



(51) International Patent Classification:

G06F 3/0484 (2013.01) G06N 20/00 (2019.01)

(21) International Application Number:

PCT/US2018/051422

(22) International Filing Date:

18 September 2018 (18.09.2018)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

62/566,026 29 September 2017 (29.09.2017) US
16/007,762 13 June 2018 (13.06.2018) US

(71) Applicant: ORACLE INTERNATIONAL CORPORATION [US/US]; 500 Oracle Parkway M/S 50P7, Redwood Shores, California 94065 (US).

(72) Inventors: ROBERTS, Tara U.; 901 Elder Circle, Austin, Texas 78733 (US). PEARMAIN, Michael; 17 Baldwin Crescent, Merrow, Guildford Surrey GU4 7XW (GB). STANESBY, Jonathan; 21 Godwin Close, Wokingham Berkshire RG41 2AH (GB). YORDY, Kevin; 568 Galapago Street, Denver, Colorado 80204 (US). SUTTHAKORN, Karen; 26 Dellands, Overton Hampshire RG25 3LB (GB).

(74) Agent: GAUDRY, Katherine et al.; 1100 Peachtree Street NE, Suite 2800, Mailstop: IP Docketing - 22, Atlanta, Georgia 30309 (US).

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

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

(54) Title: METHODS AND SYSTEMS FOR CONFIGURING COMMUNICATION DECISION TREES BASED ON CONNECTED POSITIONABLE ELEMENTS ON CANVAS

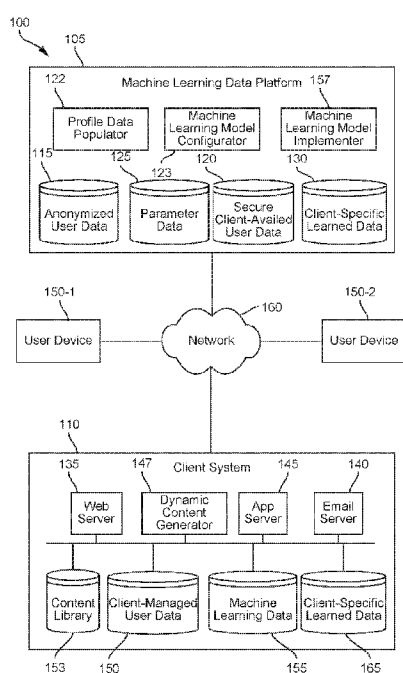


FIG. 1

(57) Abstract: Embodiments are directed at the use of a communication decision tree configured to dynamically define individual trajectories through the communication decision tree using a machine-learning technique to indicate a series of communication specifications, wherein learned data is generated by processing user data using a machine-learning technique. Content is transmitted to a user device in accordance with communication specifications identified based on the learned data and user attributes. At the same time, an interface is provided comprising a canvas on which the decision tree can be defined.

Published:

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

METHODS AND SYSTEMS FOR CONFIGURING COMMUNICATION DECISION TREES BASED ON CONNECTED POSITIONABLE ELEMENTS ON CANVAS

5 CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application claims priority to U.S. Non-Provisional Application Number 16/007,762 filed June 13, 2018, which claims the benefit of and the priority to U.S. Provisional Application Number 62/566,026, filed on September 29, 2017. Each of these applications is hereby incorporated by reference in its entirety for all purposes.

10 FIELD

[0002] Embodiments relate to configuring artificial-intelligence (AI) decision nodes throughout a communication decision tree. The decision nodes can support successive iteration of AI models to dynamically define iteration data that corresponds to a trajectory through the tree.

15 BACKGROUND

[0003] Technological advancements have improved the accessibility and complexity of multiple types of communication channels. Further, data-storage and network advancements have increased capacities, such that an increasing amount (and variety) of data can be stored at a data source for potential transmission. Therefore, a data source can be positioned to deliver many
20 types of data across any of multiple data channels at many potential times. The array of content-delivery options explodes when considering multiple, related content deliveries instead of a single distribution. Frequently, a content provider configures one or more static rules to indiscriminately provide the same content through a same communication channel to each data ingester. While the communication specification(s) may differ across receipt of different data
25 requests, the rule(s) can be configured to indiscriminately and consistently respond to data requests. Though this approach provides configuration simplicity and deterministic operation, it

fails to react to the potential variability across a population of data ingesters and thus may sub-optimally handle requests.

SUMMARY

[0004] In some embodiments, a computer-implemented method is provided. An interface is
5 availed for defining communication decision trees through which trajectories corresponding to
communication exchanges are extended. The interface includes a set of positionable visual
elements that includes a set of action-defining visual elements. Each action-defining visual
element of the set of action-defining visual elements represents a particular action that is to be
performed when a given trajectory has extended to the action-defining visual element. The set of
10 action-defining visual elements includes a switch visual element that represents a decision action
to identify a communication specification using a machine-learning technique. The set of action-
defining visual elements also includes a set of communication visual elements, each of the set of
communication visual elements representing a particular communication specification. The set of
positionable visual elements also includes a connection visual element configured to
15 directionally connect multiple positioned visual elements. Each positioned visual element of the
multiple positioned visual elements corresponds to an action-defining visual element of the set of
action-defining visual elements. The directional connection indicates an order in which particular
actions represented by the multiple positioned visual elements are to occur. The interface also
includes a canvas configured to accept positioning of one or more of the set of positionable
20 visual elements. An update to the canvas is detected. The updated canvas includes the switch
visual element positioned at a first position and a first communication visual element of the set of
communication visual elements positioned at a second position. The first communication visual
element represents a first particular communication specification. The updated canvas also
includes a second communication visual element of the set of communication visual elements
25 positioned at a third position. The second communication visual element represents a second
particular communication specification. The updated canvas further includes a set of connection
visual elements. Each of the set of connection visual elements is an instance of the connection
visual element. A first connection of the set of connection visual elements connects the switch
visual element to the first communication visual element. A second connection of the set of
30 connection visual elements connects the switch visual element to the second communication

visual element. The set of connection visual elements indicates that potential results of execution of the machine-learning technique at the switch visual element include a first result that triggers a communication transmission having the first particular communication specification and a second result that triggers a communication transmission having the second particular communication specification. A particular communication decision tree is defined based on the updated canvas that, when implemented and in response to detecting that a given trajectory associated with particular profile data has extended to a particular decision action represented by the switch visual element, triggers using the machine-learning technique and learned parameter data to process the particular profile data to produce a machine-learning result indicating which one of the first and second particular communication specifications is to be applied for a content transmission. The implementation further triggers transmitting content to a user device associated with the trajectory, the transmission being in accordance with the one of the first and second particular communication specifications as indicated in the machine-learning result.

[0005] In some embodiments, a computer-program product is provided that is tangibly embodied in a non-transitory machine-readable storage medium. The computer-program product can include instructions configured to cause one or more data processors to perform operations of part or all of one or more methods disclosed herein.

[0006] In some embodiments, a system is provided that includes one or more data processors and a non-transitory computer readable storage medium containing instructions which, when executed on the one or more data processors, cause the one or more data processors to perform operations of part or all of one or more methods disclosed herein.

BRIEF DESCRIPTION OF THE DRAWINGS

[0007] Illustrative embodiments of the present invention are described in detail below with reference to the following drawing figures:

[0008] FIG. 1 shows a block diagram of an interaction system.

[0009] FIG. 2 shows a template to be used for generating an email communication.

[0010] FIG. 3 shows a template to be used for generating an app notification communication.

[0011] FIG. 4 shows a representation of a communication decision tree.

[0012] FIG. 5 illustrates an example of a trajectory corresponding to a user device and extending through a communication decision tree.

[0013] FIG. 6 shows an exemplary interface to define a communication decision tree.

5 [0014] FIG. 7 shows an exemplary parameter-defining interface for a switch icon.

[0015] FIG. 8 shows another exemplary parameter-defining interface that includes options to effect a bias towards or against representing various content in communications.

[0016] FIG. 9 shows another exemplary parameter-defining interface that includes options to effect a bias towards or against using various communication channels to transmit
10 communications.

[0017] FIG. 10 shows a flowchart for a process for using machine learning to direct trajectories through a communication decision tree according to some embodiments of the invention.

[0018] FIG. 11 shows a flowchart for a process for defining a machine-learning-based
15 communication decision tree using an interface supporting positionable visual elements.

DESCRIPTION

[0019] In some embodiments, systems and methods are provided that avail a canvas to facilitate configuring a sequence of machine-learning implementations to partly define a communication exchange. More specifically, a canvas is availed that accepts positioning and
20 connecting of individual switch visual elements with corresponding sets of communication visual elements. A communication decision tree can then be generated based on a set of positioned and connected visual elements. The canvas can be configured to accept, for each communication visual element, an identification of one or more communication specifications associated with the communication visual element. Each switch visual element can represent a
25 machine-learning technique (to be associated with particular parameters learned through training) to be used to select a particular communication visual element from amongst a set of communication visual elements connected to the switch visual element. The canvas can be

configured to accept (e.g., for each switch visual element or generally) an identification of a target outcome (e.g., representing a user-initiated event or communication), which can direct machine-learning selections. Thus, the particular communication visual element selected using the machine-learning technique can correspond to a communication specification predicted to be relatively likely to result the target outcome (e.g., which may be represented as an event visual element in the communication decision tree).

