情况下，可以多种方式利用故障模型的结果来获得健康状况指标模型的输出。例如，分析系统108可从多个故障模型中确定最大值、中值或平均值，并将所述确定值用作健康状况指标模型的输出。

[0224] 在其它实例中，确定健康状况指标模型可涉及分析系统108将权重归属于由个别故障模型输出的个别概率。例如，来自故障集合的每一故障可被认为是同样不希望的，因此每一概率同样可在确定健康状况指标模型时进行相同加权。在其它情况下，一些故障可能被认为比其它故障更不受欢迎(例如，更具灾难性或需要更长的维修时间等)，因此那些对应的概率可能比其它概率进行更多加权。

[0225] 在又其它实例中，确定健康状况指标模型可涉及分析系统108利用一种或多种建模技术，例如回归技术。聚合响应变量可采取来自个别故障模型中的每一者的响应变量(例如，图8中的Y_{res})的逻辑分离(逻辑或)的形式。例如，与在框706确定的任何时间范围(例如，图8的时间范围814)内发生的任何测量值集合相关联的聚合响应变量可具有值1，而与在时间范围中的任一者之外发生的测量值集合相关联的聚合响应变量可具有零值。定义健康状况指标模型的其它方式也是可能的。

[0226] 在一些实施方案中，框710可能是不必要的。例如，如上文所讨论，分析系统108可确定单个故障模型，在这种情况下，健康状况指标模型可为单个故障模型。

[0227] 实际上，分析系统108可经配置以更新单个故障模型及/或总体健康状况指标模型。分析系统108可每天、每周、每月等更新模型，且可基于来自资产102或来自其它资产(例如，来自与资产102相同的群的其它资产)的历史操作数据的新部分。其它实例也是可能的。

[0228] C.部署模型及工作流

[0229] 在分析系统108定义模型-工作流对之后，分析系统108可将定义的模型-工作流对部署到一或多个资产。具体来说，分析系统108可向至少一个资产(例如资产102)发射定义的预测模型及/或对应工作流。分析系统108可周期性地或基于触发事件(例如对给定模型-工作流对的任何修改或更新)来发射模型-工作流对。

[0230] 在一些情况下，分析系统108可仅发射个性化模型或个性化工作流中的一者。例如，在分析系统108仅定义个性化模型或工作流的情况下，分析系统108可发射工作流或模

型的聚合版本以及个性化模型或工作流,或如果资产102已将聚合版本存储在数据存储装置中,那么分析系统108可能不需要发射聚合版本。总之,分析系统108可发射(1)个性化模型及/或个性化工作流、(2)个性化模型及聚合工作流、(3)聚合模型及个性化工作流或(4)聚合模型及聚合工作流。

[0231] 实际上,分析系统108可能已经针对多个资产执行了图7的框702到710的操作中的一些或全部来定义每一资产的模型-工作流对。例如,分析系统108可另外定义资产104的模型-工作流对。分析系统108可经配置以同时或按顺序向资产102及104发射相应的模型-工作流对。

[0232] D.资产的本地执行

[0233] 例如资产102等给定资产可经配置以接收模型-工作流对或其部分,并根据所接收的模型-工作流对来操作。即,资产102可在数据存储装置中存储模型-工作流对,并将由资产102的传感器及/或致动器获得的数据输入到预测模型中,且有时基于预测模型的输出执行对应工作流。

[0234] 实际上,资产102的各种组件可执行预测模型及/或对应工作流。例如,如上文所讨论,每一资产可包含本地分析装置,其经配置以存储及运行由分析系统108提供的模型-工作流对。当本地分析装置接收到特定传感器及/或致动器数据时,其可将所接收的数据输入到预测模型中,且取决于模型的输出可执行对应工作流的一或多个操作。

[0235] 在另一实例中,与本地分析装置分离的资产102的中央处理单元可执行预测模型及/或对应工作流。在又其它实例中,资产102的本地分析装置及中央处理单元可协同执行模型-工作流对。例如,本地分析装置可执行预测模型,且中央处理单元可执行工作流,反之亦然。

[0236] 在实例实施方案中,在本地执行模型-工作流对之前(或可能在首先本地执行模型工作流时),本地分析装置可个性化资产102的预测模型及/或对应工作流。无论模型-工作流对是采用聚合模型-工作流对还是个性化模型-工作流对的形式,均可能发生这种情况。

[0237] 如上文所指示,分析系统108可基于关于资产群组或特定资产的某些预测、假设及/或一般化来定义模型-工作流对。例如,在定义模型-工作流对时,分析系统108可预测、假设及/或一般化资产的有关特性及/或资产的操作状况以及其它考虑。

[0238] 无论如何,本地分析装置个性化预测模型及/或对应工作流可涉及本地分析装置确认或推翻分析系统108在定义模型-工作流对时作出的一或多个预测、假设及/或一般化中的一或多者。根据本地分析装置对预测、假设及/或一般化的评估,本地分析装置此后可修改(或在已经个性化的模型及/或工作流的情况下进一步修改)预测模型及/或工作流。以此方式,本地分析装置可帮助定义更加现实及/或准确的模型-工作流对,这可能导致更有效的资产监测。

[0239] 实际上,本地分析装置可基于许多考虑来个性化预测模型及/或工作流。例如,本地分析装置可基于由资产102的一或多个传感器及/或致动器产生的操作数据来这样做。具体来说,本地分析装置可通过以下项来个性化:(1)获得由一或多个传感器及/或致动器的特定群组产生的操作数据(例如,通过经由资产的中央处理单元间接地获得此数据,或可能直接从传感器及/或致动器本身中的某些直接获得此数据;(2)基于所获得的操作数据评估与模型-工作流对相关联的一或多个预测、假设及/或一般化;以及(3)如果所述评估指示任

何预测、假设及/或一般化是不正确的，那么相应地修改模型及/或工作流。此类操作可以各种方式执行。

[0240] 在一个实例中，本地分析装置(例如，经由资产的中央处理单元)获得由传感器及/或致动器的特定群组产生的操作数据可基于作为模型-工作流对的部分或与其一起被包含的指令。特定来说，指令可识别本地分析装置执行的一或多个测试，所述测试评估定义模型-工作流对时涉及到的一些或全部预测、假设及/或一般化。每一测试可识别本地分析装置将针对其获得操作数据的一或多个所关注的所关注的传感器及/或致动器、要获得的操作数据的量及/或其它测试考虑。因此，本地分析装置获得由传感器及/或致动器的特定群组产生的操作数据可能涉及本地分析装置根据测试指令等获得此操作数据。本地分析装置获得用于个性化模型-工作流对的操作数据的其它实例也是可能的。

[0241] 如上所述，在获得操作数据之后，本地分析装置可利用数据来评估定义模型-工作流对时涉及到的一些或全部预测、假设及/或一般化。此操作可以各种方式执行。在一个实例中，本地分析装置可将所获得的操作数据与一或多个阈值(例如，阈值及/或值的阈值范围)进行比较。通常，给定的阈值或范围可对应于用于定义模型-工作流对的一或多个预测、假设及/或一般化。具体来说，测试指令中识别的每一传感器或致动器(或传感器及/或致动器的组合)可具有对应的阈值或范围。本地分析装置接着可确定由给定的传感器或致动器产生的操作数据是高于还是低于对应的阈值或范围。本地分析装置评估预测、假设及/或一般化的其它实例也是可能的。

[0242] 此后，本地分析装置可基于评估来修改(或不修改)预测模型及/或工作流。即，如果评估指示任何预测、假设及/或一般化是不正确的，那么本地分析装置可相应地修改预测模型及/或工作流。否则，本地分析装置可在不修改的情况下执行模型-工作流对。

[0243] 实际上，本地分析装置可以多种方式来修改预测模型及/或工作流。例如，本地分析装置可(例如，通过修改值的值或范围)修改预测模型及/或工作流的一或多个参数及/或预测模型及/或工作流的触发点等等。

[0244] 作为一个非限制性实例，假设资产102的发动机操作温度不超过特定温度，分析系统108可能已经定义了资产102的模型-工作流对。结果，资产102的预测模型的部分可涉及确定第一计算值且接着仅在第一计算值超过基于假设的发动机操作温度确定的阈值时才确定第二计算值。当个性化模型-工作流对时，本地分析装置可获得由测量资产102的发动机的操作数据的一或多个传感器及/或致动器产生的数据。接着，本地分析装置可使用此数据来确定关于发动机操作温度的假设实际上是否为真(例如，发动机操作温度是否超过阈值)。如果数据指示发动机操作温度具有超过假设的特定温度的值或超过所述特定温度阈值量，那么本地分析装置可例如修改触发确定第二计算值的阈值。本地分析装置个性化预测模型及/或工作流的其它实例也是可能的。

[0245] 本地分析装置可基于额外或替代考虑来个性化模型-工作流对。例如，本地分析装置可基于一或多个资产特性(例如上文讨论的资产特性中的任一者)来这样做，所述特性可由本地分析装置确定或提供给本地分析装置。其它实例也是可能的。

[0246] 在实例实施方案中，在本地分析装置个性化预测模型及/或工作流之后，本地分析装置可向分析系统108提供预测模型及/或工作流已经个性化的指示。此指示可采取各种形式。例如，所述指示可识别本地分析装置修改的预测模型及/或工作流的方面或部分(例如，

被修改的参数及/或参数被修改为什么)及/或可识别修改的原因(例如,致使本地分析装置进行修改的基础操作数据或其它资产数据及/或原因的描述)。其它实例也是可能的。

[0247] 在一些实例实施方案中,本地分析装置及分析系统108两者均可在个性化模型-工作流对被涉及到,个性化模型-工作流对可以各种方式执行。例如,分析系统108可向本地分析装置提供测试资产102的某些状况及/或特性的指令。基于所述指令,本地分析装置可在资产102处执行测试。例如,本地分析装置可获得由特定资产传感器及/或致动器产生的操作数据。此后,本地分析装置可向分析系统108提供来自测试状况的结果。基于此类结果,分析系统108可相应地定义资产102的预测模型及/或工作流并将其发射到本地分析装置以用于本地执行。

[0248] 在其它实例中,本地分析装置可执行与执行工作流的部分相同或类似的测试操作。即,与预测模型对应的特定工作流可致使本地分析装置执行某些测试并向分析系统108发射结果。

[0249] 在实例实施方案中,在本地分析装置个性化预测模型及/或工作流(或与分析系统108一起工作以个性化预测模型及/或工作流)之后,本地分析装置可执行个性化预测模型及/或工作流而不是原始模型及/或工作流(例如,本地分析装置最初从分析系统108接收的模型及/或工作流)。在一些情况下,虽然本地分析装置执行个性化版本,但是本地分析装置可将模型及/或工作流的原始版本保留在数据存储装置中。

[0250] 通常,资产执行预测模型且基于所得输出执行工作流的操作可促进确定由模型输出的特定事件发生的可能性的原因及/或可促进防止未来发生特定事件。在执行工作流时,资产可在本地确定并采取行动来帮助防止事件发生,这在依赖分析系统108做出此类确定并提供推荐的动作是无效或的不可行(例如,当存在网络延迟时、当网络连接不良时、当资产移出通信网络106的覆盖范围时等)的情况下是有益的。

[0251] 实际上,资产可以各种方式执行预测模型,这可能取决于特定预测模型。图9是描绘可用于本地执行预测模型的本地执行阶段的一个可能实例的流程图900。将在输出资产的健康状况指标的健康状况指标模型的背景中讨论实例本地执行阶段,但是应当理解的是,相同或类似的本地执行阶段可用于其它类型的预测模型。另外,出于说明目的,实例本地执行阶段被描述为由资产102的本地分析装置实行,但是此阶段也可由其它装置及/或系统实行。所属领域的一般技术人员将明白,出于清楚及解释目的提供流程图900,且可利用操作及功能的数个其它组合来本地执行预测模型。

[0252] 如图9中所示,在框902处,本地分析装置可接收反映资产102的当前操作状况的数据。在框904处,本地分析装置可从所接收的数据识别要被输入到由分析系统108提供的模型中的操作数据集合。在框906处,本地分析装置可接着将所识别的操作数据集合输入到模型中并运行模型以获得资产102的健康状况指标。

[0253] 当本地分析装置继续接收资产102的更新的操作数据时,本地分析装置还可通过基于更新的操作数据重复框902到906的操作而继续更新资产102的健康状况指标。在一些情况下,可在每次本地分析装置从资产102的传感器及/或致动器接收新数据时或周期性地(例如,每小时、每天、每周,每月等)重复框902到906的操作。以此方式,当在操作中使用资产时,本地分析装置可经配置以动态地更新健康状况指标,可能实时更新。

[0254] 现在将进一步详细地描述图9中所说明的实例本地执行阶段的功能。在框902处,

本地分析装置可接收反映资产102的当前操作状况的数据。此数据可包含来自资产102的传感器中的一或多者的传感器数据、来自资产102的一或多个致动器的致动器数据，及/或其可包含异常状况数据以及其它类型的数据。

[0255] 在框904处，本地分析装置可从所接收的数据识别要被输入到由分析系统108提供的健康状况指标模型中的操作数据集合。此操作可以多种方式执行。

[0256] 在一个实例中，本地分析装置可基于资产102的特性(例如正针对其确定健康状况指标的资产类型或资产类别)识别用于所述模型的操作数据输入集合(例如，来自特定传感器及/或所关注致动器的数据)。在一些情况下，所识别的操作数据输入集合可为来自资产102的传感器中一些或所有者的传感器数据及/或来自资产102的致动器中的一些或所有者的的致动器数据。

[0257] 在另一实例中，本地分析装置可基于由分析系统108提供的预测模型来识别操作数据输入集合。即，分析系统108可向资产102提供用于模型的特定输入的一些指示(例如，在预测模型中或在单独的数据发射中)。识别操作数据输入集合的其它实例也是可能的。

[0258] 在框906处，本地分析装置操作接着可运行健康状况指标模型。具体来说，本地分析装置可将所识别的操作数据集合输入到模型中，所述模型进而确定并输出在未来的给定时间范围内(例如，接下来的两周)发生至少一个故障的总体可能性。

[0259] 在一些实施方案中，此操作可涉及本地分析装置将特定操作数据(例如，传感器及/或致动器数据)输入到健康状况指标模型的一或多个个别故障模型中，每一个别故障模型可输出个别概率。接着，本地分析装置可使用这些个别概率，可能根据健康状况指标模型将一些概率加权得多于其它概率，以确定在未来的给定时间范围内发生故障的总体可能性。

[0260] 在确定发生故障的总体可能性之后，本地分析装置可将发生故障的概率转换成健康状况指标，所述健康状况指标可采取反映在未来的时间范围内(例如，两周)内将不会在资产102处发生故障的可能性的单个聚合参数的形式。在实例实施方案中，将故障概率转换成健康状况指标可涉及本地分析装置确定故障概率的补。具体来说，总体故障概率可采取从0到1的值的形式；健康状况指标可通过从1减去所述数字来确定。将故障概率转换成健康状况指标的其它实例也是可能的。

[0261] 在资产本地执行预测模型之后，资产接着可基于所执行的预测模型的所得输出来执行对应工作流。通常，资产执行工作流可涉及本地分析装置致使在资产处执行操作(例如，通过向资产的机载系统中的一或多者发送指令)及/或本地分析装置致使例如分析系统108及/或输出系统110等计算系统执行远离资产的操作。如上所提及，工作流可采取各种形式，因此工作流可以各种方式执行。

[0262] 例如，可使资产102在内部执行修改资产102的一些行为的一或多个操作，例如修改数据采集及/或发射方案、执行本地诊断工具、修改资产102的操作状况(例如，修改速度、加速度、风扇速度、螺旋桨角度、进气口等或经由资产102的一或多个致动器执行其它机械操作)，或输出可能是健康状况相对较低指标或推荐的预防性动作的指示，所述预防性动作应当在资产102的用户接口处对资产102执行或对外部计算系统执行。

[0263] 在另一实例中，资产102可向通信网络106上的系统(例如输出系统110)发射致使系统实行操作的指令，例如产生工作令或订购特定零件以用于维修资产102。在又另一实例

中，资产102可与远程系统（例如分析系统108）进行通信，所述远程系统接着促进致使操作远离资产102发生。资产102本地执行工作流的其它实例也是可能的。

[0264]　E.模型/工作流修改阶段

[0265]　在另一方面中，分析系统108可实行修改阶段，在所述修改阶段期间，分析系统108基于新的资产数据修改部署的模型及/或工作流。可对聚合及个性化模型及工作流两者执行此阶段。

[0266]　特定来说，当给定资产（例如，资产102）根据模型-工作流对操作时，资产102可向分析系统108提供操作数据及/或数据源112可向分析系统108提供与资产102相关的外部数据。至少基于此数据，分析系统108可修改资产102的模型及/或工作流及/或其它资产（例如资产104）的模型及/或工作流。在修改其它资产的模型及/或工作流时，分析系统108可共享从资产102的行为学习得到的信息。

[0267]　实际上，分析系统108可以多种方式进行修改。图10是描绘可用于修改模型-工作流对的修改阶段的一个可能实例的流程图1000。出于说明目的，实例修改阶段被描述为由分析系统108实行，但是此修改阶段也可由其它系统实行。所属领域的一般技术人员将明白，出于清楚及解释目的提供流程图1000，且可利用操作的数个其它组合来修改模型-工作流对。

[0268]　如图10中所示，在框1002处，分析系统108可接收分析系统108从其识别特定事件的发生的数据。数据可为源自资产102的操作数据或来自数据源112的与资产102相关的外部数据等数据。所述事件可采取上文讨论的任何事件的形式，例如资产102处的故障。

[0269]　在其它实例实施方案中，事件可采取新的组件或子系统被添加到资产102的形式。另一事件可采取“领先指示符”事件的形式，这可能涉及资产102的传感器及/或致动器产生与在模型定义阶段期间在图7的框706处识别的数据不同（可能相差阈值差）的数据。此差可指示资产102具有高于或低于与资产102类似的资产的正常操作状况的操作状况。又另一事件可采取后面是一或多个领先指示符事件的事件的形式。

[0270]　基于所识别的特定事件的发生及/或基础数据（例如，与资产102相关的操作数据及/或外部数据），分析系统108接着可修改聚合、预测模型及/或工作流及/或一或多个个性化预测模型及/或工作流。特定来说，在框1004处，分析系统108可确定是否修改聚合预测模型。分析系统108可由于许多原因而确定修改聚合预测模型。

[0271]　例如，如果所识别的特定事件的发生是包含资产102在内的多个资产的第一次发生（例如在来自资产队伍的资产处第一次发生特定故障或第一次将特定的新组件添加到来自资产队伍的资产），那么分析系统108可修改聚合预测模型。

[0272]　在另一实例中，如果与所识别的特定事件的发生相关联的数据不同于用于最初定义聚合模型的数据，那么分析系统108可进行修改。例如，所识别的特定事件的发生可能在先前不与特定事件的发生相关联的操作状况下发生（例如，特定故障可能在先前未在所述特定故障下测量的相关联传感器值下发生）。修改聚合模型的其它原因也是可能的。

[0273]　如果分析系统108确定修改聚合预测模型，那么分析系统108可在框1006处这样做。否则，分析系统108可前进到框1008。

[0274]　在框1006处，分析系统108可至少部分地基于在框1002处所接收的与资产102相关的数据来修改聚合模型。在实例实施方案中，可以各种方式（例如上文参考图5的框510讨论

的任何方式)修改聚合模型。在其它实施方案中,聚合模型也可以其它方式修改。

[0275] 在框1008处,分析系统108可接着确定是否修改聚合工作流。分析系统108可由于许多原因而修改聚合工作流。

[0276] 例如,分析系统108可基于在框1004处是否修改聚合模型及/或在分析系统108处是否存在其它改变来修改聚合工作流。在其它实例中,尽管资产102执行聚合工作流,但如果在框1002处发生所识别的事件发生,那么分析系统108可修改聚合工作流。例如,如果工作流旨在帮助促进防止事件(例如,故障)的发生且工作流被正确执行,但是事件仍然发生,那么分析系统108可修改聚合工作流。修改聚合工作流的其它原因也是可能的。

[0277] 如果分析系统108确定修改聚合工作流,那么分析系统108可在框1010处这样做。否则,分析系统108可前进到框1012。

[0278] 在框1010处,分析系统108可至少部分地基于在框1002处所接收的与资产102相关的数据来修改聚合工作流。在实例实施方案中,可以各种方式(例如上文参考图5的框514讨论的任何方式)修改聚合工作流。在其它实施方案中,聚合模型也可以其它方式修改。

[0279] 在框1012到框1018处,分析系统108可经配置以至少部分地基于在框1002处接收的与资产102相关的数据(例如,针对资产102及104中的每一者)修改一或多个个性化模型及/或(例如,针对资产102中的一个或资产104)修改一或多个个性化工作流。分析系统108可以类似于框1004到1010的方式来这样做。

[0280] 然而,修改个性化模型或工作流的原因可能与聚合情况的原因不同。例如,分析系统108可进一步考虑首先用于定义个性化模型及/或工作流的基础资产特性。在特定实例中,如果所识别的特定事件的发生是具有资产102的资产特性的资产的此特定事件的第一次发生,那么分析系统108可修改个性化模型及/或工作流。修改个性化模型及/或工作流的其它原因也是可能的。

[0281] 为了说明,图6D是修改的模型-工作流对630的概念说明。具体来说,模型-工作流对说明630是来自图6A的聚合模型-工作流对的修改版本。如所示,修改的模型-工作流对说明630包含来自图6A的模型输入602的原始列,且包含用于模型计算634、模型输出范围636及工作流操作638的修改列。在此实例中,修改的预测模型具有来自传感器A的单个输入数据,且具有两个计算值:计算值I及III。如果修改模型的输出概率小于75%,那么执行工作流操作1。如果输出概率在75%与85%之间,那么执行工作流操作2。且如果输出概率大于85%,那么执行工作流操作3。其它实例修改模型-工作流对在本文中是可能的且予以考虑。

[0282] 返回到图10,在框1020处,分析系统108接着可向一或多个资产发射任何模型及/或工作流修改。例如,分析系统108可向资产102(例如,数据引起修改的资产)发射修改的个性化模型-工作流对且向资产104发射修改的聚合模型。以此方式,分析系统108可基于与资产102的操作相关联的数据来动态地修改模型及/或工作流,并将此类修改分配给多个资产,例如资产102所属的队伍。因此,其它资产可从源自资产102的数据中受益,因为其它资产的本地模型-工作流对可基于此数据被改进,由此帮助创建更加准确及稳健的模型-工作流对。

[0283] 虽然上述修改阶段被讨论为由分析系统108执行,但是在实例实施方案中,资产102的本地分析装置可另外或替代地以与上文讨论的类似的方式实行修改阶段。例如,在一个实例中,当资产102通过利用由一或多个传感器及/或致动器产生的操作数据来操作时,

本地分析装置可修改模型-工作流对。因此，资产102的本地分析装置、分析系统108或其一些组合可在资产相关状况改变时修改预测模型及/或工作流。以此方式，本地分析装置及/或分析系统108可基于其可用的最近的数据连续地调整模型-工作流对。

[0284] F.动态执行模型/工作流

[0285] 在另一方面中，资产102及/或分析系统108可经配置以动态调整执行模型-工作流对。特定来说，资产102及/或分析系统108可经配置以检测触发关于资产102及/或分析系统108是否应执行预测模型及/或工作流的责任的改变的某些事件。

[0286] 在操作中，资产102及分析系统108两者均可执行代表资产102的模型-工作流对的全部或部分。例如，在资产102从分析系统108接收模型-工作流对之后，资产102可将模型-工作流对存储在数据存储装置中，但是接着可依赖于分析系统108集中执行模型-工作对中的部分或全部。特定来说，资产102可向分析系统108提供至少传感器及/或致动器数据，分析系统108接着可使用此数据来集中执行资产102的预测模型。基于模型的输出，分析系统108接着可执行对应工作流，或分析系统108可向资产102发射模型的输出或使资产102本地执行工作流的指令。

[0287] 在其它实例中，分析系统108可依赖于资产102来本地执行模型-工作流对中的部分或全部。具体来说，资产102可本地执行预测模型中的部分或全部，并将结果发射到分析系统108，所述结果接着可致使分析系统108集中执行对应工作流。或资产102也可本地执行对应工作流。

[0288] 在又其它实例中，分析系统108及资产102可共享执行模型-工作流对的责任。例如，分析系统108可集中执行模型及/或工作流的部分，而资产102本地执行模型及/或工作流的其它部分。资产102及分析系统108可发射来自它们相应的已执行的责任的结果。其它实例也是可能的。

[0289] 在某个时间点，资产102及/或分析系统108可确定应当调整模型-工作流对的执行。即，一或两者可确定执行责任应当被修改。此操作可以各种方式发生。

[0290] 图11是描绘可用于调整模型-工作流对的执行的调整阶段的一个可能实例的流程图1100。出于说明目的，实例调整阶段被描述为由资产102及/或分析系统108实行，但是此修改阶段也可由其它系统实行。所属领域的一般技术人员将明白，出于清楚及解释目的提供流程图1100，且可利用操作的数个其它组合来调整模型-工作流对的执行。

[0291] 在框1102处，资产102及/或分析系统108可检测调整因子(或潜在地多个调整因子)，其指示需要对模型-工作流对的执行进行调整的状况。此类状况的实例包含通信网络106的网络状况或资产102及/或分析系统108的处理状况等等。实例网络状况可包含网络延迟、网络带宽、资产102与通信网络106之间的链路的信号强度，或网络性能的一些其它指示等等。实例处理状况可包含处理容量(例如，可用的处理能力)、处理使用量(例如，所消耗的处理能力的量)或处理能力的一些其它指示等等。

[0292] 实际上，检测调整因子可以各种方式执行。例如，此操作可涉及确定网络(或处理)状况是否达到一或多个阈值或状况是否以某种方式改变。检测调整因子的其它实例也是可能的。

[0293] 特定来说，在一些情况下，检测调整因子可涉及资产102及/或分析系统108检测到资产102与分析系统108之间的通信链路的信号强度低于阈值信号强度的指示或以某个变

化率降低的指示。在此实例中，调整因子可指示资产102即将“离线”。

[0294] 在另一种情况下，检测调整因子可另外或可替代地涉及资产102及/或分析系统108检测到网络延迟高于阈值延迟或以某个变化率增加的指示。或所述指示可为网络带宽低于阈值带宽或以某个变化率下降。在这些实例中，调整因子可指示通信网络106滞后。

[0295] 在又其它情况下，检测调整因子可另外或替代地涉及资产102及/或分析系统108检测到处理容量低于特定阈值或以某个变化率降低及/或处理使用量高于阈值或以某个变化率增加的指示。在此类实例中，调整因子可指示资产102(及/或分析系统108)的处理能力较低。检测调整因子的其它实例也是可能的。

[0296] 在框1104处，基于检测到的调整因子，可调整本地执行责任，这可以多种方式发生。例如，资产102可能已经检测到调整因子，且接着确定本地执行模型-工作流对或其部分。在一些情况下，资产102接着可向分析系统108发射资产102在本地执行预测模型及/或工作流的通知。

[0297] 在另一实例中，分析系统108可能已经检测到调整因子，且接着向资产102发射指令以致使资产102本地执行模型-工作流对或其部分。基于所述指令，资产102接着可在本地执行模型-工作流对。

[0298] 在框1106处，可调整集中执行责任，这可以多种方式发生。例如，可基于分析系统108检测到资产102在本地执行预测模型及/或工作流的指示来调整集中执行责任。分析系统108可以多种方式检测此指示。

[0299] 在一些实例中，分析系统108可通过从资产102接收资产102在本地执行预测模型及/或工作流的通知来检测所述指示。通知可采取各种形式，例如二进制或文本，且可识别资产在本地执行的特定预测模型及/或工作流。

[0300] 在其它实例中，分析系统108可基于资产102的所接收的操作数据来检测指示。具体来说，检测指示可涉及分析系统108接收资产102的操作数据，且接着检测所接收的数据的一或多个特性。根据所接收的数据的一或多个检测到的特性，分析系统108可推断出资产102在本地执行预测模型及/或工作流。

[0301] 实际上，可以各种方式来执行检测所接收的数据的一或多个特性。例如，分析系统108可检测所接收的数据的类型。特定来说，分析系统108可检测数据源，例如产生传感器或致动器数据的特定传感器或致动器。基于所接收的数据的类型，分析系统108可推断出资产102在本地执行预测模型及/或工作流。例如，基于检测到特定传感器的传感器标识符，分析系统108可推断出资产102在本地执行预测模型及对应工作流，其导致资产102从特定传感器采集数据并向分析系统108发射所述数据。

[0302] 在另一情况中，分析系统108可检测所接收的数据的量。分析系统108可将所述量与某个数据阈值量进行比较。基于所述量达到阈值量，分析系统108可推断出资产102正在本地执行致使资产102采集等于或大于阈值量的数据量的预测模型及/或工作流。其它实例也是可能的。

[0303] 在实例实施方案中，检测所接收的数据的一或多个特性可涉及分析系统108检测所接收的数据的一或多个特性中的特定变化，例如所接收的数据的类型的变化、所接收的数据量的变化或接收数据的频率的变化。在特定实例中，所接收的数据的类型的变化可涉及分析系统108检测其正在接收的传感器数据源的改变(例如，正在产生被提供给分析系统