[0020] A machine-learning model can be defined for each represented switch visual element. The machine-learning model can be trained using previous trajectories pertaining to other communication decision trees (e.g., but capitalizing on the other communication decision trees having communication visual elements that correspond to similar or same communication specifications as those represented by communication visual elements in the model being trained). For example, the model can be trained by determining – for the trajectories routed so as to trigger a communication having a particular communication specification – what subsequent user-initiated events were represented by those trajectories (e.g., and what portion of the trajectories represented an occurrence of a client-identified target outcome). The model can further or alternatively be trained using trajectories as they emerge that pertain to the generated communication decision tree.

[0021] In some instances, a model can be trained using a data set that reflects previous events (e.g., trajectories through a same or different communication decision tree and/or other indication of an event sequence) and is augmented with new data. The new data may have recently become available (e.g., via newly received form input or metadata detection) but may correspond to a variable type estimated to be static or predictably changing. For example, if a user's age is identified at time X , the user's age at time $X-3\text{ years}$ can be calculated, while an accuracy of a retrospective estimate of an interest or location variable over an extended time period may be less reliable. The training can then determine whether various attributes represented in the new data was predictive of whether particular events were going to occur.

[0022] The interface can be configured to accept indications as to biases that are to be applied at various machine-learning stages. For example, with respect to a given switch element that is connected to a particular set of communication visual elements, a client may interact with a slider control visually associated with a first visual element to indicate that path selections are to

be boosted towards (or dampened from) the first visual element. Metadata that feeds into the machine-learning model can be set based on the interaction to enable effecting of a corresponding bias. In some instances, the metadata can correspond to an unlearned hyperparameter that is then used to adjust or constrain a learned parameter (e.g., weight). In some instances, the metadata be used to define a post-processing adjustment to be applied to a result generated by the machine-learning model. In some instances, a client or system implements a bias towards a given communication visual element when training data corresponding to a communication specification represented by the element is relatively low (e.g., generally and/or in association with a given communication stage).

[0023] A communication decision tree generated based on visual-element configurations indicated using a canvas can include a set of nodes. Each node can correspond to (for example) a detected event or a branching node that corresponds to a communication-specification decision and that is connected to multiple next nodes representing a communication corresponding to one or more particular communications specifications. Each individual trajectory through the communication decision tree can correspond to an individual user and/or one or more particular user devices. Each individual trajectory can extend across a particular subset of the set of nodes, where nodes in the subset are representative of specific actions initiated by the user and/or initiated at a particular user device of the one or more particular devices, specific characteristics of a communication transmitted to the a particular user device of the one or more particular devices; and/or a decision to be made as to a specification of an upcoming communication. For example, a specification of an upcoming communication can indicate when it is to be transmitted, to which device it is to be transmitted, over which type of communication channel it is to be transmitted, and/or what type of content it is to include. In some instances, natural language processing can be used to assign one or more categories to each of one or more content objects transmitted in a training set and/or to each of one or more content objects available for transmission. A communication specification may then identify a particular category of content to be transmitted.

[0024] Each communication-specification decision can be made based on current profile data corresponding to the user and/or particular user devices, a machine-learning model and one or more learned parameters for the machine-learning model. The parameter(s) can be learned based

on user data associated with a set of other users and that indicates, for each of the set of other users, one or more attributes of the other user and/or a set of events (e.g., user-initiated actions or characteristics of communications transmitted to the user). The parameter(s) can further be learned based on a trajectory target (e.g., identified by a client) that corresponds to a particular node within the communication decision tree and/or a particular action initiated by the user.

[0025] A communication decision tree can be configured to include multiple branching nodes (e.g., each corresponding to a switch visual element), such that multiple communication-specifications decisions may be made for a single trajectory. In some instances, each decision is made using a same type of machine-learning algorithm (e.g., a supervised regression algorithm).

However, the algorithm may be differentially configured at each of the branching nodes, such that the branching nodes differ with respect to (for example) the types of input processed for each trajectory and/or the learned parameters to be used to process input corresponding to a trajectory. In various instances, the algorithms for different branching nodes may be trained to optimize a same or different variable (e.g., based on an identification of a same target node or different target nodes). Not only may the branching nodes vary with regard to the types of input that the algorithm is configured to process, but additionally the types of profile data available to potentially be processed for a given user can vary (e.g., profile data may accumulate over time due to interaction monitoring). Further, the learned data associated with any given node can change in time (due to continuous and/or repeated learning).

[0026] As one example, a trajectory for a user can be initialized upon detecting that profile data corresponding to the user includes information for at least a predefined set of fields. The profile data can be collected using one or more web servers over one or more sessions associated with the user and/or retrieved from a remote data source. In some instances, a user device automatically detects at least some of the profile data and communicates it to the web server(s) (e.g., via automatically populated header information in a communication that identifies, for example, a unique device identifier, MAC address, browser type, browser version, operating system type, operating system version, device type, language to which the device is set, etc.). In some instances, a communication includes data that represents user input (e.g., text entered into a web form, link selections, page navigation, etc.), which can then be logged as profile data.

[0027] Initializing the trajectory can include identifying a first node within the communication decision tree, which can include a first branching node. The first decision node can correspond to a decision as to identify which of multiple content objects (e.g., identifying various groups of items and/or information associated with a web site) to transmit within an email communication to a device of the user. The first decision node can also correspond to a decision as to when – within a two-day period – to send the email. The decisions can be made based on first learned data that indicates – for particular types of users – what types of objects and/or communication times are most likely to lead to a target outcome. For example, the target outcome can include an occurrence where the user activating a link within the email to access a page on the web site and/or the user interacting with the web site in a manner that corresponds to a conversion (e.g., a purchase of an item represented on the web site), and the first learned data can indicate that predictive factors as to which of three content objects will be more effective at resulting in the target outcome include whether a user most frequently uses a mobile device (versus a laptop or computer), the user's age, and previous email-interaction indications as to for which types of content objects the user clicked on a link.

[0028] Once the email is sent, the trajectory can extend to a node representing the transmitted content until a next event is detected. The next event can include (for example) activation of a link within the web site, thereby indicating that the user is engaged in a current session with the web site. Upon detecting this event, the trajectory can extend to a second decision node to determine how to configure a requested web page on the web site (e.g., in terms of whether to include dynamic content objects and/or how to arrange various content objects). In this example, second learned data indicates – for particular types of users – what configurations are most likely to lead to a same target outcome. For example, the second learned data can indicate that predictive factors as to which of four configurations will be more effective at resulting in the target outcome include whether a user most frequently uses a mobile device (versus a laptop or computer), a browser type, a current location of a user device, and a current time of day at the user's location. Once the webpage (configured in accordance with the decision made at the second decision node) is sent, the trajectory can extend to a node representing the configuration of the transmitted webpage. The trajectory can continue to extend upon detecting various user-

initiated, system-initiated or external events (e.g., a passing of a predefined time interval since a previous event).

[0029] In this example, the target outcome remains the same across multiple decisions.

However, rather than identifying a static workflow of actions to perform – or even rather than
5 determining a user-specific complete sequence of actions to perform – techniques disclosed
herein base decisions pertaining to individual actions on current profile data, current learned data
and current event detections. Machine learning is executed iteratively throughout a life cycle of a
particular trajectory to identify piecemeal actions to be performed. This approach can facilitate
high utilization of data (e.g., as expansions and/or evolutions of learned data and/or profile data
10 can be utilized in mid-process decisions), which can promote achieving the target objective.
Further, the approach enables a change (e.g., initiated by a client) to a definition and/or
constraint of a machine-learning technique to take a quick effect (e.g., as the change can still
influence trajectories having already been initiated). For example, a client may change a target
objective from conversion to retaining a user device on a web site for at least a threshold session
15 duration. Modified parameters for machine-learning models associated with various branching
nodes can be determined and immediately effected, so as to affect previously initiated
trajectories that subsequent reach the node(s).

[0030] Communication decisions (and/or partial directing through a communication decision

tree) can be based on anonymized or partially anonymized data, either or both of which can be
20 built from anonymized, partially anonymized or non-anonymized data provided by one or more
providers or clients. For example, a remote user-data management system can receive partially

anonymized or non-anonymized data from one or more data providers and can obscure or
eliminate fields in individual records according to data-privacy rules and/or can aggregate field
values across a user sub-population to comply with data-privacy rules. As described herein,

25 anonymized or partially anonymized data is data that has been stripped of PII and/or aggregated
such that individual data values cannot be, beyond a certain probability, associated with
particular people or users. Thus, the anonymized or partially anonymized data can lack or
obscure sufficient data values to prevent identifying a particular person as being a particular user
or to prevent identifying a particular person as having at least a threshold probability as being a
30 user. For example, the anonymized or partially anonymized data may lack a name, email address,

IP address, physical address and/or phone number from profile data. The anonymized or partially anonymized data may include or exclude certain demographic data, such as an age, city, occupation, etc. In some instances, anonymized or partially anonymized data is useful to collect so as to comply with privacy rules, regulations and laws, while also being able to process some of the data. The anonymized or partially anonymized data can include information gathered from devices based on IP address range, zip code, date, categories of prior online inquiry, race, gender, age, purchase history, and/or browsing history, etc., which may have been gathered according to a variety of privacy policies and laws that restrict the flow of personally identifiable information (PII), etc.

[0031] In some instances, the anonymized or partially anonymized data is used to generate and/or update learned data (e.g., one or more parameters) associated with individual machine-learning configurations. This type of training need not (for example) require or benefit from data fields such as contact information, so data records can be stripped of these fields. As another example, one or more sub-populations can be generated based on values for a particular field, and specific values for that field may thereafter be replaced with an identifier of a sub-population.