108的数据的传感器及/或致动器的变化)。

[0304] 在一些情况下,检测所接收的数据中的变化可涉及分析系统108将近期接收的数据与过去(例如,当前时间之前的一小时、一天、一周等)接收的数据进行比较。无论如何,基于检测到所接收的数据的一或多个特性的变化,分析系统108可推断出资产102在本地执行预测模型及/或工作流,其导致对由资产102提供到分析系统108的所述数据的此变化。

[0305] 另外,分析系统108可基于在框1102处检测到调整因子来检测资产102在本地执行预测模型及/或工作流的指示。例如,如果分析系统108在框1102处检测到调整因子,那么分析系统108可向资产102发射致使资产102调整其本地执行责任的指令,且相应地,分析系统108可调整自己的集中执行责任。检测指示的其它实例也是可能的。

[0306] 在实例实施方案中,可根据对本地执行责任的调整来调整集中执行责任。例如,如果资产102现在在本地执行预测模型,那么分析系统108可相应地停止集中执行预测模型(且可或可不中止执行对应工作流)。另外,如果资产102在本地执行对应工作流,那么分析系统108可相应地停止执行工作流(且可或可不停止集中执行预测模型)。其它实例也是可能的。

[0307] 实际上,资产102及/或分析系统108可连续地执行框1102到1106的操作。有时,可调整本地及集中执行责任以促进优化模型-工作流对的执行。

[0308] 另外,在一些实施方案中,资产102及/或分析系统108可基于检测到调整因子来执行其它操作。例如,基于通信网络106的状况(例如,带宽、延迟、信号强度或网络质量的另一指示),资产102可在本地执行特定工作流。分析系统108可基于分析系统108检测通信网络的状况提供特定工作流,所述特定工作流可能已经存储在资产102上,或可为已经存储在资产102上的工作流的修改版本(例如,资产102可在本地修改工作流)。在一些情况下,除其它可能的工作流操作外,特定工作流可包含增加或减少采样速率的数据采集方案及/或增加或减少被发射到分析系统108的数据的发射率或数量的数据发射方案。

[0309] 在特定实例中,资产102可确定通信网络的一或多个检测到的状况已经达到相应的阈值(例如,指示网络质量不良)。基于此确定,资产102可在本地执行工作流,所述工作流包含根据数据发射方案发射数据,所述数据发射方案减少资产102发射到分析系统108的数据量及/或频率。其它实例也是可能的。

[0310] <u>V.实例方法</u>

[0311] 现在转到图12,描绘了说明用于定义及部署可由分析系统108执行的聚合预测模型及对应工作流的实例方法1200的流程图。对于方法1200及下文讨论的其它方法,由流程图中的框所说明的操作可根据上述讨论来执行。另外,上文讨论的一或多个操作可被添加到给定的流程图。

[0312] 在框1202处,方法1200可涉及分析系统108接收多个资产(例如资产102及104)的相应操作数据。在框1204处,方法1200可涉及分析系统108基于所接收的操作数据来定义与多个资产的操作相关的预测模型及对应工作流(例如,故障模型及对应工作流)。在框1206处,方法1200可涉及分析系统108向多个资产中的至少一个资产(例如,资产102)发射预测模型及对应工作流以供至少一个资产进行本地执行。

[0313] 图13描绘了用于定义及部署可由分析系统108执行的个性化预测模型及/或对应工作流的实例方法1300的流程图。在框1302处,方法1300可涉及分析系统108接收多个资产

的操作数据，其中多个资产至少包含第一资产(例如资产102)。在框1304处，方法1300可涉及分析系统108基于所接收的操作数据来定义与多个资产的操作相关的聚合预测模型及聚合对应工作流。在框1306处，方法1300可涉及分析系统108确定第一资产的一或多个特性。在框1308处，方法1300可涉及分析系统108基于第一资产的一或多个特性以及聚合预测模型及聚合对应工作流来定义与第一资产的操作相关的个性化预测模型或个性化对应工作流中的至少一者。在框1310处，方法1300可涉及分析系统108向第一资产发射所定义的至少一个个性化预测模型或个性化对应工作流以供第一资产进行本地执行。

[0314] 图14描绘了用于动态地修改可由分析系统108执行的模型-工作流对的执行的实例方法1400的流程图。在框1402处，方法1400可涉及分析系统108向资产(例如，资产102)发射与资产的操作相关的预测模型及对应工作流以供资产进行本地执行。在框1404处，方法1400可涉及分析系统108检测资产在本地执行预测模型或对应工作流中的至少一者的指示。在框1406处，方法1400可涉及分析系统108基于检测到的指示来通过预测模型或对应工作流中的至少一者的计算系统来修改集中执行。

[0315] 类似于方法1400，用于动态地修改模型-工作流对的执行的另一方法可由资产(例如资产102)执行。例如，此方法可涉及资产102从中央计算系统(例如，分析系统108)接收与资产102的操作相关的预测模型及对应工作流。所述方法还可涉及资产102检测指示与调整预测模型及对应工作流的执行相关联的一或多个状况的调整因子。所述方法可涉及基于检测到的调整因子(i)修改预测模型或对应工作流中的至少一者的资产102的本地执行，及(ii)向中央计算系统发射资产102在本地执行预测模型或对应工作流中的至少一者的指示以促进中央计算系统通过预测模型或对应工作流中的至少一者的计算系统修改集中执行。

[0316] 图15描绘了例如由资产102的本地分析装置本地执行模型-工作流对的实例方法1500的流程图。在框1502处，方法1500可涉及本地分析装置经由网络接口接收与经由本地分析装置的资产接口耦合到本地分析装置的资产(例如，资产102)的操作相关的预测模型，其中预测模型由远离本地分析装置定位的计算系统(例如，分析系统108)基于多个资产的操作数据来定义。在框1504处，方法1500可涉及本地分析装置经由资产接口接收用于资产102的操作数据(例如，由一或多个传感器及/或致动器产生且可经由资产的中央处理单元间接地或直接从一或多个传感器及/或致动器接收的操作数据)。在框1506处，方法1500可涉及本地分析装置基于用于资产102的所接收的操作数据的至少部分来执行预测模型。在框1508处，方法1500可涉及本地分析装置基于执行预测模型来执行对应于预测模型的工作流，其中执行工作流包含经由资产接口致使资产102执行操作。

[0317] VI.结论

[0318] 上文已经描述了所揭示的创新的实例实施例。然而，所属领域的技术人员将会理解，可对所描述的实施例进行改变及修改而不背离由权利要求界定的本发明的真实范围及精神。

[0319] 另外，就本文中所描述的实例涉及由例如“人”、“操作者”、“用户”或其它实体等行动者执行或发起的操作来说，这仅仅是出于实例及解释的目的。权利要求不应被解释为要求此类行动者采取行动，除非在权利要求语言中明确叙述。

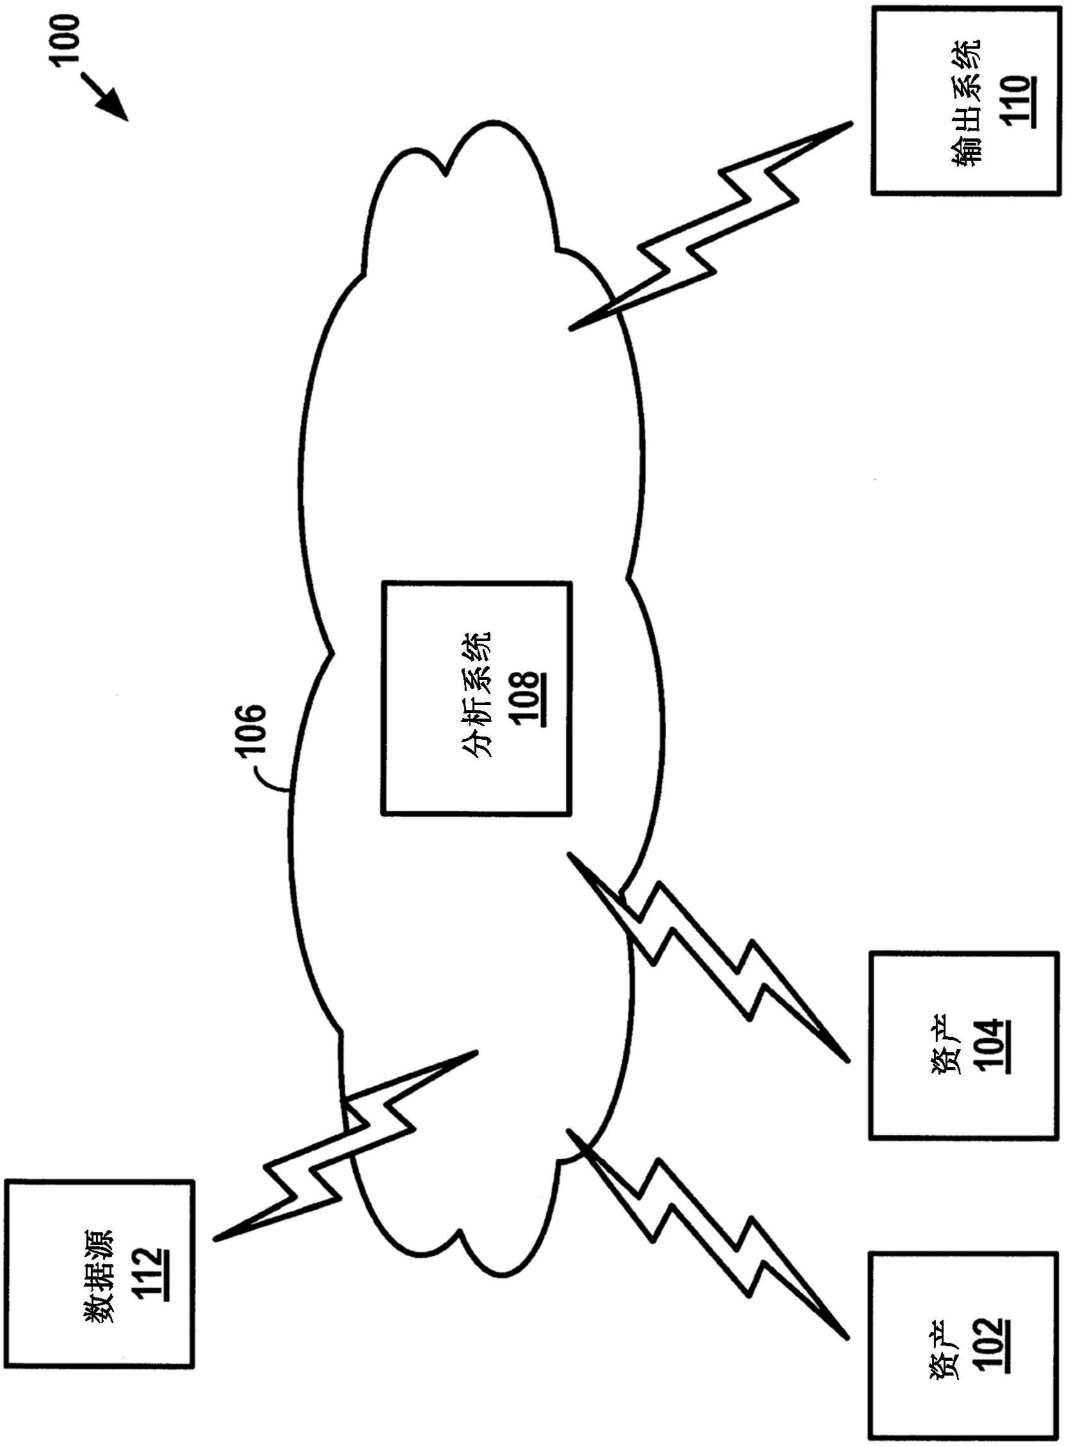

图1

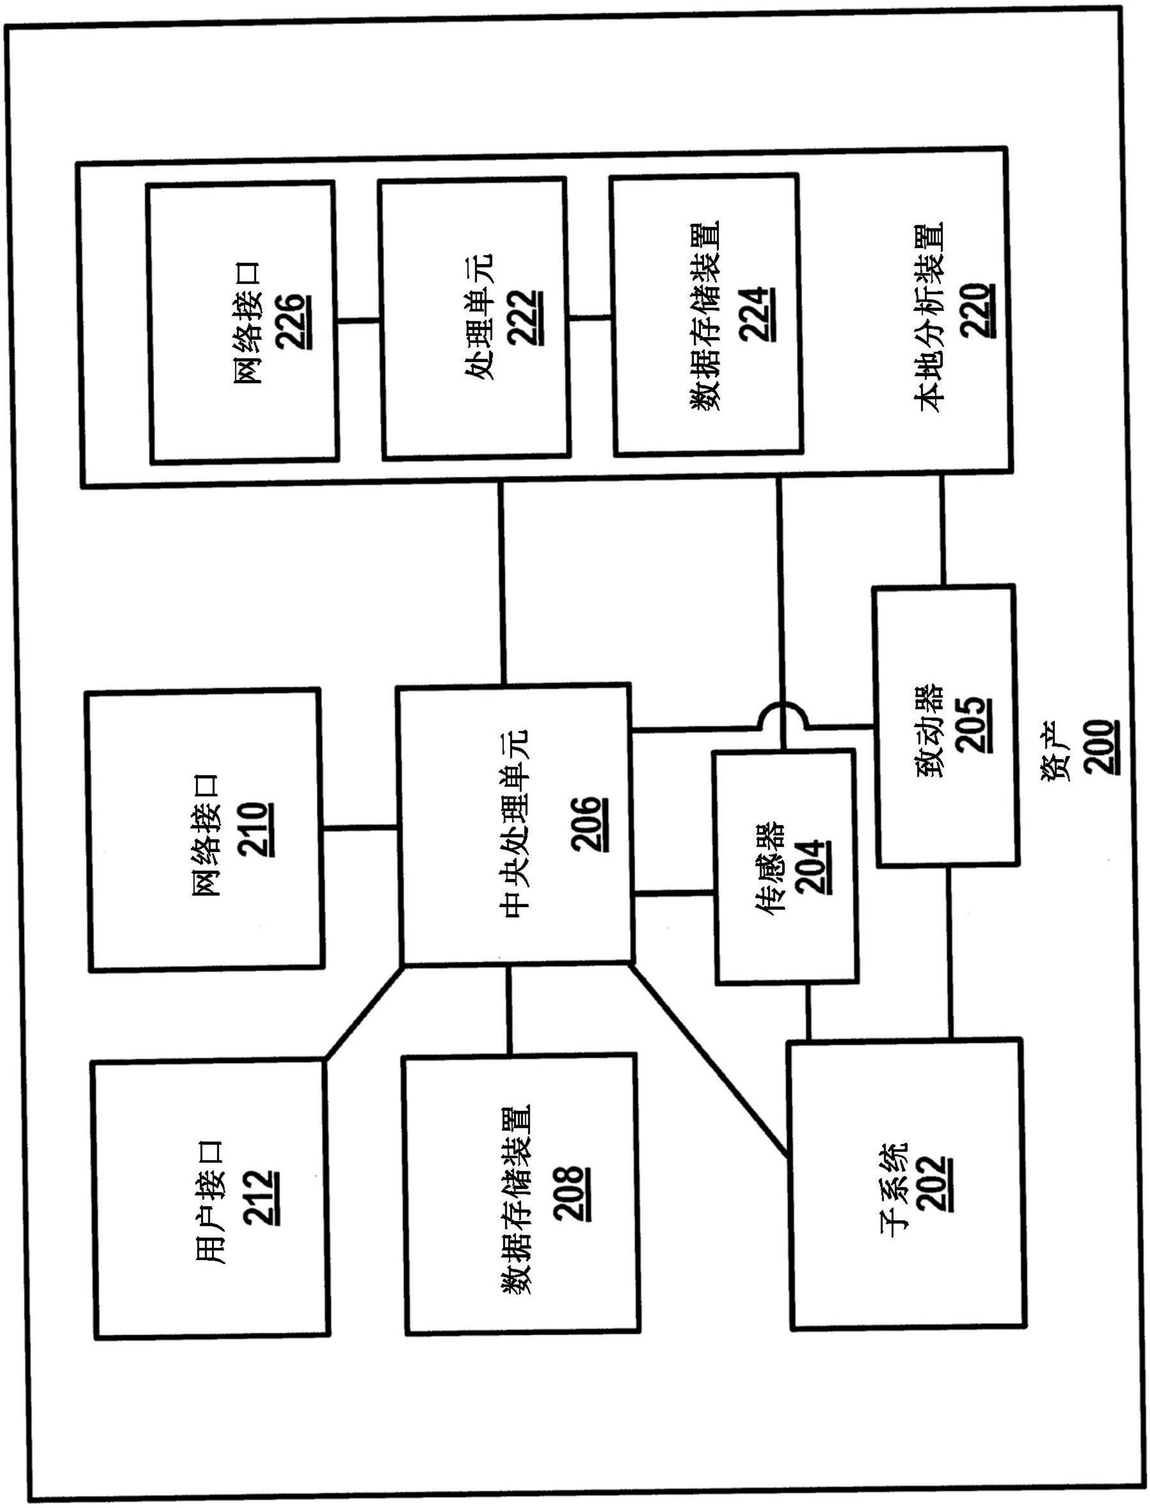
网络接口
226
处理单元
222
数据存储装置
224
本地分析装置
220
网络接口
210
中央处理单元
206
传感器
204
致动器
205
资产
200
用户接口
212
数据存储装置
208
子系统
202