[0032] In some instances, profile data corresponding to a particular user for which decisions are being made include the anonymized or partially anonymized data. For example, a system can detect that a trajectory has reached a branching node and request data from a user-data management system (e.g., using a device identifier or other identifier associated with the trajectory). The system can return profile data that includes (for example) one or more specific non-anonymized field values, one or more field values that have been generalized (e.g., assigned to a category), and/or eliminated field values. The non-anonymized field data may be included in profile data when (for example) such field values were supplied by (e.g., upon having been collected using data-collecting features built into a webpage and/or via a transmission from the client) or otherwise accessible to (e.g., via a data-sharing agreement) a client for which a decision is being made. The system may also return population data (e.g., that can itself be learned and/or can evolve over time) that indicates relationships between field values, which may be used to estimate values or categories for missing field values.

[0033] FIG. 1 shows a block diagram of an interaction system 100. A machine learning data platform 105 can include one or more cloud servers and can be configured to receive user data from one or more client systems 105. The user data can include anonymized or partially anonymized user data (stored in an anonymized user data store 115) and/or secure client-availed user data (stored in a secure client-availed user data store 120), which may be less anonymized than anonymized user data or not anonymized. When secure client-availed user data is received, it may be securely stored in association with an identifier of a client, such that other clients cannot gain access to the data. The data may be stored in a multi-tenant cloud storage system such that multiple clients can log in to a central location to access a server or collection of servers, but where the specific access to data is controlled depending on which client has authenticated to the cloud storage system. Anonymized or partially anonymized user data may, or may not, be particularly configured for various clients (e.g., depending on which data the client supplied and/or data-sharing agreements associated with the client). Thus, a profile data populator 122 at machine learning data platform 105 can generate profile data corresponding to one or more individual users for particular clients and can customize which field values are included in the profile data for individual clients.

[0034] In some instances, profile data populator 122 enhances a profile data set to supplement client-availed user data with partially anonymized user data, which can (as aggregated) define client-specific learned data (stored in a client-specific learned data store 130) for a given user. For example, data from a profile in the client-availed data can be mapped to one or more data sets in the anonymized or partially anonymized user data, such that richer data sets can be used in the machine-learning analyses. The mapping may occur using overlapping data (e.g., an IP address, if included in the anonymized or partially anonymized user data, a purchase time, a pseudo-random user identifier assigned by a client, etc.).

[0035] A machine learning model configurer 123 can configure a given machine learning model based on (for example) an identified target outcome, available training data, one or more client-identified constraints and/or potential actions as indicated by a communication decision tree and/or by a client. Configuring the machine learning model can include defining one or more parameters for a particular instance of the model (e.g., the instance being associated with a particular branching node, client and/or time period).

[0036] Each parameter can be indicative of a relationships and/or correlation between user attributes (stored in a learned parameter data store 125). The parameter(s) can include a weight that indicates how and/or an extent to which a first user attribute is predictive of a second user attribute that corresponds to an indication as to whether and/or an extent to which a target outcome occurred. The weight may be defined along a discrete or continuous value range and/or can be binary.

[0037] As one example, the parameter(s) can indicate which attributes from amongst a set of attributes are predictive of future occurrence of a particular type of conversion event. For example, it may be determined that having visited a webpage associated with a “travel” tag more than twice in the last month was a predictor of buying a piece of luggage. As another example, it may be determined that having visited a movie-review webpage within a given day was a predictor for later purchasing an online rental of a movie. Indirect associations and trends may also be learned, such as identifying there is an inverse correlation between an age of the user and an average time spent online each day. Each parameter may be associated with a strength and/or confidence of a relationship, optionally with serial associations between the data points gathered and the conclusions being made, where each association in serial carries a certain probability that the data at the start of the association is accurate for what it says and a certain other probability that the association itself is accurate.

[0038] The configuring may, but need not, be performed using client-availed profile data and/or to produce client-specific parameters. The client-specific parameter(s) may be, for example, a modified version of the parameter(s) generated using the anonymized or partially anonymized profile data.

[0039] Various machine-learning techniques may be used to generate learned data. For example, a machine-learning technique may use decision-tree learning, association-rule learning, an artificial neural network, deep learning, inductive logic programming, a support vector machine, clustering, a Bayesian network, reinforcement learning, representation learning, similarity and metric learning, sparse dictionary learning, a genetic algorithm, or rule-based machine learning. In some instances, a machine-learning technique includes an ensemble technique, which learns inter-ensemble weights to apply to results produced from various underlying techniques (such as two or more of those previously mentioned). The inter-ensemble

weights may be identified based on (for example) accuracy, speed and/or resource usage associated with the underlying techniques.

[0040] Training a machine-learning technique (to identify one or more parameters) can include identifying how a set of observed inputs (e.g., content of a marketing email, content of a promotion, and/or the configuration of a web site) relates to a set of corresponding outputs (e.g., an outcome, such as the presence or absence of certain conversion event, for a corresponding marketing email, a corresponding promotion, and/or a corresponding web site configuration). These observed observations can be used to identify modeled relationships and/or trends, with a goal of predicting candidate factual information (e.g., a predicted next input to be received or a predicted output based on certain inputs) that has not yet occurred based on factual information leading up to the candidate factual information. Each prediction can carry a confidence or probability, and chains of predictions have a combined confidence or probability.

[0041] Thus, machine learning model configurator 123 can identify model parameters for particular client systems 110 based on (for example) target outcomes, client-specific profile data and/or machine-learning techniques. Client-specific learned data can be selectively shared with a client system having provided the underlying client-availed profile data. Client system 110 can include a system that hosts one or more web sites, hosts one or more apps and/or causes emails to be transmitted. For example, client system 110 can include a web server 135 that receives and responds to HTTP requests for pages on one or more domains and an email server 140 that delivers emails to users' email addresses. Client system 110 may further or alternatively include an app server 145 to receive and respond to requests received via an application executing on a user device. Thus, one or more servers at client system 110 can be configured to detect requests from one or more user devices 150-1, 150-2 and/or trigger transmission of content to one or more user devices 150-1, 150-2. User devices 150-1, 150-2 may include, for example, a computer, smart phone, tablet, etc. It will be appreciated that, in various circumstances, a single user device may be associated with a single user or more than one users. Further, a single user may be associated with a single user device or more than one user devices.

[0042] Web server 135 and/or app server 145 may store indications of requests for content (e.g., a webpage or app page) from a content library 153 as user data in a client-managed user data store 150. The stored data may include automatically detected information (e.g., a request

time) along with information included in the request (e.g., a device identifier, IP address, requested webpage, user-entered input, etc.). Storing the data may include updating a profile to include the data. Web server 135, email server 140 and/or app server 145 may further store data in client-managed user data store 150 that indicates which content was distributed to particular user devices (e.g., by identifying a transmission time, user-device identifier, content-object identifier(s), and/or type of communication).

[0043] Client system 110 can transmit at least part of the user data from client-managed user data store 150 to machine learning data platform 105, which can store it in secure client-availed user data store 120. The transmission(s) may occur periodically, during a request for client-specific learned data, at regular time intervals, etc. In some instances, client system 110 at least partly anonymizes some or all of the user data (e.g., by omitting or obscuring values for at least some fields) before transmitting it to machine learning data platform (e.g., such that it is stored as anonymized or partially anonymized user data at the platform). In some instances, the data is not at least partly anonymized, such that the data is either stored in secure client-availed user data store 120 or is at least partially anonymized at machine learning data platform 105. For some datasets, the anonymized or partially anonymized data is received from a third party, after being stripped of PII, and stored by client system 110 without ever having access to the non-anonymized data. In some embodiments, the anonymized or partially anonymized data is natively anonymized or partially anonymized. In these embodiments, websites may run embed scripts on their web sites that, when executed, gather anonymized or partially anonymized data about accesses of the web sites by users. The scripts may gather only information that may be gleaned without knowing a user's personal information and stored in a data cloud that ensures that user identity can never be deduced beyond a certain probability.

[0044] Client system 110 can store machine-learning data in a machine learning data store 155. In some instances, the machine learning data includes an indication of one or more decisions made at a branching node for a given trajectory, one or more content specifications identified using a communication decision tree and/or one or more parameters. The machine-learning data can be requested from, received from and/or derived from data from machine learning platform 105. For example, in some instances, machine learning model configurator 123 causes parameters generated for and/or applicable to a client to be transmitted to client system 110. As

another example, a machine learning model implementor 157 can apply machine learning model configured with particular parameters to particular profile data to identify one or more particular communications specifications to define a communications action to be taken for a client (and/or a next node of a trajectory) that corresponds to the profile data. Machine learning model

5 implementor 157 can then cause an indication of the identified communications action and/or the next node to be transmitted in association with an identifier of a trajectory, user and/or user device.

[0045] Identifying a next node and/or communications specification(s) can include running a machine learning model (associated with a current branching node) using particular profile data
10 and one or more learned parameters. A result can indicate (for example) which of various content-presentation characteristics is associated with a high (e.g., above-threshold) or highest probability of leading to a particular target outcome (e.g., target conversion). In some instances, the analysis includes identifying one or more content-presentation characteristics associated with a highest probability of leading to a particular conversion target outcome. In some instances, the
15 analysis balances the probabilities of leading to a particular conversion results with a predefined cost metric associated with various content-presentation characteristics.

[0046] In some instances, running the machine learning model using the parameters (e.g., at machine learning data platform 105 or client system 110) can include (for example) performing a regression analysis using the profile data and parameters to generate a number that can be
20 compared to one or more thresholds. The one or more thresholds can define two or more ranges (e.g., open-ended or closed ranges), with each range corresponding to a particular next node and/or communications action. In some instances, running the machine learning model using the parameters can include processing at least part of the profile data and at least part of the parameters to produce a result that can be compared to (e.g., via calculation of a difference,
25 calculation of a cost using a cost function, etc.) each of a set of reference data variables (e.g., single values, a vector, a matrix, a time series, etc.) – each being associated with a particular next node and/or communications action and each potentially defined at least in part based on a parameter. A node or communication associated with a reference data variable for which the comparison indicated a closest correspondence can be selected.