图2

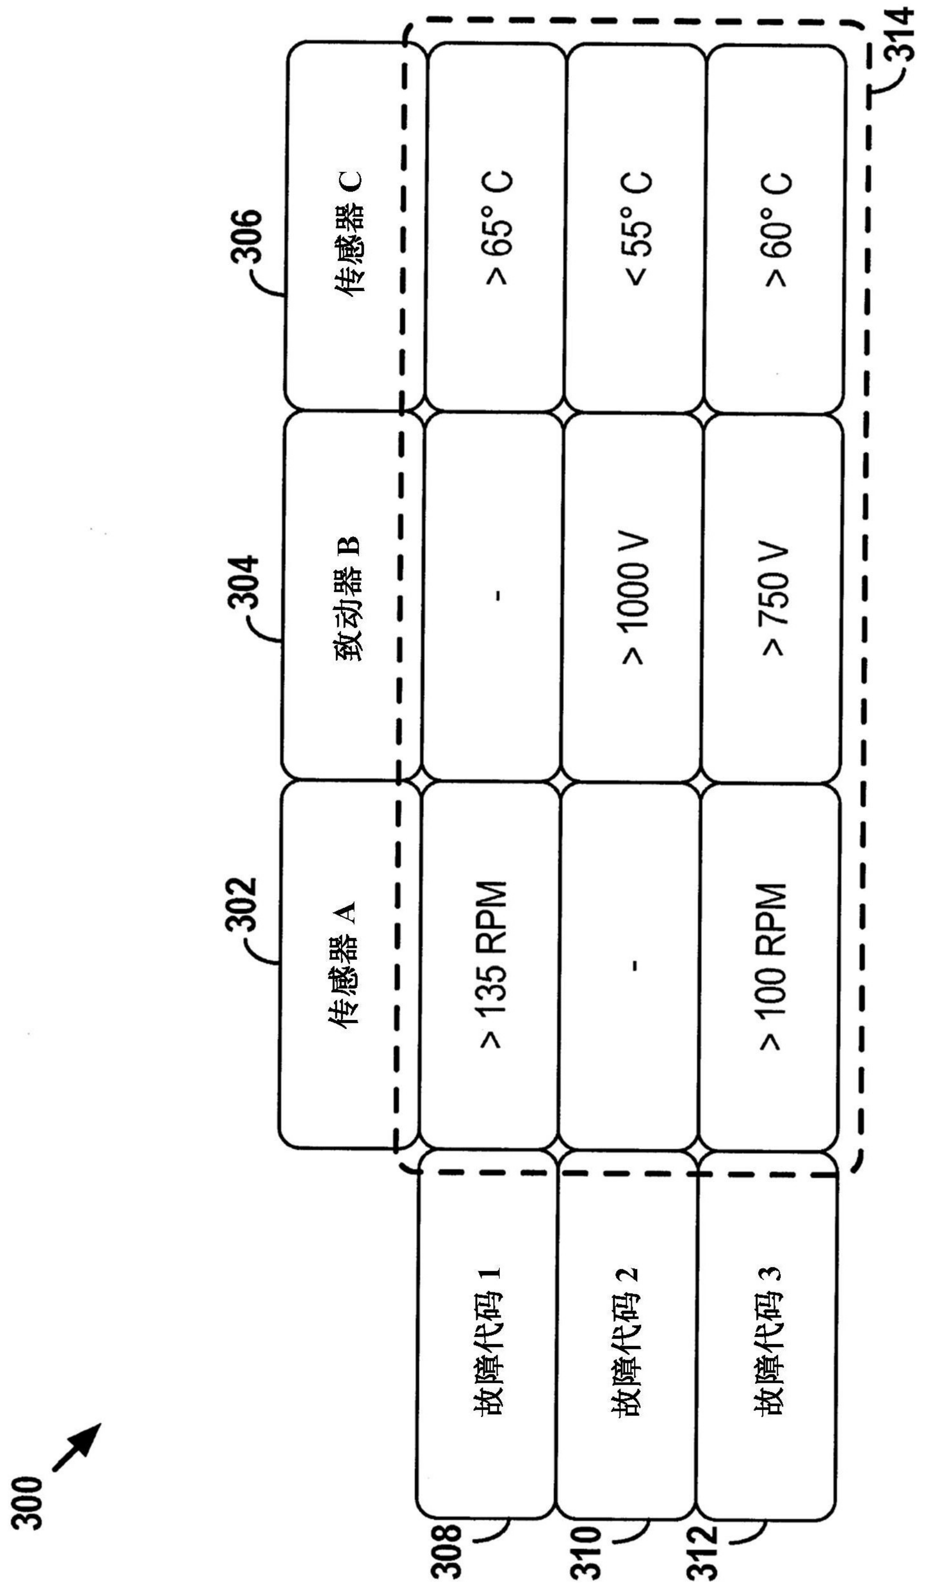
300
302
304
306
传感器 A
致动器 B
传感器 C
> 135 RPM
-
> 65° C
-
> 1000 V
< 55° C
> 100 RPM
> 750 V
> 60° C
故障代码 1
故障代码 2
故障代码 3
308
310
312
314

图3

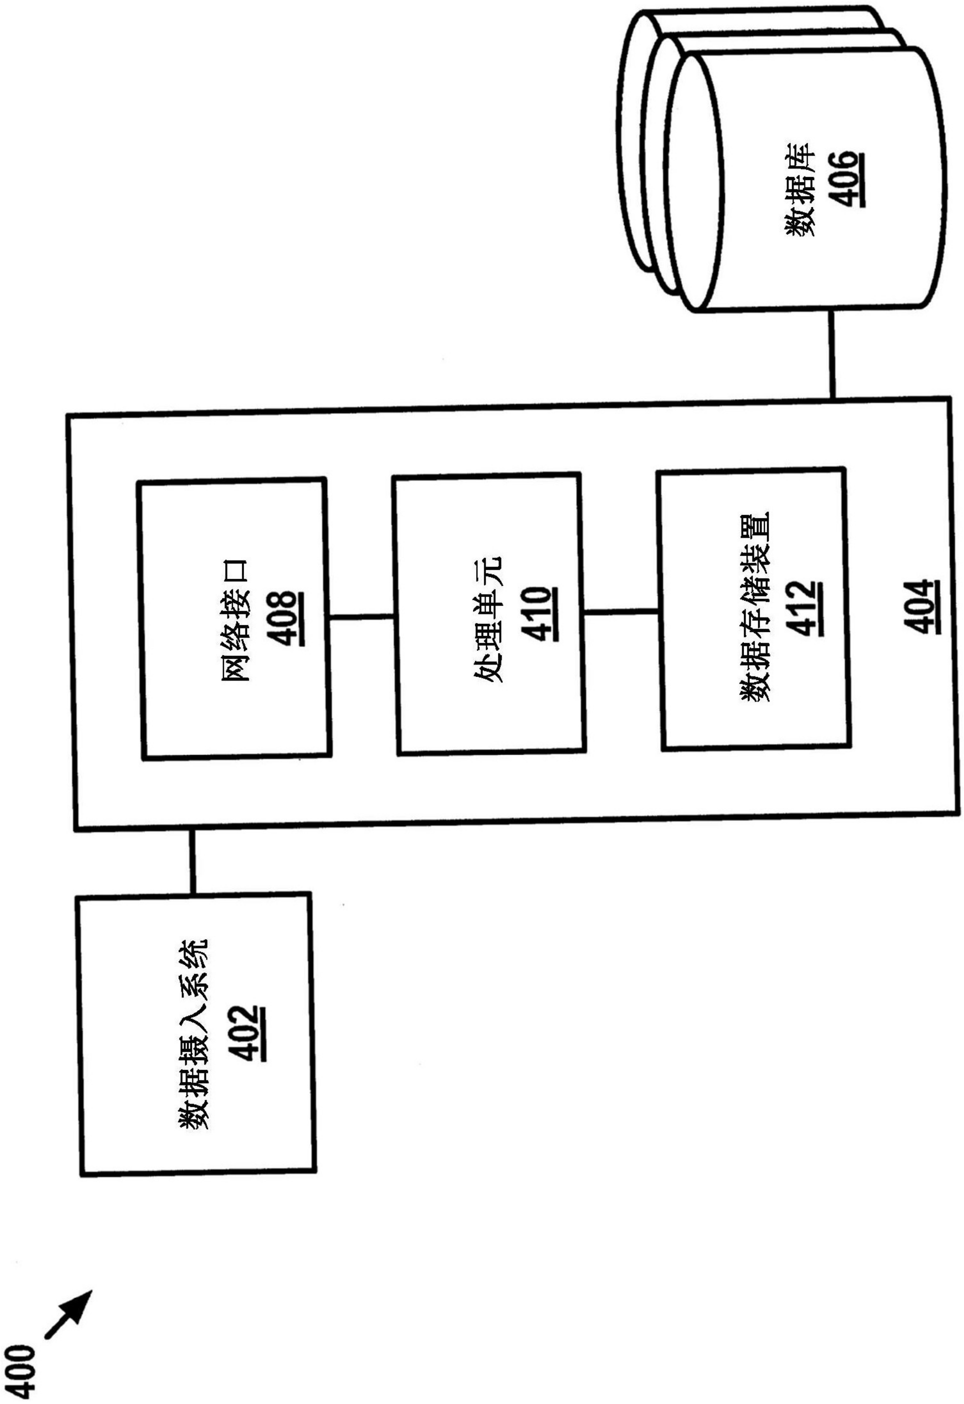

图4

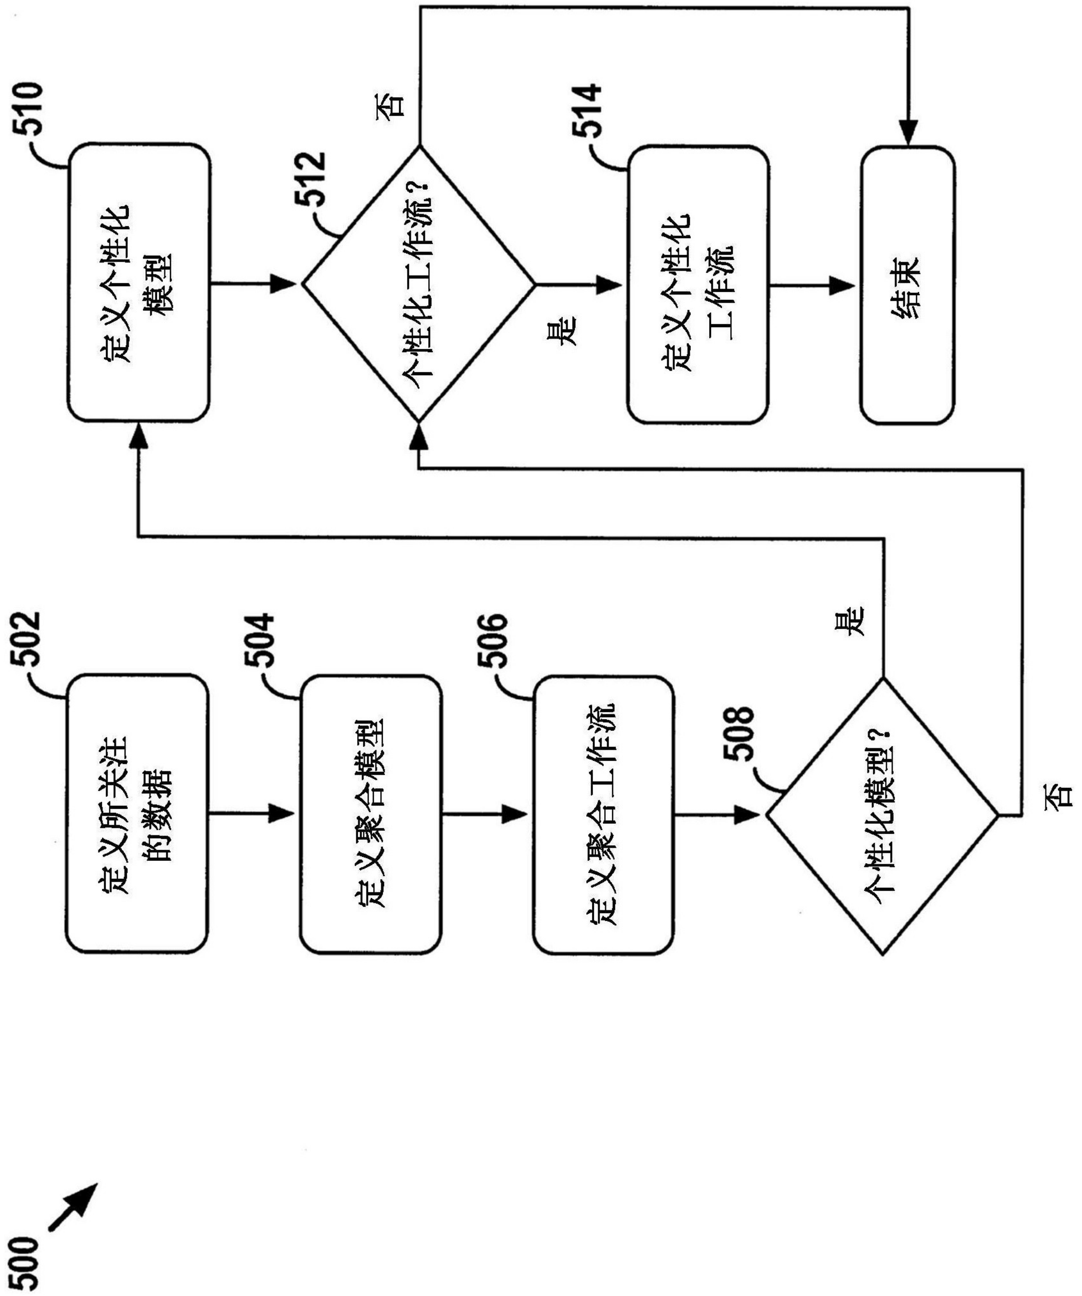
500
502
定义所关注的数据
504
定义聚合模型
506
定义聚合工作流
508
个性化模型？
是
否
510
定义个性化模型
512
个性化工作流？
是
否
514
定义个性化工作流
结束

图5

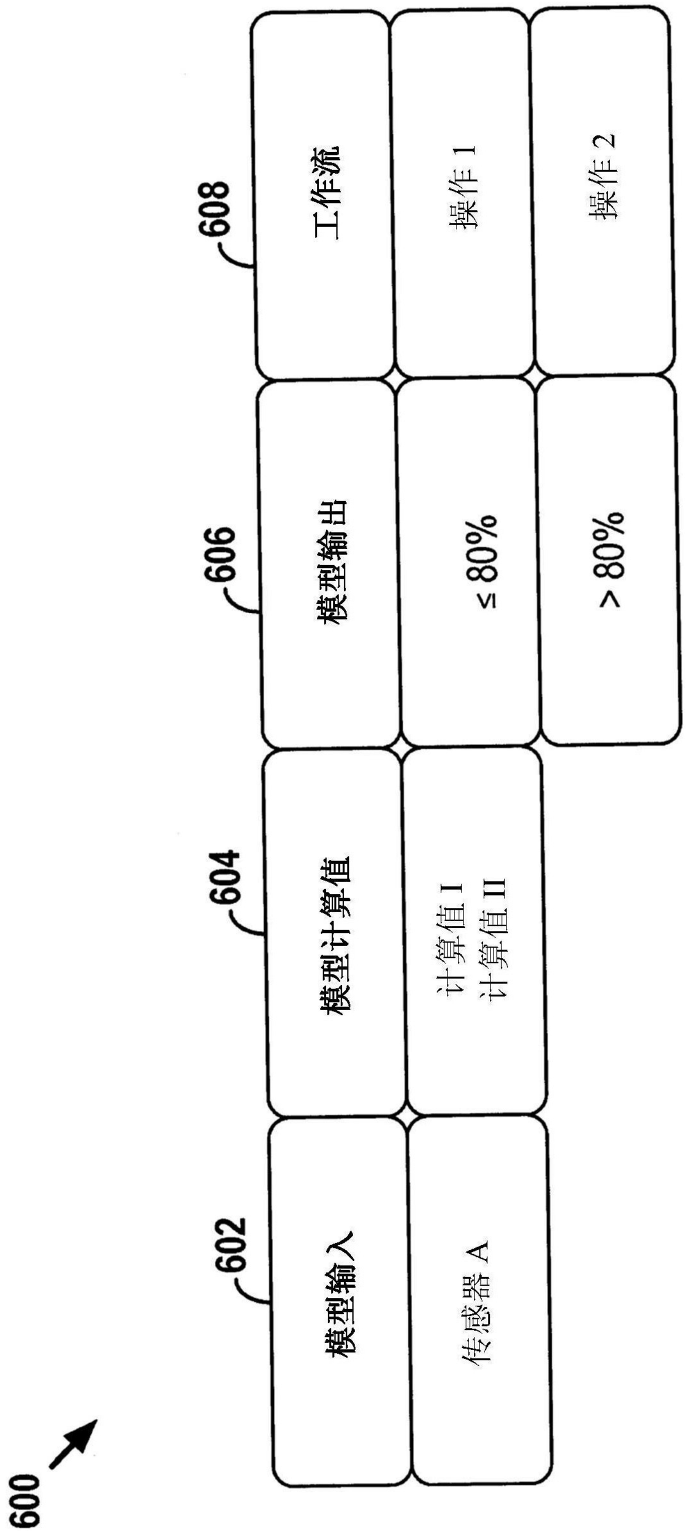
600
602
模型输入
传感器 A
604
模型计算值
计算值 I
计算值 II
606
模型输出
≤80%
>80%
608
工作流
操作 1
操作 2

图6A

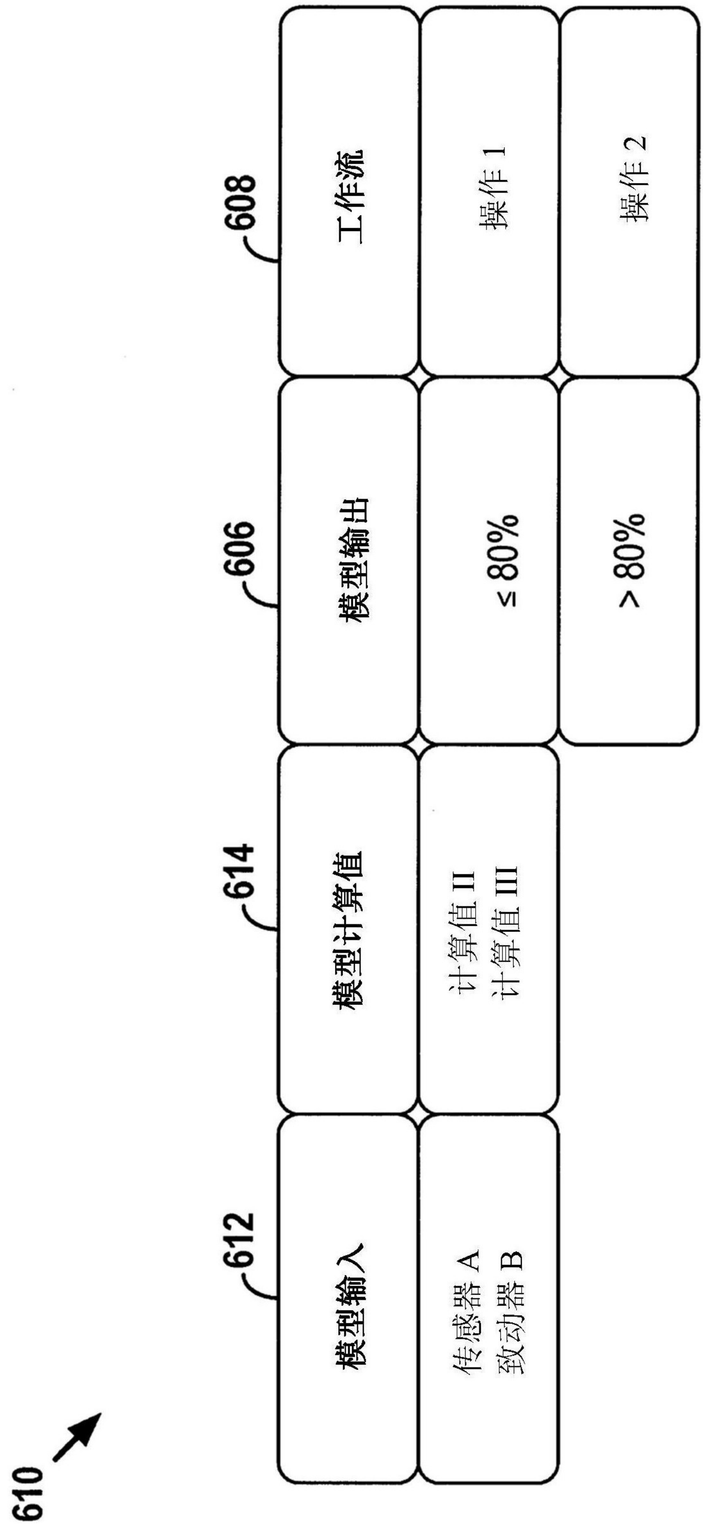
610
612
模型输入
传感器 A
致动器 B
614
模型计算值
计算值 II
计算值 III
606
模型输出
≤80%
>80%
608
工作流
操作 1
操作 2

图6B

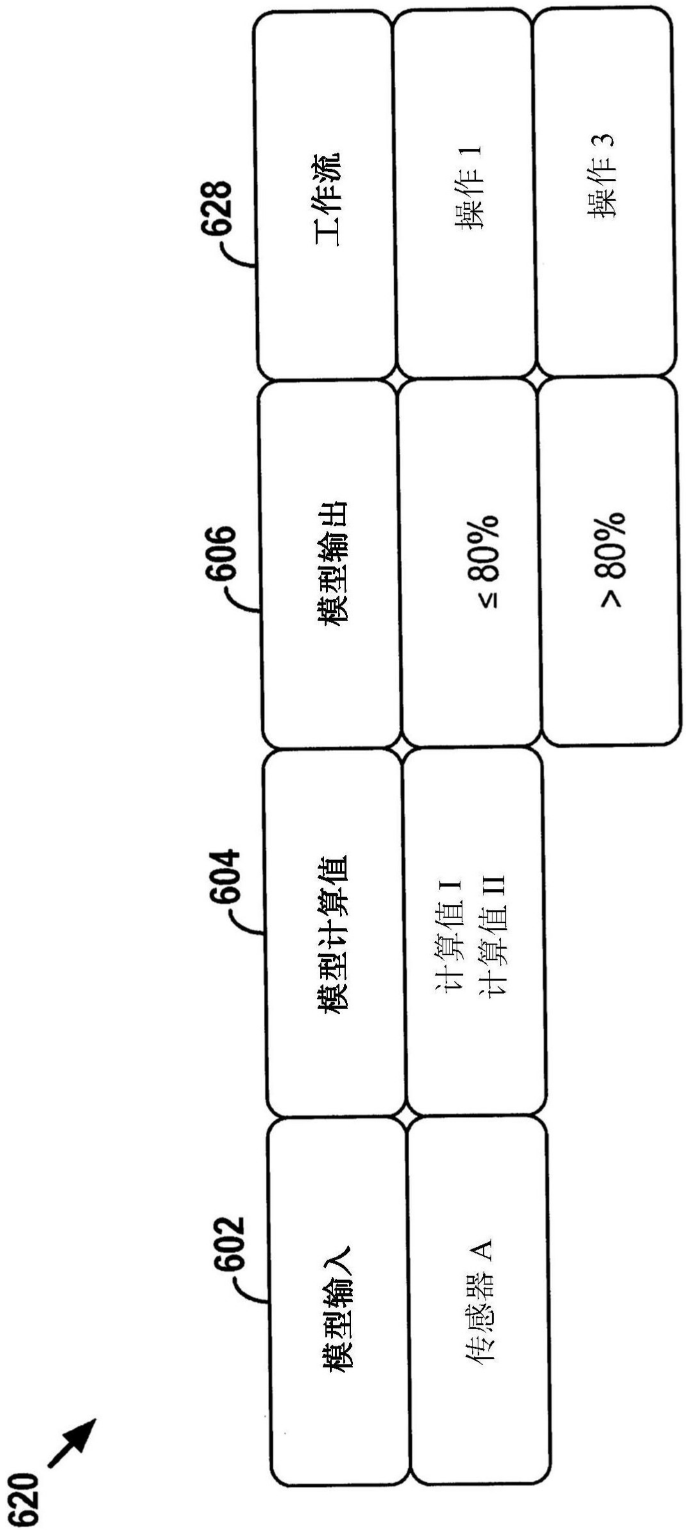
620
602
模型输入
传感器 A
604
模型计算值
计算值 I
计算值 II
606
模型输出
≤80%
>80%
628
工作流
操作 1
操作 3