[0047] A dynamic content generator 147 can trigger a presentation of a content object in accordance with the selected communication specification(s). To generate an appropriate instruction, dynamic content generator 147 may first identify what communication channel is to be used to transmit the object, the type of object that is to be transmitted, a version of content that is to be transmitted and/or when the content object is to be transmitted. The identification can be determined based on (for example) a result of an implementation of a machine learning model, a configuration of a machine learning model (e.g., which may restrain potential options with respect to one or more of these options), and/or one or more parameters.

[0048] Dynamic content generator 147 can identify a type of communication (e.g., email, SMS message, pop-up browser window or pushed app alert) to be transmitted, which can inform (for example) which of web server 135, email server 140 and/or app server 145 is to transmit the communication. The identification can be made explicitly (e.g., based on a machine-learning result, parameter, and/or machine-learning-model configuration) or implicitly (e.g., due to a selected content object being of a particular type).

[0049] Identifying the content object can include selecting from amongst a set of existing content objects or generating a new content object. The content object can include (for example) a webpage, an object within a webpage, an image, a text message, an email, an object within an email and/or text. In some instances, a result of executing a configured machine-learning model on profile data identifies a particular content object. In some instances, a result identifies a characteristic of content (e.g., having a particular metadata category) and/or identifies a particular technique for selecting content. For example, a result may indicate that a “tools” item is to be featured in a content object and/or that a communication is to include four content objects that correspond to four different (though unspecified) categories. In such instances, dynamic content generator 147 can (for example) select from amongst a set of potential content objects using a selection technique that is (for example) indicated via a result of the machine-learning implement, via a parameter, and/or via a predefined setting. For example, a selection technique may indicate that a selection technique is to include a pseudo-random selection technique, a technique to identify a most recently added object, a technique to identify a highest-conversion object within a set of potential content objects (e.g., having one or more attributes as indicated in a machine-learning result).

[0050] In some instances, a time at which a communication is to be transmitted is explicitly identified (e.g., based on a machine-learning result, parameter, and/or machine-learning-model configuration). In some instances, a communication is transmitted immediately, upon receiving a next request for content (e.g., corresponding to a given web site or app) from a user device associated with a machine-learning result, or in accordance with a predefined transmission schedule.

[0051] In some instances, each specification corresponding to a communication is identified (e.g., during a task and/or using a machine-learning model, a machine-learning configuration, a parameter, a client rule, etc.) at or before the communication is transmitted. Thus, all or some client-controlled configuration of the communication and/or its transmission can be performed prior to transmission of the communication. In some instances, at least one specification corresponding to a communication is identified (e.g., during a task and/or using a machine-learning model, a machine-learning configuration, a parameter, a client rule, etc.) after the communication is transmitted. Thus, at least some client-controlled configuration of the communication and/or its transmission can be performed after transmission of the communication. This post-transmission configuration can thus be based upon learned data and/or profile data that was not available prior to the transmission of the communication. For example, additional profile data corresponding to a user may become available between a first time at which an email was transmitted and a second time at which the email is opened and rendered. The transmitted email can include a script that executes when the email is to be rendered. The script can cause a request to be issued to identify device properties, such as a layout and/or application type. The script can pass these properties along with a request for content to be presented to a server. Thus, the server can select content and/or identify one or more display configurations based on specific rendering information, current profile data and/or current parameters to direct a selection of specific content.

[0052] As an additional or alternative example, the communication may contain one or more references or links to pages that, when opened (e.g., in a web browser), render content for display. The pages targeted by the links may include some content that was determined, by the machine learning engine, before or at the time the communication was generated. The pages can further be configured to include content that is to be selected or generated when a request for

rendering the pages is detected (e.g., when a script detects activation of a link) and/or when the pages are being generated or rendered (e.g., as indicated by executing a script as part of loading the page). In some instances, a script in the email identifies the content configuration at the time of rendering or at the time that rendering is requested. In some instances, a script executing on the linked page identifies the content configuration.

[0053] As one example, a client system may offer online purchases of food delivery. It may be detected that a particular user had looked at a menu for a given restaurant at 2pm. The client system may retrieve a set of user attributes from a profile data for the user from its client-managed user data. Client-specific learned data may indicate that there is a 76% chance that the user will make a purchase from the restaurant if an email including a discount code is sent in the evening to the user (e.g., as compared to a lower probability associated with other types of communication and other times). In response to determining that the 76% chance is above a 65% threshold for sending a discount threshold, email server 140 transmits an email to the user device. The email includes a script that, when executed, identifies the restaurant and discount to be presented. The user opens the email the next day at 10am. The code is executed to request the restaurant and discount from the client system. The client system has since received updated public learned correlation data. The client system inputs the time, the user's location (as she is now at work) and prior purchase information to a decision tree built based on the learned data. It is determined that the discount is to be 10% (e.g., to maintain a threshold likelihood of conversion) and the restaurant is to be a deli near the user's work (e.g., to maximize a likelihood of conversion), whereas – had the user opened the email the night before, different user attributes and learned data would have resulted in a 15% discount (e.g., to maintain the threshold likelihood) from an Indian restaurant near the user's home (e.g., to maximize the likelihood). The email includes a link to order from the deli. When the user clicks on the link, the web server determines what content is to be presented – specifically, which food items are to be recommended. The recommendations are based on even more recently updated public learned correlation data, which indicate that salads and sandwiches should be recommended over soup and entrees, as the former options have been recently popular (predicted to be popular due to the warmer weather). Thus, this example illustrates how content presentations can be dynamically customized for a given user based on very recent learned data and user attributes.

[0054] Machine learning data platform 105 can generate updated client data based on (for example) any communications received from a user device (e.g., responsive to a workflow action). For example, the updated client data can include one or more new fields generated based on data in a header or payload of a received communication, an indication as to whether (e.g.,
5 and when) a particular event was detected, and/or a current or final stage of the workflow to which the profile is assigned. Machine learning data platform 105 can avail the updated client data (e.g., along with corresponding profile identifiers) to client system 110, which can store the updated data in client-specific learned data store 165. Client system 110 may, but need not, separately store the updated data from underlying profile(s).

10 [0055] It will be appreciated that, in some instances, some or all of machine learning data platform can be incorporated within client system 110. In some instances, client system 110 communicates with machine learning data platform during iterations of a communication decision tree. For example, client system 110 (e.g., web server 135 or app server 145 at client system 110) may detect a flag (e.g., included in a URL) in a request for web content or app
15 content received from a user device, where the flag indicates its association with a machine-learning-based workflow). Client system 110 may then alert machine learning model implementor 157 of the request, so that a trajectory can be appropriately updated.

[0056] Machine learning data platform, client system 110 and user devices 150-1, 150-2 can communicate over a network 160, which can include, for example, the Internet, a local area
20 network, a wide area network, and so on. It will be appreciated that various alternatives to the depicted and described embodiments are contemplated. For example, some or all of the machine learning may be performed at client system 110. Client system 110 may periodically receive anonymized or partially anonymized user data to process using a machine-learning technique.

[0057] Other techniques for using and configuring communication decision trees are detailed
25 in U.S. Application Number 16/007,677, filed on June 13, 2018 (entitled “Directed Trajectories through Communication Decision Tree using Iterative Artificial Intelligence”), and U.S. Application Number XXXX, filed on June 13, 2018 (entitled “Machine-Learning Based Processing of De-Obfuscated Data for Data Enrichment”). Each of these applications is hereby incorporated by reference in its entirety for all purposes.

[0058] FIGS. 2 and 3 illustrate interfaces 200 and 300 for configuring templates 202 and 302 for communications configured to be partly configured upon detecting a rendering process or at rendering. The configuring can include executing a configured machine-learning model using current learned configurations of the model and current profile data. Template 202 shown in
5 FIG. 2 includes a template to be used for generating an email communication, and template 302 shown in FIG. 3 includes a template to be used for generating an app notification communication.

[0059] Template 202 includes static text (e.g., text 205) and interaction features (e.g., button 210). Template 202 further represents a particular layout, in which three items are to be linearly
10 represented above text 205. Template 202 also include dynamic components (e.g., dynamic text 215 and dynamic image 220) that are configured to be identified when rendering of the email is requested or occurring. Thus, when an email communication is transmitted, the static components can be transmitted along with code configured to (upon detecting a request to render the email) locally identify at least part of current profile data, request at least part of current
15 profile data, request identification of dynamic components, receive or retrieve dynamic components (e.g., identified using current profile data, current anonymized or partially anonymized data and/or current learned parameters) and/or generate a complete email based on the template and dynamic components. The generated email can then be presented.

[0060] Template 302 includes a static layout and multiple dynamic text components (e.g., a
20 dynamic title section 310. Template 302 can be configured to be transmitted with a script that facilitates dynamically identifying each dynamic text component. For example, the script can – upon detecting a request to present the notification (e.g., in response to opening an app, clicking on a notification app element, etc.) – locally identify at least part of current profile data, request at least part of current profile data, request identification of dynamic text components, receive or
25 retrieve dynamic text components (e.g., identified using current profile data, current anonymized or partially anonymized data and/or current learned parameters) and/or generate a complete notification based on the template and dynamic text components. The generated notification can then be presented. Interface 300 shows an example of a dynamically generated notification 315 this includes the static layout and particular dynamic text.

[0061] FIG. 4 shows a representation of a communication decision tree 400. Communication decision tree 400 includes a starting node 405, at which each trajectory begins. A particular trajectory can be (in this example) initialized upon detecting that a user has completed two particular actions (e.g., initialized two web-site sessions, purchased two items from a web site, navigated to at least two webpages on a web site, etc.).

[0062] Communication decision tree 400 includes three branching nodes 410, 415 and 420 – each of which branches to connect to three nodes representing three different actions. A trajectory can automatically and immediately extend from initial node 405 to a first branching node 410, which triggers a first decision to be made. Specifically, the first decision can include identifying a communication channel to use to send an alert of a web-site feature. The alert can include an automatically presented static header that indicates (for example) that a product or discount (generally) is available in association with the web site. The alert may further be associated with dynamic content (e.g., that specifically identifies one or more products and/or a discount) that is to be identified at a second branching node 415 upon detecting a request to open the notification.

[0063] First branching node 410 is connected to a first action node 425a that represents an email communication channel, a second action node 425b that represents an SMS-message communication channel, and a third action node 425c that represents an app-based communication channel (where a notification would be pushed to and/or by an app installed at a user device).

[0064] The first decision can be made using a machine-learning model configured based upon one or more first parameters. The one or more first parameters can be dynamically determined based on anonymized and/or partially anonymized user data and/or client-specific data. For example, anonymized and/or partially anonymized user data may indicate – for each of various user sub-populations (as defined based on one or more user attributes) – how effective an alert transmission sent via each of the three types of communications channels was at triggering the user to initiate a session at a corresponding web site (e.g., as determined based on using tracking links within the alerts) and complete a transaction during the session. The anonymized and/or partially anonymized user data may correspond to many different web sites and/or web sites having one or more particular characteristics. The client-specific data can include data tracked by

a given client for the particular web site of interest and can data that specifically identifies each user to which various alerts were transmitted and the result. The client-specific data may thus be richer in some respects relative to the anonymized and/or partially anonymized data, but the number of users represented in the client-specific data may be smaller than that represented in the anonymized and/or partially anonymized data. Further, the client-specific data may lack pertinent attribute combinations. For example, a given client may not have previously used app-based alerts, which may have reduced an accuracy with which a machine-learning model could predict potential effects of such alerts.

[0065] The machine-learning model (configured with the first parameters) can use profile data associated with the trajectory to determine which communication channel to user. The profile data can include client-collected profile data (e.g., using metadata, cookies and/or inputs associated with previous HTML requests from a user device associated with the trajectory). The profile data may further include other profile data requested and received from a remote user-profile data store, which may collect and manage profile data from multiple web hosts, clients, etc.

[0066] Upon identifying the communication channel, the trajectory extends to the corresponding action node (425a, 425b or 425c). An alert is then sent using the corresponding communication channel. The alert can be configured to automatically identify limited content and to cause the trajectory to extend to second branching node 410 upon detecting a request to open the alert. A decision can then be made at second branching node 410 to determine specific content to be presented in a body of the alert.

[0067] Thus, second branching node 415 is connected to a first notification content node 430a that represents content that identifies a product most recently viewed by the user at the web site, a second notification content node 430b that represents content that identifies four of the products most viewed (across users) at the web site over the last week, and a third notification content node 430c that represents content that includes an identification of a discounts. The second decision can be made using the machine-learning model configured based upon one or more second parameters. Thus, in some (but not all) instances, a general type of machine-learning model used at various branching nodes to make decisions can be the same, though

particular configurations (e.g., indicating weights to be assigned to various user attributes, which user attributes are to be considered at all and/or target outcomes) can differ.

[0068] The one or more second parameters can be dynamically determined based on anonymized and/or partially anonymized user data and/or client-specific data. However, each of the anonymized and/or partially anonymized user data and/or the client-specific data may have changed since making the first decision, which can contribute to differences between the first and second parameters. Further, the potential actions considered at second branching node 415 differs from those considered at first branching node 410. Therefore, the first and second configurations can be different. Additionally, the profile data that is processed can differ between the first and second branching nodes. For example, a client-associated application may have been installed at a user device between processing performed at the first and second branching nodes (e.g., such that application-based notifications are an option at the second branching node but were not at the first).

[0069] Upon identifying the content, the trajectory extends to the corresponding content node (430a, 430b or 430c). The corresponding content is then transmitted to the user device, such that it can be presented at the user device.

[0070] The content can include one or more tracking links to a webpage at the web site. Upon detecting that a tracking link has been activated, the trajectory can extend to a third branching node 420. A decision can then be made at third branching node 415 to determine specific content to be presented at the requested webpage.

[0071] Thus, third branching node 420 is connected to a first webpage content node 435a that represents content that identifies four representative products – each associated with a different category; a second webpage content node 435b that represents content that identifies four representative products – each associated with a same category; and a third webpage content node 435c that represents content that identifies a single product predicted to be of interest to a given user based on previous webpage-interaction data.

[0072] The third decision can be made using the machine-learning model configured based upon one or more third parameters. The third parameter(s) can differ from the first parameter(s) and/or the second parameter(s) as a result of temporal changes to anonymized and/or partially

anonymized user data, the client-specific data and/or as a result of differences of the potential actions. Additionally, the profile data processed at third branching node 420 can be different than that processed at first branching node 410 and/or second branching node 415 (e.g., as a result of detecting new metadata in communications from the user device and/or receiving new
5 information corresponding to the profile from a remote system).

[0073] Upon identifying the content, the trajectory extends to the corresponding content node (435a, 435b or 435c). The corresponding content is then transmitted to the user device, such that it can be presented at the user device within a corresponding webpage.

[0074] It will be appreciated that, while communication decision tree 400 depicted in FIG. 4
10 shows a single decision being made at each communication stage (when a notification is to be transmitted, when a body of a notification is to be presented, and when a webpage is to be presented), multiple decisions may instead be made using a machine-learning model. For example, at branching node 410, a decision may be made as to what communication channel to use and when to transmit a notification (e.g., by identifying a time within a time period or a time
15 from amongst a set of potential times). As another example, a separate decision may be made before or after the communications-channel decision to identify a transmission time. Thus, a machine-learning model may be configured to generate multiple outputs or multiple machine-learning models can have multiple configurations (each corresponding to different parameters and/or hyperparameters, each trained separately and/or each producing a separate type of output).

[0075] FIG. 5 illustrates an example of a trajectory 500 corresponding to a user device and
20 extending through communication decision tree 400. In this instance, a machine-learning result made at first branching node 410 indicated that an email communication channel was to be used to send a notification, such that trajectory 500 extended to first action node 425a. An email notification is then transmitted to the user device. A request for email content is detected,
25 indicating that a user is attempting to view the email, such that trajectory 500 extends to second branching node 415. There, a decision is made to include content that includes an identification of a discounts in the email. Thus, trajectory 500 extends to third notification content node 430c, and the corresponding content is transmitted to the user device.

[0076] A request for a webpage corresponding to a targeted link within the email is then detected, such that trajectory 500 extends to third branching node 420. A machine-learning result is generated that indicates that the webpage is to include content that identifies four representative products – each associated with a different category. Therefore, trajectory 500 extends to first email content node 435a, at which the corresponding webpage content is transmitted to the user device.

[0077] In the depicted instance, the decisions at the first branching node, the second branching node and the third branching node are indicated as having been made at 5pm on a first day, 12pm on a second day, and 6pm on the second day. Corresponding actions are then immediately performed. It will be appreciated that action times may further be decided in accordance with a machine-learning model execution, client rule or other technique.

[0078] It will be further appreciated that identifying the machine-learning-based decision can include implementing one or more additional constraints and/or factors. Alternatively or additionally, the machine-learning-based decision can be further modified based on one or more additional constraints and/or factors. For example, U.S. Application Number 14/798,293, filed on July 13, 2015, (which is hereby incorporated by reference in its entirety for all purposes) further details additional techniques to dynamically identify communication characteristics, which may be further combined with machine-learning techniques disclosed herein.

[0079] Various interfaces can be provided that facilitate defining the communication decision tree and initializing various machine-learning models. Specifically, an interface can enable a client to define a structure of the communication decision tree and/or – for each decision node – one or more hyperparameters of a machine-learning model to be executed at the node. It will be noted that a machine-learning model can be defined based on one or more hyperparameters and one or more parameters. Each of the one or more hyperparameters includes a variable that is not learned via training of the machine-learning model, while the one or more parameters include one or more variables that are learned via training of the machine-learning model. Thus, an interface can be configured to allow a client to specify hyperparameters that indicate (for example) a number of branching nodes, actions corresponding to each branch connected to each branching node, other inter-node connections, one or more constraints to be observed during execution of individual machine-learning models, and so on. In some instances, a hyperparameter indicates an

extent to which a tendency towards a particular decision is to be boosted or dampened at a branching node of a communication decision tree. Other examples of hyperparameters include a number of layers to be included in a communication decision tree, an indication as to any convergence to occur in a communication tree, a learning rate, and/or an optimization variable.

5 For example, a hyperparameter can include a optimization variable that identifies a particular target node within a communication decision tree. The machine-learning model can then be trained to learn parameters that optimize decisions for resulting in trajectories that cross through or end at the particular target node. **[0080]** Various interfaces may be provided to receive input from a client to influence configuration and/or implementation of a machine-learning model. For
10 example, an interface can be configured to receive input that indicates a target of optimization (e.g., target node within a communication decision tree), an identification of one or more inputs to be processed by the model (e.g., one or more profile attributes to be processed generally throughout the communication decision tree and/or at specific nodes), a general bias towards or against a trajectory direction, a bias towards or against a trajectory direction to be implemented
15 at one or more particular nodes within the tree, and so on.