图6C

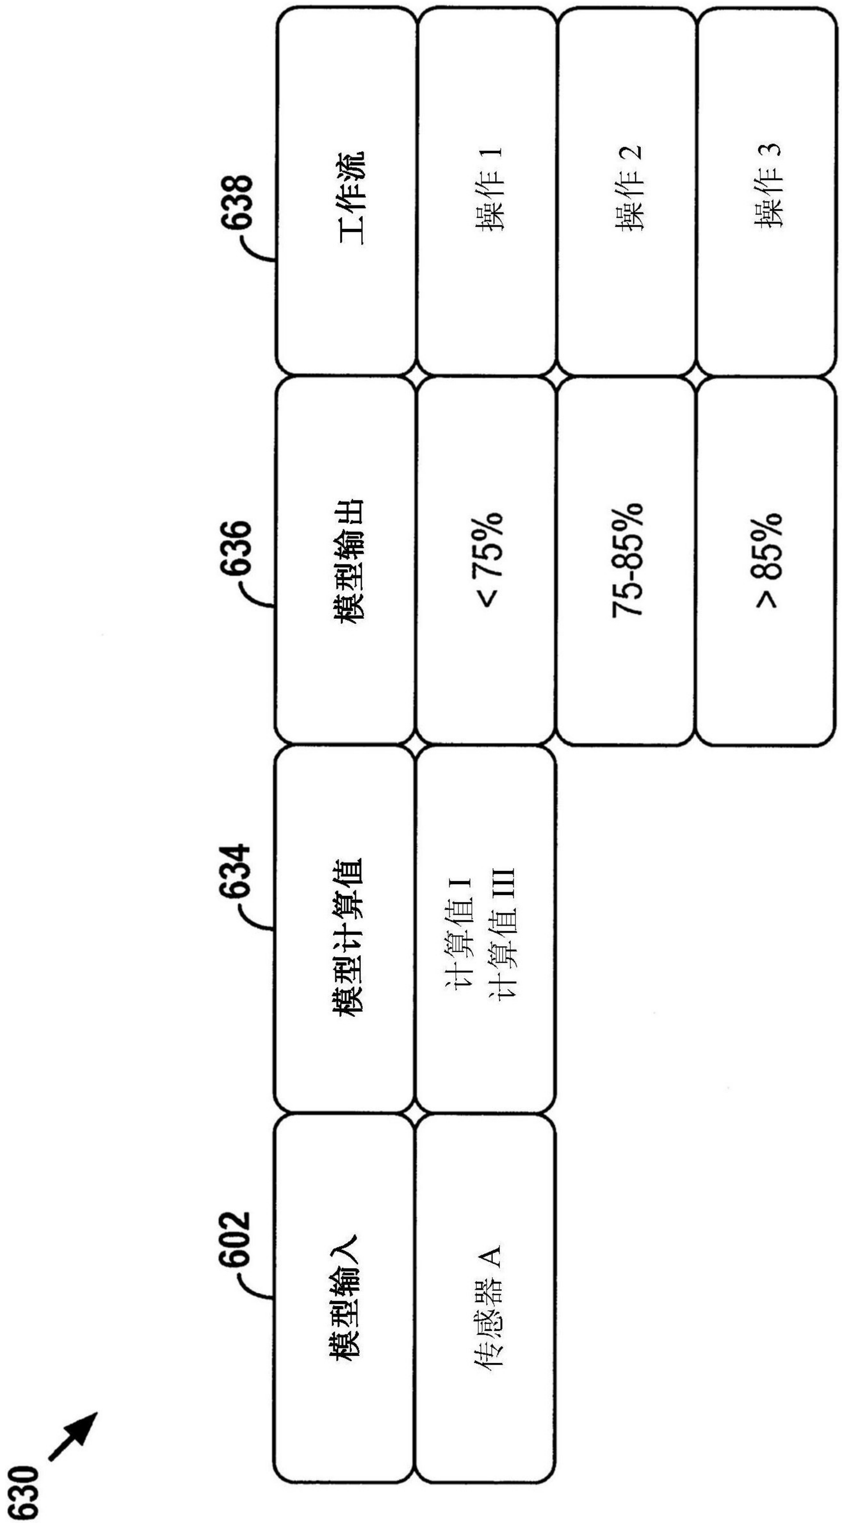
630
602
模型输入
传感器 A
634
模型计算值
计算值 I
计算值 III
636
模型输出
< 75%
75-85%
> 85%
638
工作流
操作 1
操作 2
操作 3

图6D

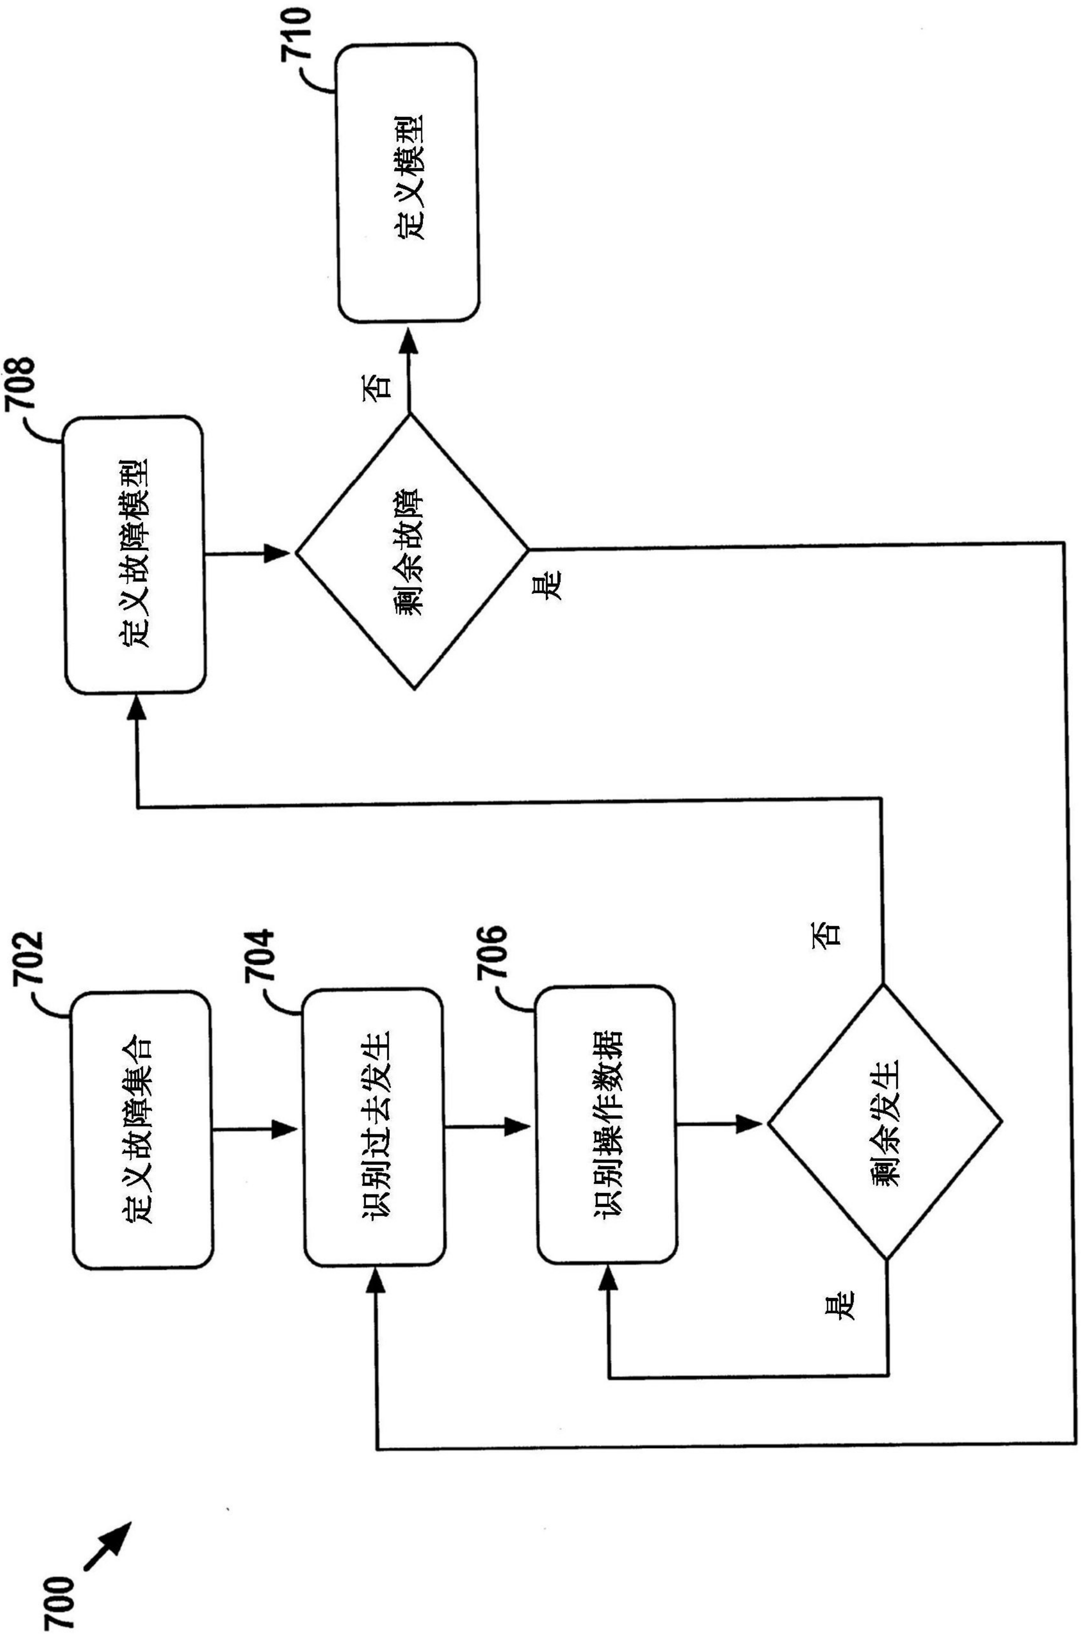
700
702
定义故障集合
704
识别过去发生
706
识别操作数据
剩余发生
是
否
708
定义故障模型
剩余故障
是
否
710
定义模型

图7

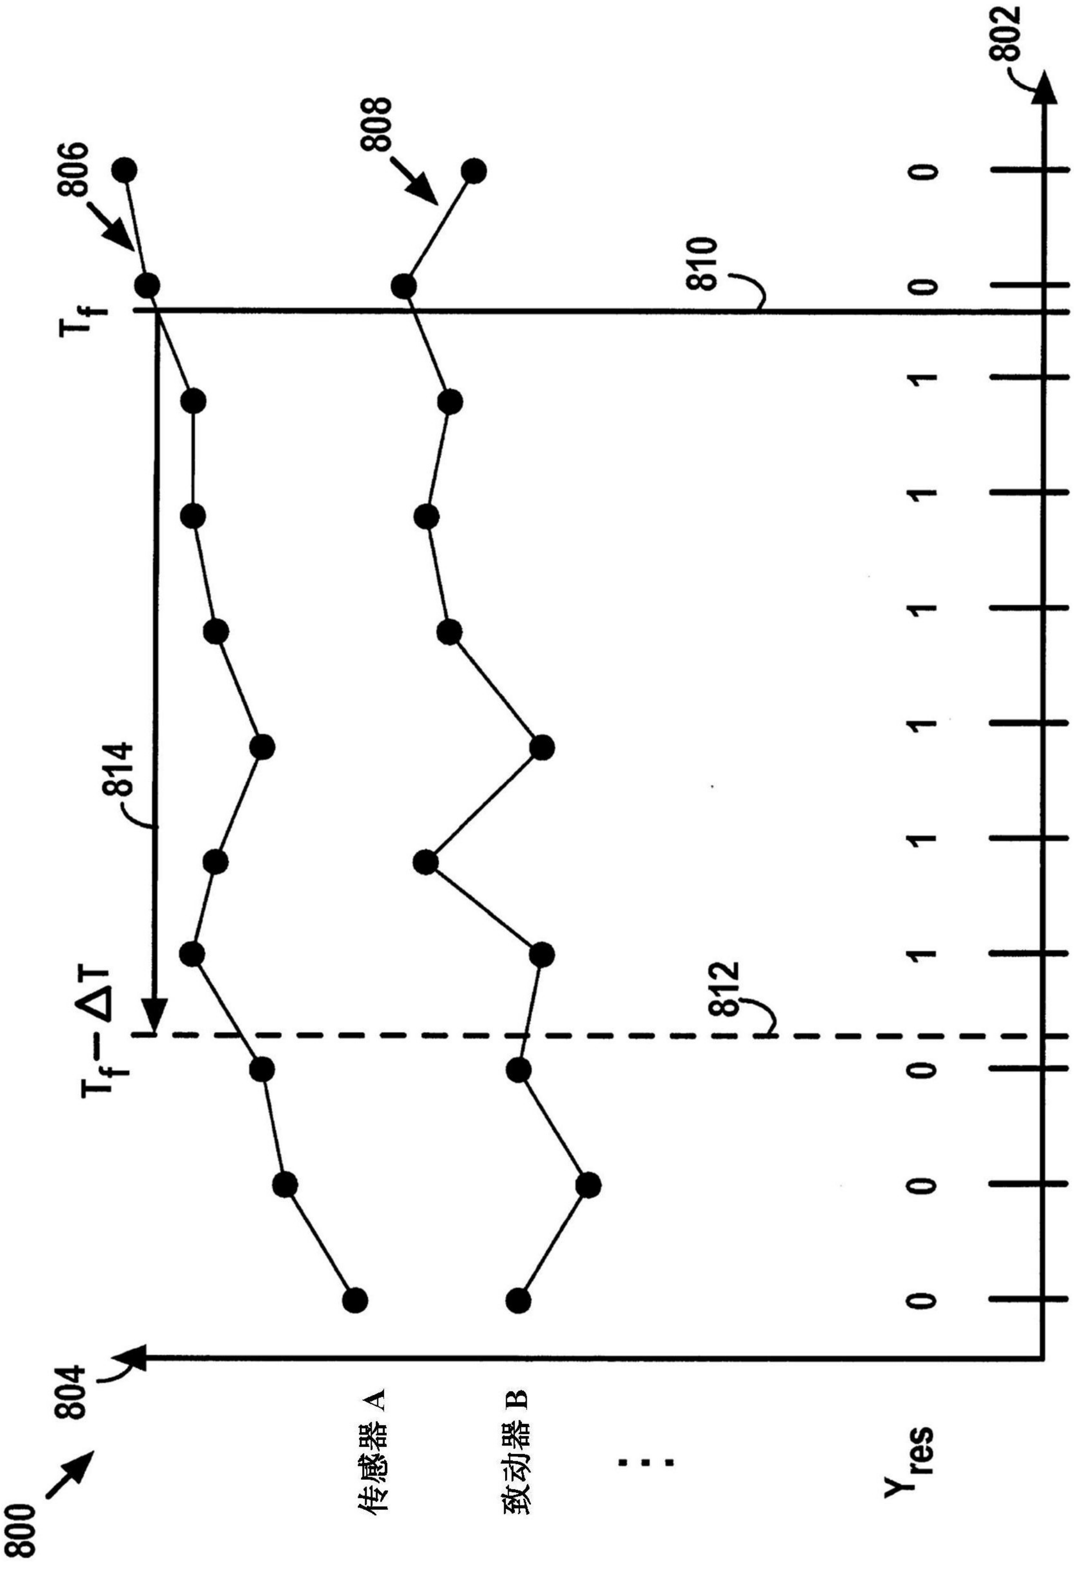
800
804
802
806
808
810
812
814
T_f
T_f − ΔT
传感器 A
致动器 B
...
Y_res
0 0 0 1 1 1 1 1 1 0 0
T_i T_i+1 T_i+2 T_i+3 T_i+4 T_i+5 T_i+6 T_i+7 T_i+8 T_i+9 T_i+10