[0081] FIG. 6 shows an exemplary interface 600 to define a communication decision tree. Specifically, interface includes a canvas 605 on which representations of various nodes can be positioned and connected. Interface 600 can include a set of icons that can be selected and positioned on canvas 605 to represent specific sequential operations. The set of icons can include
20 a start icon 610 representing a start of the communication decision tree. Start icon 610 can be associated with configuration logic that can receive a definition of a condition that, when satisfied, indicates that a trajectory through the communication decision tree is to be initiated.

[0082] The set of icons can further include an end icon 615. The communication decision tree can be defined to indicate that a given trajectory is complete upon reaching end icon 615. A
25 client can then connect action-defining icons and/or event-detection icons between a positioned start icon 610 and a positioned end icon 615 to represent various operations and assessments that are to be performed during trajectory observations.

[0083] An action-defining icon included in the set of icons can be a switch icon 620. Switch icon 620 corresponds to a branching node, at which a branch is selected or “switched to”. The
30 selection can be made using a configured machine-learning model and profile data. In many

instances, switch icon 620 is connected to multiple other communication icons, such as an email icon 625 indicating that an email is to be transmitted to a user device, a text-message icon 630 indicating that a text or SMS message is to be transmitted to a user device, and/or an app-message icon 635 indicating that an alert is to be indicated via an app installed at a user device.

5 **[0084]** An event-detection icon included in the set of icons can include target-detection icon 637, which represents that an event that corresponds to a target outcome for one or more machine-learning techniques has been detected. Target-detection icon 637 and/or another event-detection icon can indicate (for example) that a notification has been opened, a targeted link included in a notification has been activated, a user device associated with a trajectory has
10 initiated a session with a web site, a product (e.g., any product or a specific product) has been purchased on a web site, additional profile information corresponding to the trajectory has been provided, and so on.

[0085] Interface 600 can include a connection tool 640 that can be used to connect multiple icons in a directional manner. Each connection can indicate that the communication decision tree
15 is configured to allow a trajectory to extend across the connected node in the indicated direction. However, each connection can be associated with a condition, such that a trajectory only extends across the connection when the condition is satisfied. For example, a connection can be configured such that a condition is satisfied when a determination is made at a branching node (connected at one end of the connection) to perform an action represented by a communication
20 icon (connected at another end of the connection). As another example, a condition may be configured to be satisfied upon detecting a particular type of interaction in association with a trajectory-associated user device.

[0086] Each action-defining icon can be associated with one or more action parameters that define a detailed action to be performed when a trajectory has reached the icon. For example, a
25 parameter-defining interface may presented as part of interface 600 upon clicking on an icon and/or a parameter-defining interface can be opened in a pop-up window upon right-clicking on and/or double-clicking the icon.

[0087] In some instances, each action-defining icon and/or event-detection icon corresponds to a widget or piece of code that can be independently executed. Canvas 605 can serve as a

communication fabric, such that a result produced by one widget (e.g., an indication from a machine-learning model corresponding to a switch icon that a communication is to be transmitted in accordance with a particular communication specification) can be availed to another widget (e.g., a widget corresponding to a communication icon corresponding to the particular communication specification). Thus, canvas 605 can extend trajectories in response to widget results and orchestrate aspects of communication exchanges.

[0088] While not shown in FIG. 6, it will be appreciated that, in some instances, multiple switch icons 620 can be positioned on canvas 605. Each switch icon 620 can correspond to a separate instance of a machine-learning model that can be separately configured and operated.

[0089] FIG. 7 shows an exemplary parameter-defining interface 700 for a switch icon. Parameter-defining interface 700 includes a field for a Stage Label that accepts text input. The text input can subsequently displayed next to the associated icon in the interface for defining the communication decision tree. A description can also be entered via text input, which can be displayed (for example) in the interface for defining the communication decision tree in response to detecting a single click or double click in association with the icon.

[0090] For switch icons that are configured to identify a selection or action specification and/or that are configured to implement a machine-learning model, parameter-defining interface 700 can include a field to define a target outcome. For example, a pull-down menu may identify a set of events that are being tracked and are available for identification as a target outcome. The target outcome can include an action initiated at a user device, a system-initiated notification, etc. For example, a target outcome can include detecting that a link within a communication availed to the user device was clicked, that a communication availed to the user device was opened, that a communication resulted in a purchase made in association with the user device (i.e., a conversion), that a chat session was initiated, that a form was completed, etc.

[0091] For switch icons that are configured to identify a selection or action specification and/or that are configured to implement a machine-learning model, parameter-defining interface 700 can further include one or more fields that indicate potential results to be identified. For example, interface 700 includes fields that correspond to htree paths or branches that extend from the icon. In this instance, a stage-label name of another action-defining icon is identified for

each path. In some instances, path information is automatically updated at parameter-defining interface 700 upon detecting that a switch is connected to one or more action-defining icons at the interface for defining the communication decision tree. It will also be appreciated that parameter-defining interface 700 can include an option to add an additional path, remove a path, etc.

[0092] In some instances, one of the paths can be identified as a default path. Trajectories may then generally be routed to the default path unless (for example) a machine-learning model predicts that another path will have at least a threshold degree of greater probability of resulting in the target outcome, traversal through another path will produce additional data for the path that is of a threshold value (e.g., as indicated by a predicted improvement in confidences of subsequent predictions), etc. In some instances, whether a default path is selected depends on a confidence associated with a probability of the target outcome occurring (e.g., unless it is predicted that another path has at least a 60% probability of resulting in a target outcome and that the probability has a confidence of at least 50%).

[0093] In some instances, a switch icon can be configured to select a path and/or next action (or lack thereof) and a time to extend the path to a next icon (e.g., and perform any next action). The time can be selected from amongst multiple actions and/or along an open or closed continuum. In the depicted instance, parameter-defining interface 700 includes a maximum time at which the trajectory is extended to a next action-defining icon. Thus, here, the trajectory is to be extended no later than one day after the trajectory has reached the switch icon unless decision logic executed in association with the switch icon indicates that another time period is sufficiently more advantageous (e.g., due to a higher probability of resulting in a target outcome and/or to increased training data).

[0094] A machine-learning technique and/or other selection technique can be configured to identify a path from amongst multiple potential paths that is associated with a highest probability of resulting in a target outcome. In some instances, the technique further introduces some degree of noise and/or variability such that a path deemed to be sub-optimal are occasionally selected to facilitate continue training of an underlying model.

[0095] In some instances, a client may have a reason to introduce a bias towards or against selection of a particular path. For example, a particular path may be costly (e.g., computationally and/or financially) to use relative to another path. As another example, a particular path may have high availability relative to another path. As yet another example, a client may desire to quickly gain information as to an efficacy of a given path so as to inform subsequent resource-allocation decisions.

[0096] Thus, parameter-defining interface 700 can include one or more options to effect a bias towards or against individual paths. In the depicted instance, a slider is provided for each path. When the slider is positioned towards the right “Boost” side, the path-selection technique can be adjusted to be biased towards a corresponding path. When the slider is positioned towards the left “Constrain” side, the path-selection technique can be adjusted to be biased against a corresponding path. Such boosting and/or constraining options may have imposed limits, such that (for example) an effect of moving the slider to the left-most constrain position is not to prevent a selection of a corresponding path. Such limits can allow a machine-learning model to continue to collect data pertaining to various options and continue to modify one or more parameters through learning. When there are only two options, a single interface component may be provided to identify relative bias towards one option versus the other. Meanwhile, option-specific boosting/constraining options can provide more intuitive controls when there are more than two options.

[0097] FIG. 8 shows another parameter-defining interface 800 that includes options to effect a bias towards or against representing various content in communications. In the depicted instance, nine content items (each representing a corresponding product) are represented. A slider is provided in visual association with a representation of each content item. When the slider is positioned towards the right “Boost” side, a content selection (e.g., which can correspond to selecting between multiple paths representing different content or can correspond to selecting content subsequent to identifying a communications channel) can be adjusted to be biased towards a corresponding content item. When the slider is positioned towards the left “Constrain” side, the path-selection technique can be adjusted to be biased against a corresponding item.

[0098] In the depicted instance, the slider is positioned to a left-most position. This triggers a “Never Offer” option to be presented. In some instances, if the Never Offer option is not selected, the first content item may at least occasionally still be selected.

[0099] Based on the relative biases indicated by the sliders and historical communication counts, a system can predict a number of times that individual content items will be represented in a given day. Thus, as a client moves one or more sliders, interface 800 may automatically update estimated counts as to a number of times that individual content items will be presented (e.g., per day) given the slider positions.

[0100] It will be appreciated that different types of biases can further be identified and effected. For example, one or more sliders may be provided to indicate biases related to when a communication is transmitted. For example, a slider may indicate an extent to which to bias decisions towards an immediate transmission (and/or towards transmission at another time, such as at a capped time).

[0101] Effecting a bias (e.g., towards or against a type of communication channel, towards or against representing particular types of content, towards or against transmitting a communication at a particular time, etc.) can include (for example) modifying one or more weights in a machine-learning models and/or one or more thresholds. In some instances, effecting a bias includes performing a post-processing (e.g., to redistribute a portion of the results to improve an extent to which a distribution of communication attributes matches a target distribution indicated based on the bias(es).