图8

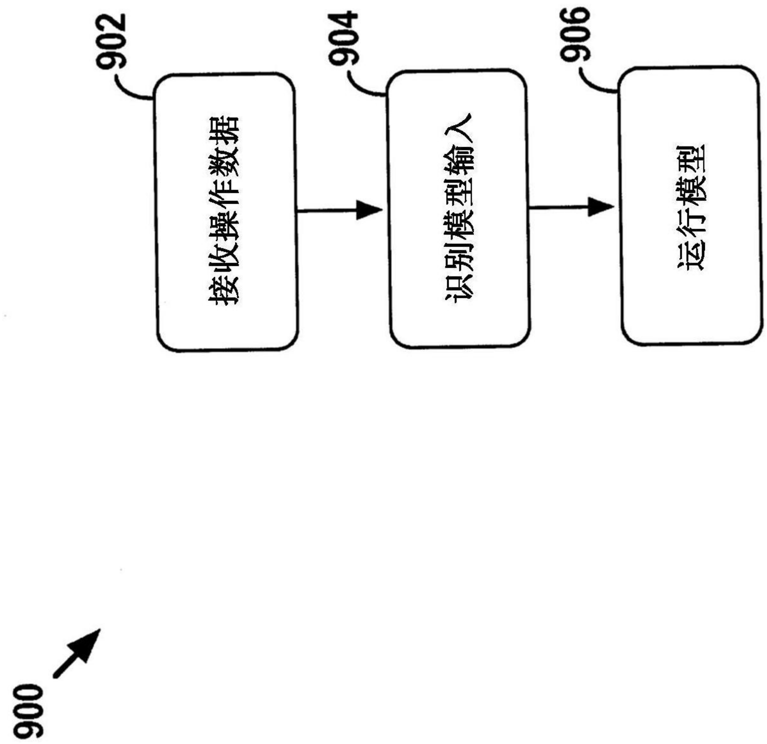

图9

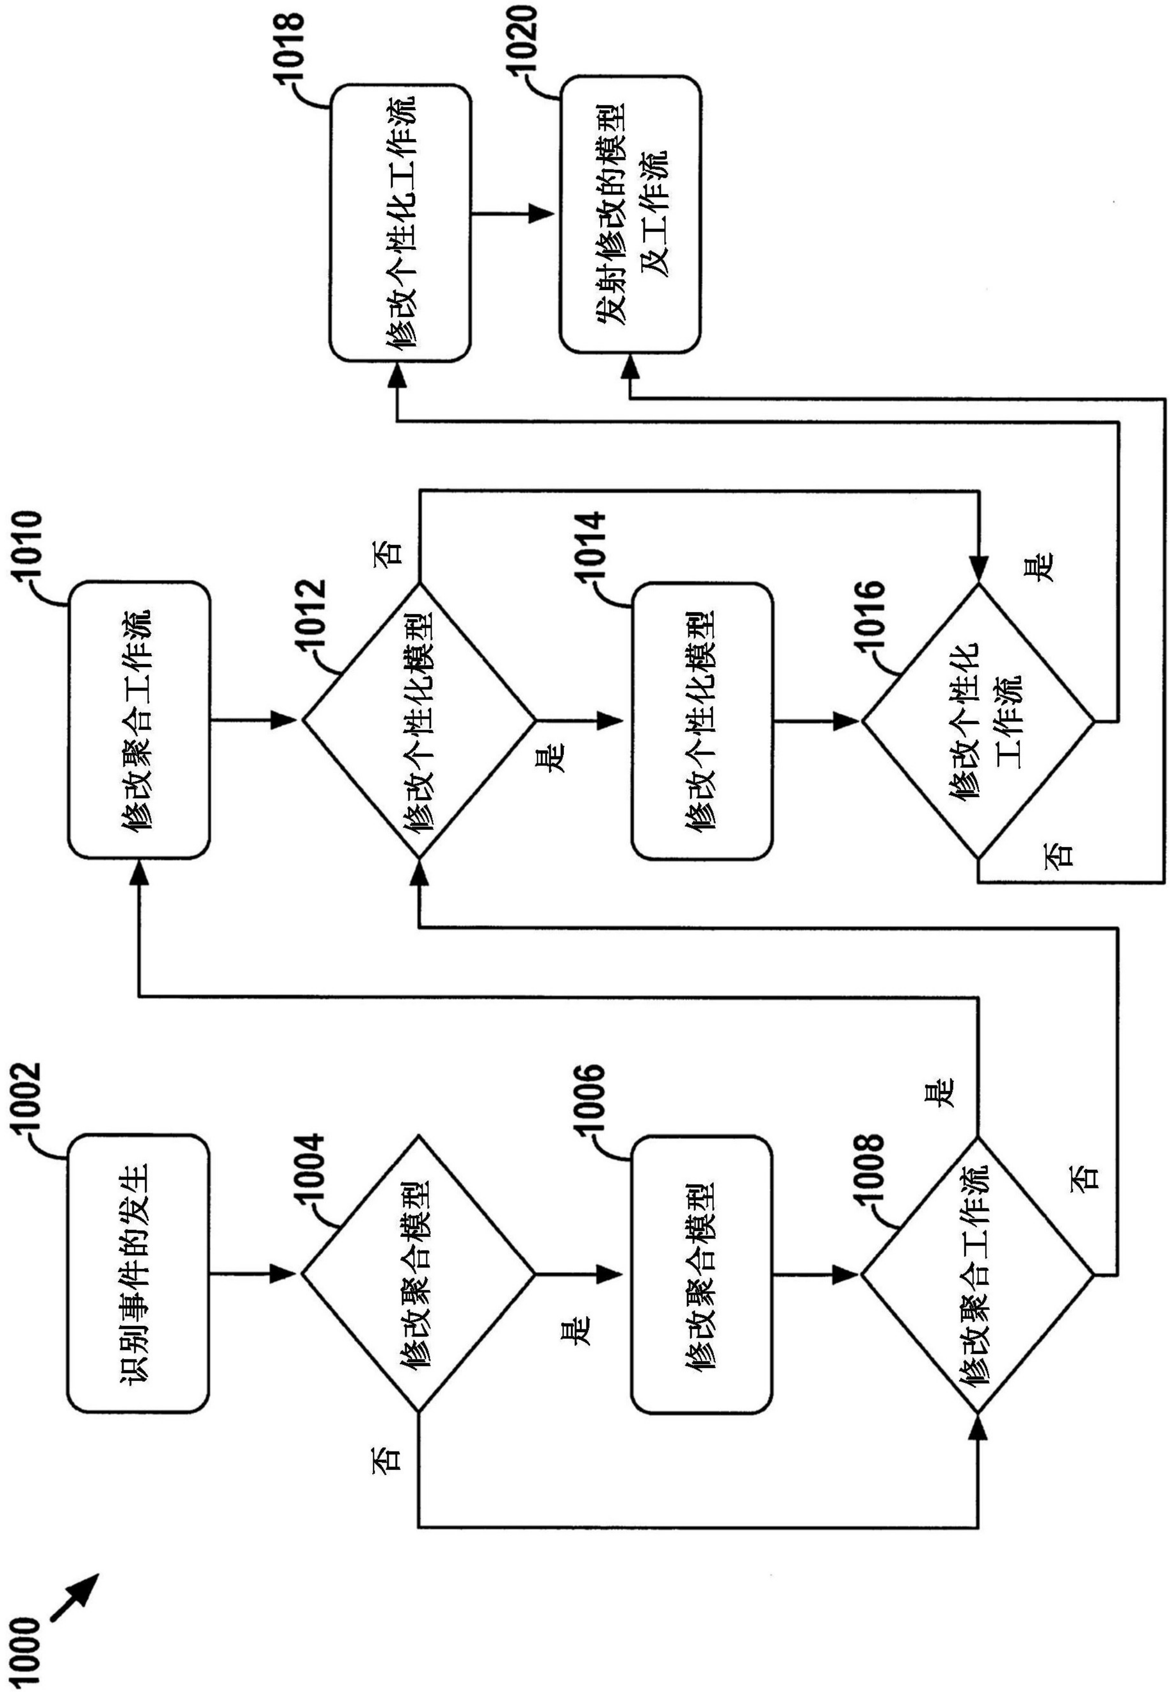
1000
1002
识别事件的发生
1004
修改聚合模型
否
是
1006
修改聚合模型
1008
修改聚合工作流
是
否
1010
修改聚合工作流
1012
修改个性化模型
否
是
1014
修改个性化模型
1016
修改个性化工作流
是
否
1018
修改个性化工作流
1020
发射修改的模型及工作流

图10

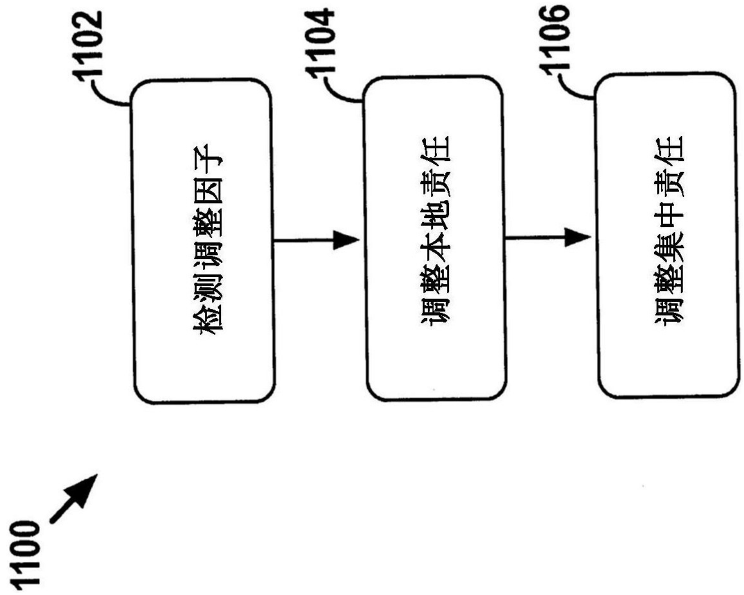

图11

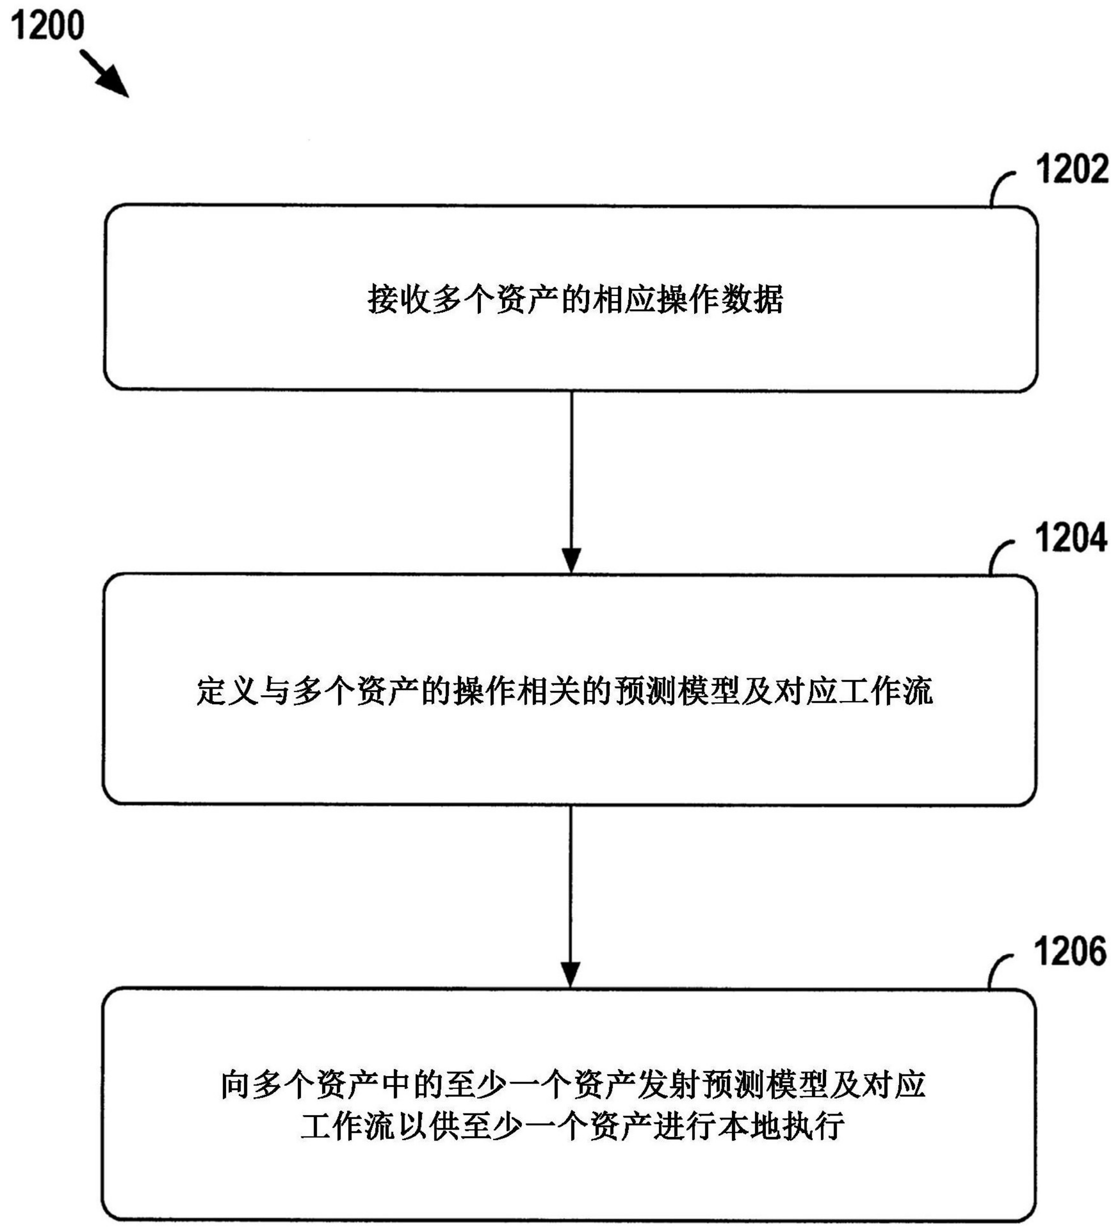
1200
1202
接收多个资产的相应操作数据
1204
定义与多个资产的操作相关的预测模型及对应工作流
1206
向多个资产中的至少一个资产发射预测模型及对应工作流以供至少一个资产进行本地执行

图12

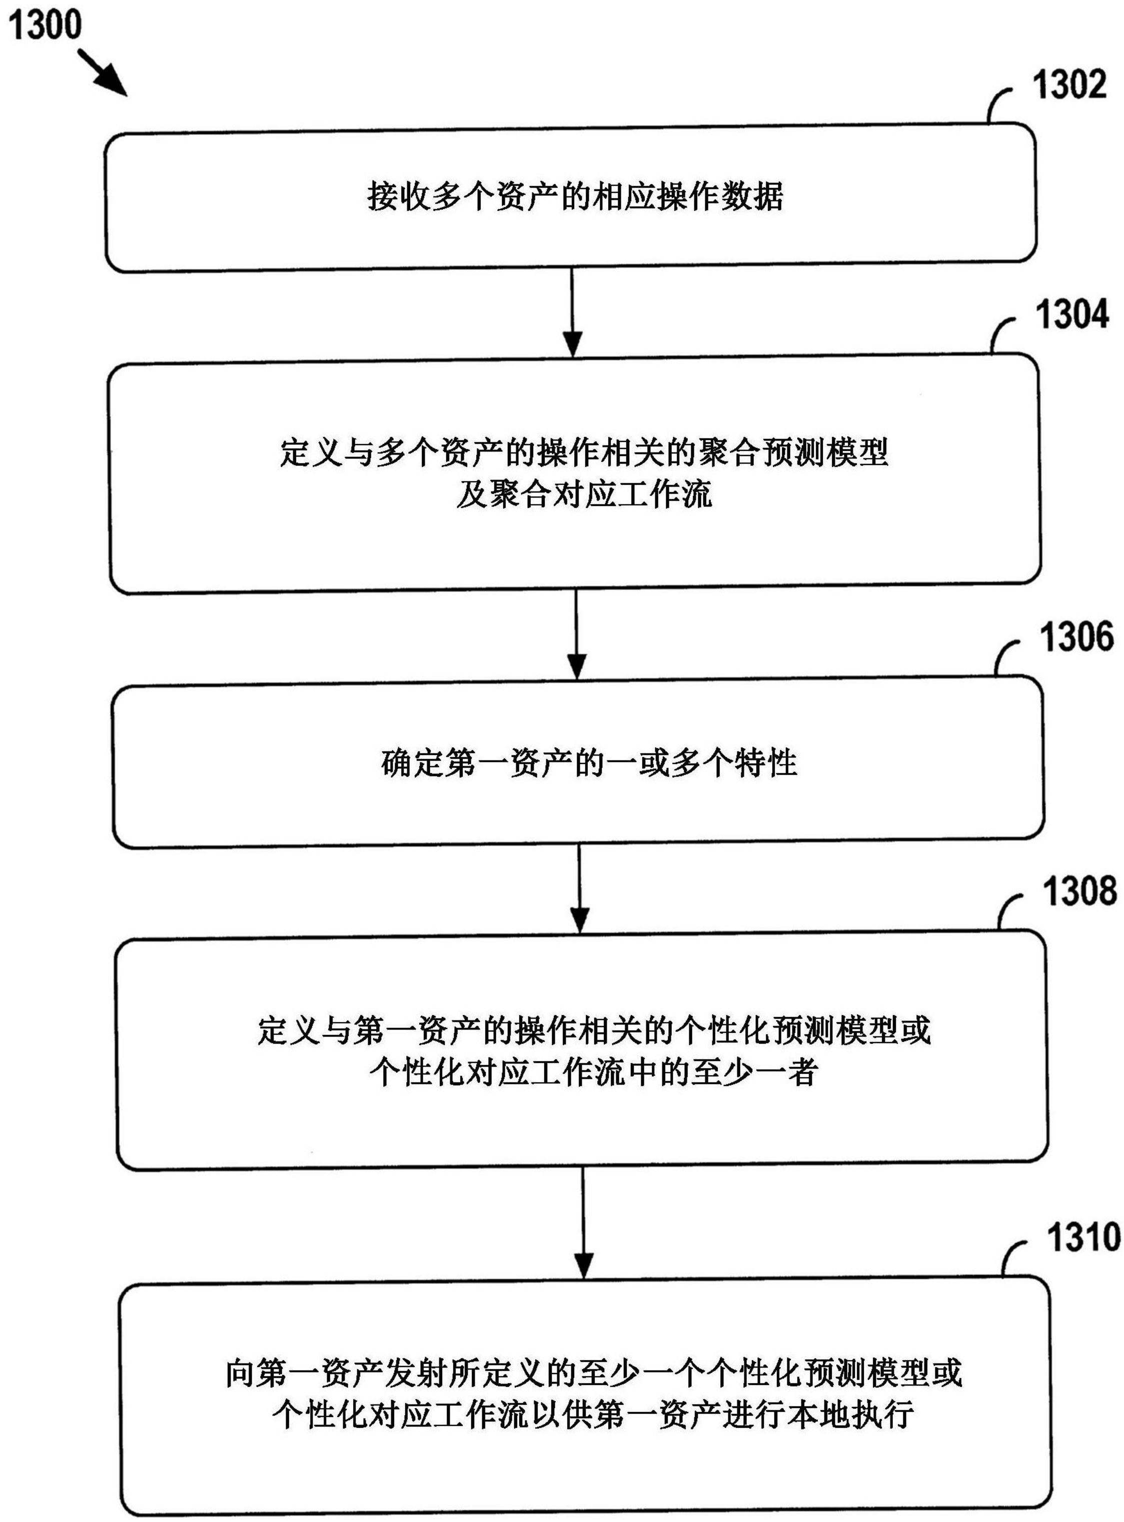
1300
1302
接收多个资产的相应操作数据
1304
定义与多个资产的操作相关的聚合预测模型
及聚合对应工作流
1306
确定第一资产的一或多个特性
1308
定义与第一资产的操作相关的个性化预测模型或
个性化对应工作流中的至少一者
1310
向第一资产发射所定义的至少一个个性化预测模型或
个性化对应工作流以供第一资产进行本地执行

图13

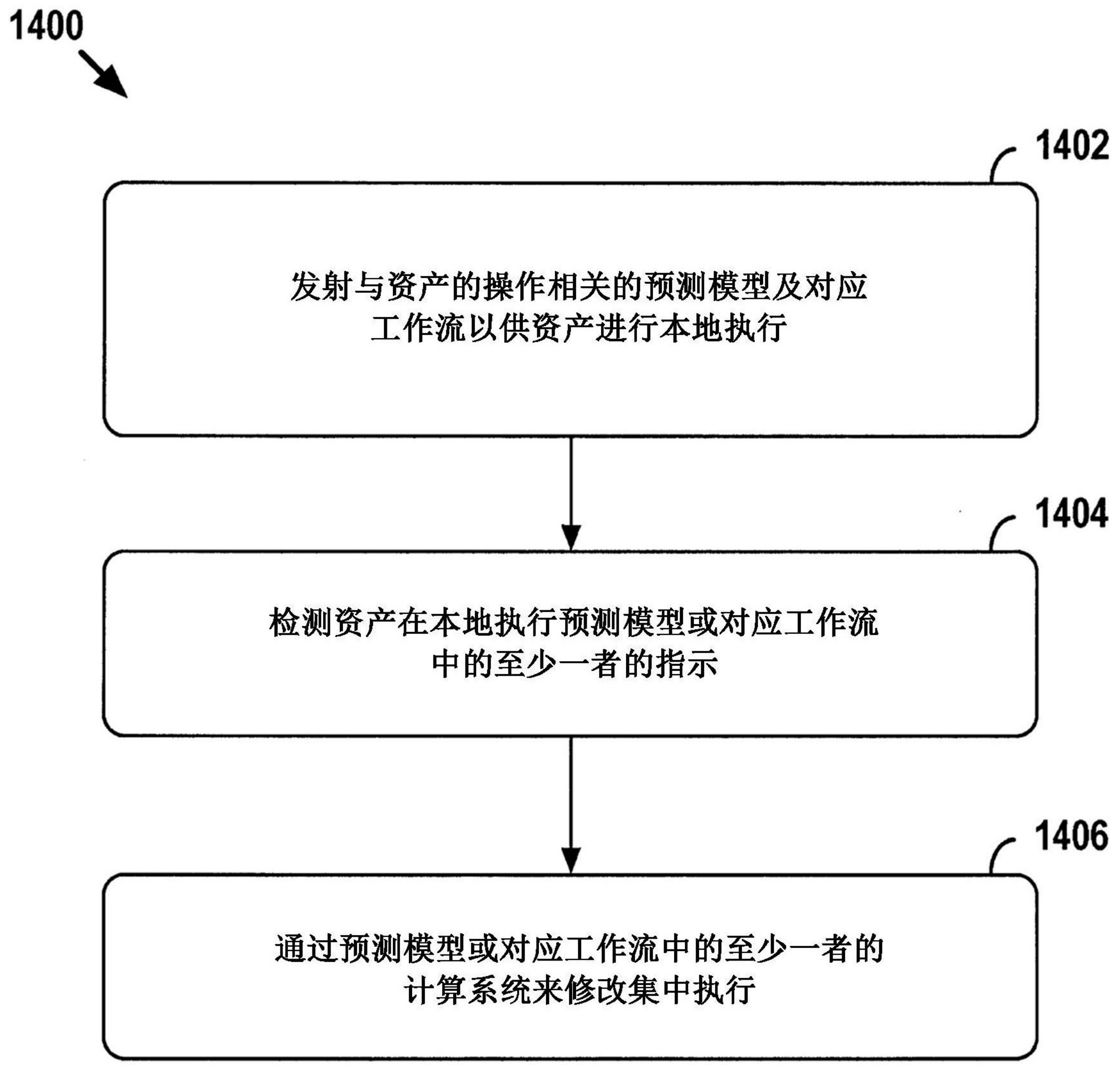
1400
1402
发射与资产的操作相关的预测模型及对应工作流以供资产进行本地执行
1404
检测资产在本地执行预测模型或对应工作流中的至少一者的指示
1406
通过预测模型或对应工作流中的至少一者的计算系统来修改集中执行

图14

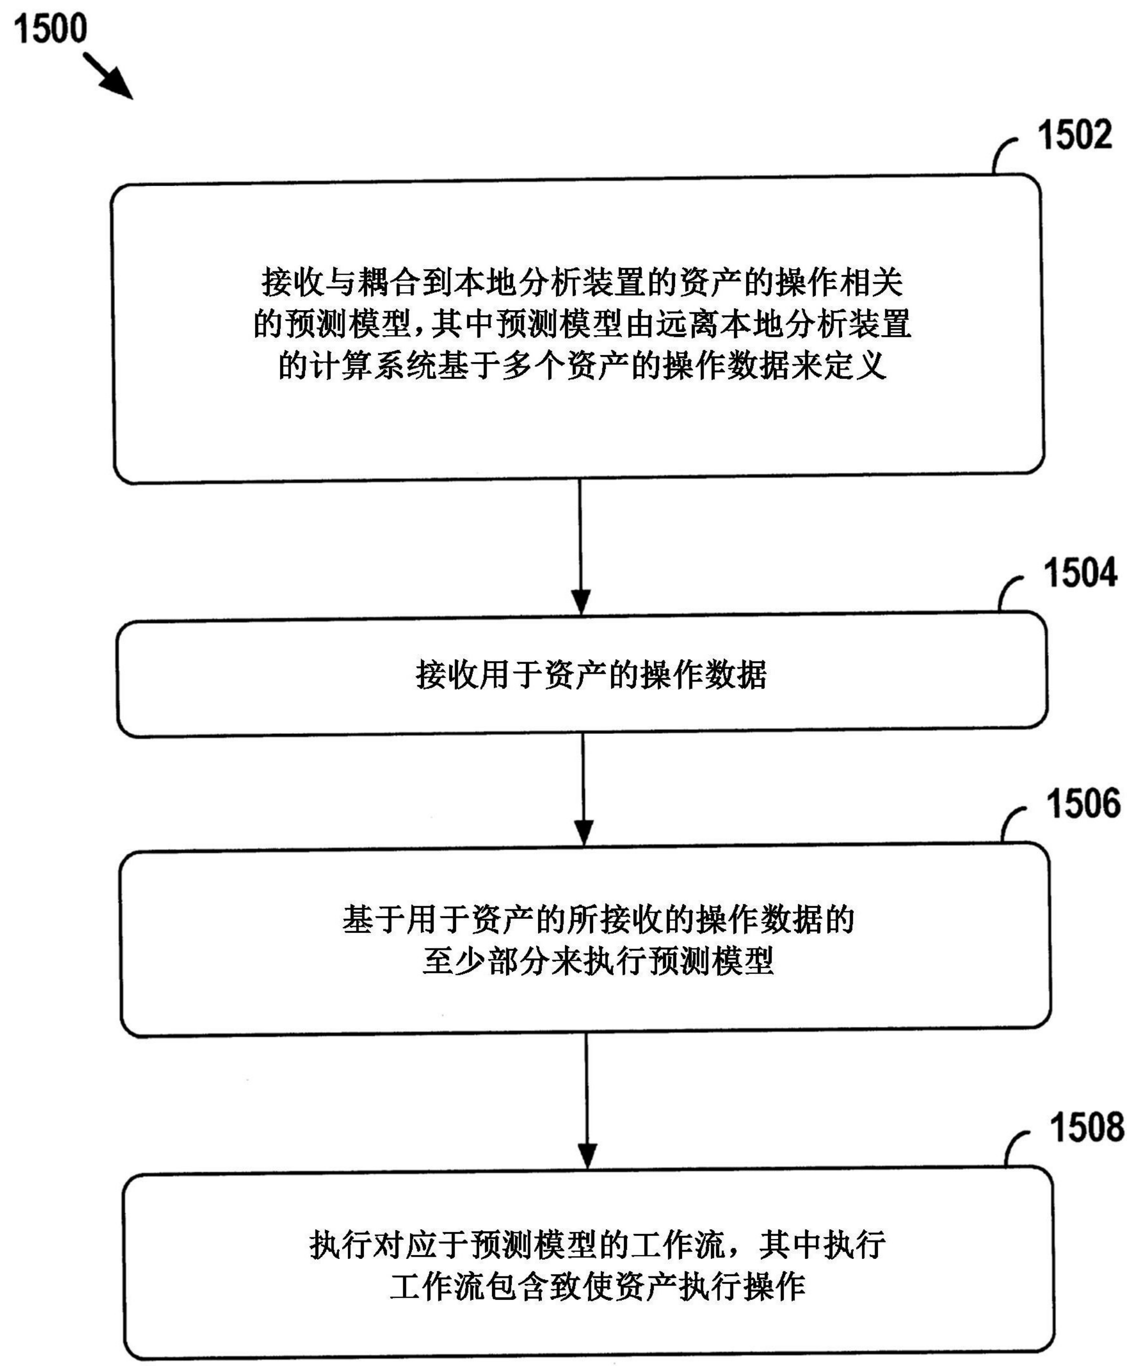
1500
1502
接收与耦合到本地分析装置的资产的操作相关的预测模型，其中预测模型由远离本地分析装置的计算系统基于多个资产的操作数据来定义
1504
接收用于资产的操作数据
1506
基于用于资产的所接收的操作数据的至少部分来执行预测模型
1508
执行对应于预测模型的工作流，其中执行工作流包含致使资产执行操作

图15