[0102] FIG. 9 shows yet another parameter-defining interface 900 that includes options to effect a bias towards or against using various communication channels to transmit communications. In the depicted instance, three communication channels are represented: email, app-based notification and SMS message. A slider is provided in visual association with a representation of each channel. When the slider is positioned towards the right “Boost” side, a content transmission can be adjusted to be biased towards using the corresponding type of channel. When the slider is positioned towards the left “Constrain” side, the path-selection technique can be adjusted to be biased against a corresponding channel.

[0103] Interface 900 further shows a time-series representation indicating a number of communications that have been transmitted using each channel within a recent time period and further indicating a number of communications scheduled for transmission using each channel.. across an upcoming time period. A current time is represented by the vertical line. The communications can be scheduled in accordance with a selection technique that selects between multiple potential transmission times (e.g., which may be included in a same or different machine-learning model relative to one selecting a communication channel). Thus, a client can view scheduled loads across various channels and determine whether to adjust any biases set for or against a channel.

[0104] FIG. 10 shows a flowchart for a process 1000 for using machine learning to direct trajectories through a communication decision tree according to some embodiments of the invention. Process 1000 begins at block 1005 where a data structure representing a communication decision tree is accessed. The communication decision tree can be configured to dynamically define individual trajectories through the communication decision tree using a machine-learning technique to indicate a series of communication specifications. More specifically, the communication decision tree can include a set of nodes. A given trajectory can be extended across nodes in response to detecting an event indicating that such extension is to occur. An event can include (for example) detecting a particular type of action or communication from a user event or can include identifying a particular decision (corresponding to a node identification) at a trajectory-managing system or machine learning data platform. The set of nodes can include a set of branching nodes. Each branching node of the set of branching nodes can correspond to an action point configured to identify a direction for a given trajectory and/or to identify a particular action to be initiated at a trajectory-managing system or machine learning data platform. A branching node can be configured to identify the direction or action using a configured machine learning model.

[0105] At block 1010, it is detected that a trajectory (associated with a particular user and/or particular user device) has extended to reach a branching node of the communication decision tree. The trajectory may have been extended to the branching node as a result of detecting a particular type of event initiated at the user device (e.g., a communication indicating that the user device is engaged in a session at a client-associated web site, a communication indicating that the

user has completed a profile form submission, etc.) and/or as a result of completing a particular system-initiated action.

[0106] At block 1015, learned data that has been generated by processing other user data is retrieved. The learned data can include data generated while training a machine-learning technique. It will be appreciated that the training may occur during a separate time relative to using the machine-learning technique to direct one or more trajectories, or the training and utilization of the machine-learning technique may be performed concurrently. The other user data can include one or more trajectories through a same or different communication decision tree. The other user data can indicate profile data and/or attributes corresponding to one or more users and can further indicate various events detected and/or initiated in association with individual trajectories. Thus, for example, the other user data may indicate a probability of detecting a particular type of event (e.g., identified by a client as a target outcome) when various circumstances exist.

[0107] At block 1020, one or more user attributes associated with a user corresponding to the trajectory (detected as extending to the branching node) are retrieved. The user attribute(s) can include (for example) a type of user device; a geographical location of a user device; a type of browser being used at the user device; an operating system being used at the user device; a partial or complete history of an interaction between the user device and a particular web site; an interaction between the user device and one or more other web sites, cookie data associated with the user device; historical data indicating types of notifications (e.g., types of emails, text messages and/or app messages) that were opened at the user device, that resulted in activation of an included link, etc. The one or more particular user attributes can be collected and/or retrieved locally and/or requested and received from a remote source.

[0108] At block 1025, one or more communication specifications are identified based on the learned data and the one or more user attributes. For example, the learned data can include one or more parameters of a machine-learning model (e.g., a regression model). The machine-learning model may further be defined based on one or more hyperparameters. The machine-learning model can then be configured to process the user attribute(s) using the parameter(s), hyperparameter(s) and/or underlying structure. A result of an implementation of the model may identify a selection from amongst multiple available options that is predicted to be the most

successful in achieving a target outcome. The multiple available options may correspond to (for example) different types of communication channels to be used, different types of content to be transmitted, and/or different timings of transmission. In some instances, the multiple available options share one or more other communication specifications.

5 **[0109]** At block 1030, transmission of content to a user device associated with the trajectory is triggered. The content transmission is performed in accordance with the one or more communication specifications.

10 **[0110]** At block 1035, it is determined whether the trajectory has extended to reach another branching node within the communication decision tree. If so, process 1000 returns to block 1015 and blocks 1015-1035 are repeated. However, the repeated iteration of block 1015 may include retrieving different learned data generated by processing other user data (e.g., potentially, but not necessarily, in combination with at least some of the user data). The different learned data may have been generated using a same or different configuration of the machine-learning technique (e.g., having same or different values and/or types of parameters and/or

15 hyperparameters). The repeated iteration of block 1020 can include retrieving at least one other user attribute. The repeated iteration of block 1025 can include identifying at least one other communication specification (and/or from amongst a different set of potential communication specifications) based on the different learned data and the at least one other user attribute. The at least one other communication specification can be identified using a same or different type of

20 underlying model. And the repeated iteration of block 1030 can include triggering another transmission of other content in accordance with the at least one other communication specification.

25 **[0111]** When it is determined that the trajectory has not extended to reach another branching node, process 1000 proceeds to block 1040 to determine whether the trajectory is complete. The determination can be made by determining whether a current end of a trajectory is associated with a trajectory that lacks an extending connection. If it is determined that the trajectory is complete, processing of the trajectory can be terminated. If it is determined that the trajectory is not complete, process 1000 can return to block 1035 to await a determination that the trajectory has reached another branching node (e.g., as a result of a user-initiated action or external event).

[0112] Thus, process 1000 facilitates repeatedly using differently configured machine-learning models to identify specifications corresponding to different stages in a communication exchange. At the different stages, the models can use different profile data (e.g., values for different fields or values that have changed in time) and/or different model parameters (e.g., learned based on different inputs and/or outputs pertaining to the models and/or based on temporal changes). This iterative application of machine-learning models facilitates dynamically directing communication exchanges for individual users.

[0113] FIG. 11 shows a flowchart for a process 1100 for defining a machine-learning-based communication decision tree using an interface supporting positionable visual elements. Process 1100 begins at block 1105 where an interface is availed that includes a set of visual elements and a canvas for element positioning. Each of the set of visual elements can be positionable on the canvas. For example, the interface may be configured to allow a user to click on a representation of a visual element and – while maintaining the click – drag a cursor to another position on the canvas to drop the visual element at the other position. As another example, a representation can be selected (e.g., via a click or double-click) and another input (e.g., another click or double-click) received while the cursor is at another position can cause the visual element to be positioned at the other position.

[0114] The set of visual elements can include a set of action-defining visual elements. Each action-defining visual element of the set of action-defining visual elements can a particular action that is to be performed when a given trajectory has extended to the action-defining visual element. The set of action-defining visual elements can include a switch visual element that represents a decision action (e.g., made using a machine-learning model) to identify a communication specification using a machine-learning technique. The set of action-defining visual elements can further include a set of communication visual elements. Each of the set of communication visual elements can represent a particular communication specification (e.g., a type of communication channel, specific content, transmission time, etc.). The set of visual elements can also include a connection visual element configured to directionally connect multiple positioned visual elements. Each positioned visual element of the multiple positioned visual elements can correspond to an action-defining visual element of the set of action-defining

visual elements. The directional connection can indicate an order in which particular actions represented by the multiple positioned visual elements are to occur.

[0115] At block 1110, an update to the canvas is detected. The updated canvas can include the switch visual element being positioned at a first position within the canvas, a first

5 communication visual element of the set of communication visual elements positioned at a second position within the canvas, and a second communication visual element of the set of communication visual elements being positioned a third position within the canvas. The first communication visual element can represent a first particular communication specification, and the second communication visual element can represent a a second particular communication
10 specification.

[0116] The updated canvas can further include a set of connection visual elements. Each of the set of connection visual elements can include an instance of the connection visual element. A first connection of the set of connection visual elements can be positioned to connect the switch visual element to the first communication visual element. A second connection of the set of
15 connection visual elements can be positioned to connect the switch visual element to the second communication visual element. The set of connection visual elements can indicate that potential results of execution of the machine-learning technique at the switch visual element include a first result that triggers a communication transmission having the first particular communication specification and a second result that triggers a communication transmission having the second
20 particular communication specification.

[0117] At block 1115, a particular communication decision tree is defined based on the updated canvas. At block 1120, it is detected that a given trajectory associated with particular profile data has extended to a particular decision action represented by the switch visual element. In response to the detection, at block 1125, the machine-learning technique (configured with
25 learned parameter data and/or static data) is used to process the particular profile data to produce a machine-learning result. The learned parameter data can include data learned during a separate or ongoing training of a machine-learning model based on a set of trajectories associated with other users and/or associated with a same or different communication decision tree. The processing of the particular profile data using the machine-learning technique can indicate which

one of the first and second particular communication specifications is to be applied for a content transmission.

[0118] Thus, at block 1130, content is transmitted to a user device associated with the trajectory. The transmission is performed in accordance with the one of the first and second particular communication specifications as indicated in the machine-learning result. For example, the first and second communication visual elements may correspond to different types of communication channels. Block 1125 may then include identifying one of the two types of communication channels, and the content can be transmitted via the identified channel.

[0119] Thus, the canvas facilitates defining configurations for a communication decision tree. However, a client need not define a communication exchange that applies to all users and/or that includes merely one or more deterministic rules. Rather, the interface supports generally identifying options of various communication specifications, an order of communication events and/or constraints. Specification communication specifications can then be automatically and dynamically generated using machine-learning techniques. This approach can facilitate configuring a communication system to abide by client priorities but can allow the communication system to dynamically adapt to characteristics of particular users, resource loads, recent interaction patterns, etc.

[0120] It will be appreciated that variations of the disclosed techniques are contemplated. For example, a branching node may use another type of artificial-intelligence model that is not a machine-learning model to select a communication specification to be used for a communication. As another example, an interface may be configured to accept a selection of a particular type or a more general type of artificial-intelligence model that is to be used at a trajectory stage corresponding to a switch element. As yet another example, an interface may be configured to allow an indication of what data (e.g., in terms of corresponding to one or more communication decision trees, one or more time periods, and/or one or more user-population segments) is to be used to train a machine-learning model corresponding to one, more or all switch elements positioned on a canvas.

[0121] It will be appreciated that technology disclosed herein can be used to support various types of decision trees. For example, nodes in the tree and/or visual elements represented on a

canvas can (in some instances) correspond to elements that generally are associated with logic that evaluates whether a given condition is satisfied (e.g., a particular type of inter-device communication is detected, a non-client-associated application indicates that an action was performed, a particular time has passed) and, upon detecting satisfaction, a particular action is performed. For a subset of the nodes and/or visual elements, the conditioned particular action can include executing a machine-learning model based on profile data to select from amongst a set of connected nodes (or visual elements) to proceed, such that another particular action associated with the selected node (or visual element) can be performed. For example, machine-learning-based selection of trajectory paths may be integrated into an If This Then That environment. Rather than having branching nodes connected to nodes identifying communication specifications, the branches could (for example) identify different applications to use to store data. Thus, a decision framework can be established to enable an artificial-intelligence applet and/or plugin to communicate with one or more other applets or back through a canvas.

[0122] It will further be appreciated that, while some disclosures herein indicate that a target outcome can be used to shape machine-learning training and execution, more complicated instances are considered. For example, a negative outcome (e.g., an unsubscribe request or complaint) can alternatively or additionally be identified and used. In some instances, a score can be assigned to various results based on a quantity or extent to which one or more target results and/or one or more negative results occurred. The score can then be used for training and implementing one or more machine-learning models.

[0123] Specific details are given in the above description to provide a thorough understanding of the embodiments. However, it is understood that the embodiments can be practiced without these specific details. For example, circuits can be shown in block diagrams in order not to obscure the embodiments in unnecessary detail. In other instances, well-known circuits, processes, algorithms, structures, and techniques can be shown without unnecessary detail in order to avoid obscuring the embodiments.

[0124] Implementation of the techniques, blocks, steps and means described above can be done in various ways. For example, these techniques, blocks, steps and means can be implemented in hardware, software, or a combination thereof. For a hardware implementation, the processing units can be implemented within one or more application specific integrated

circuits (ASICs), digital signal processors (DSPs), digital signal processing devices (DSPDs), programmable logic devices (PLDs), field programmable gate arrays (FPGAs), processors, controllers, micro-controllers, microprocessors, other electronic units designed to perform the functions described above, and/or a combination thereof.

5 **[0125]** Also, it is noted that the embodiments can be described as a process which is depicted as a flowchart, a flow diagram, a data flow diagram, a structure diagram, or a block diagram. Although a flowchart can describe the operations as a sequential process, many of the operations can be performed in parallel or concurrently. In addition, the order of the operations can be re-arranged. A process is terminated when its operations are completed, but could have additional
10 steps not included in the figure. A process can correspond to a method, a function, a procedure, a subroutine, a subprogram, etc. When a process corresponds to a function, its termination corresponds to a return of the function to the calling function or the main function.

[0126] Furthermore, embodiments can be implemented by hardware, software, scripting languages, firmware, middleware, microcode, hardware description languages, and/or any
15 combination thereof. When implemented in software, firmware, middleware, scripting language, and/or microcode, the program code or code segments to perform the necessary tasks can be stored in a machine readable medium such as a storage medium. A code segment or machine-executable instruction can represent a procedure, a function, a subprogram, a program, a routine, a subroutine, a module, a software package, a script, a class, or any combination of instructions,
20 data structures, and/or program statements. A code segment can be coupled to another code segment or a hardware circuit by passing and/or receiving information, data, arguments, parameters, and/or memory contents. Information, arguments, parameters, data, etc. can be passed, forwarded, or transmitted via any suitable means including memory sharing, message passing, ticket passing, network transmission, etc.

25 **[0127]** For a firmware and/or software implementation, the methodologies can be implemented with modules (e.g., procedures, functions, and so on) that perform the functions described herein. Any machine-readable medium tangibly embodying instructions can be used in implementing the methodologies described herein. For example, software codes can be stored in a memory. Memory can be implemented within the processor or external to the processor. As used herein
30 the term “memory” refers to any type of long term, short term, volatile, nonvolatile, or other

storage medium and is not to be limited to any particular type of memory or number of memories, or type of media upon which memory is stored.

[0128] Moreover, as disclosed herein, the term "storage medium" can represent one or more memories for storing data, including read only memory (ROM), random access memory (RAM), magnetic RAM, core memory, magnetic disk storage mediums, optical storage mediums, flash memory devices and/or other machine readable mediums for storing information. The term "machine-readable medium" includes, but is not limited to portable or fixed storage devices, optical storage devices, wireless channels, and/or various other storage mediums capable of storing that contain or carry instruction(s) and/or data.

[0129] While the principles of the disclosure have been described above in connection with specific apparatuses and methods, it is to be clearly understood that this description is made only by way of example and not as limitation on the scope of the disclosure.

WHAT IS CLAIMED IS:

1. A computer-implemented method comprising:
 availing an interface for defining communication decision trees through which
trajectories corresponding to communication exchanges are extended, the interface including:

5 a set of positionable visual elements that includes:

 a set of action-defining visual elements, each action-defining
visual element of the set of action-defining visual elements representing a
particular action that is to be performed when a given trajectory has extended to
the action-defining visual element, the set of action-defining visual elements
10 including:

 a switch visual element that represents a decision action to
identify a communication specification using a machine-learning
technique; and

 a set of communication visual elements, each of the set of
15 communication visual elements representing a particular communication
specification; and

 a connection visual element configured to directionally connect
multiple positioned visual elements, each positioned visual element of the
multiple positioned visual elements corresponding to an action-defining visual
20 element of the set of action-defining visual elements, wherein the directional
connection indicates an order in which particular actions represented by the
multiple positioned visual elements are to occur; and

 a canvas configured to accept positioning of one or more of the set of
positionable visual elements;

25 detecting an update to the canvas, wherein the updated canvas includes:

 the switch visual element positioned at a first position;

 a first communication visual element of the set of communication visual
elements positioned at a second position, the first communication visual element
representing a first particular communication specification;

a second communication visual element of the set of communication visual elements positioned a third position, the second communication visual element representing a second particular communication specification; and

a set of connection visual elements, each of the set of connection visual elements being an instance of the connection visual element, a first connection of the set of connection visual elements connecting the switch visual element to the first communication visual element, a second connection of the set of connection visual elements connecting the switch visual element to the second communication visual element, the set of connection visual elements indicating that potential results of execution of the machine-learning technique at the switch visual element include a first result that triggers a communication transmission having the first particular communication specification and a second result that triggers a communication transmission having the second particular communication specification, and

defining a particular communication decision tree based on the updated canvas that, when implemented and in response to detecting that a given trajectory associated with particular profile data has extended to a particular decision action represented by the switch visual element, triggers:

using the machine-learning technique and learned parameter data to process the particular profile data to produce a machine-learning result indicating which one of the first and second particular communication specifications is to be applied for a content transmission; and

transmitting content to a user device associated with the trajectory, the transmission being in accordance with the one of the first and second particular communication specifications as indicated in the machine-learning result.

2. The method of claim 1, wherein the interface or another interface further includes an input component configured to receive a specification of an optimization variable, wherein the method further comprises:

detecting new data at the input component that identifies a particular optimization variable; and

training a machine-learning model using the particular optimization variable to generate at least part of the learned parameter data;

wherein using the machine-learning technique and learned parameter data includes running the machine-learning model configured with the learned parameter data.

5 3. The method of claim 1, wherein the interface or another interface further includes one or more input components configured to receive an identification of one or more profile attributes to be processed by the machine-learning technique, wherein the particular profile data selectively includes the one or more profile attributes.

10 4. The method of claim 1, further comprising:
receiving an indication of a bias towards or against a specific communication attribute; and
implementing an adjustment to the machine-learning technique that corresponds to the bias, wherein the machine-learning result is produced using the machine-learning technique having had the adjustment implemented.

15 5. The method of claim 1, further comprising:
accessing a set of trajectories, each trajectory of at least some of the set of trajectories identifying a path along at least part of a different communication decision tree; and
performing a training task using the particular communication decision tree and the set of trajectories to generate the learned parameter data.

20 6. The method of claim 1, wherein each of the first particular communication specification and the second particular communication specification identifies a different type of communication channel, and wherein the content is transmitted via one of the different types of communication channel as indicated by the machine-learning result.

25 7. The method of claim 1, wherein:
the updated canvas further includes:
a second switch visual element positioned at a third position;

a third communication visual element positioned at a fourth position, the third communication visual element representing a third particular communication specification; and

a fourth communication visual element positioned a fifth position, the fourth communication visual element representing a fourth particular communication specification;

the set of connection visual elements further includes:

a third connection connecting the second switch visual element to the third communication visual element and a fourth connection connecting the second switch visual element to the fourth communication visual element, the third and fourth communication visual elements indicating that potential results of execution of the machine-learning technique at the second switch visual element include a third result that triggers a communication transmission having the third particular communication specification and fourth result that triggers a communication transmission having the fourth particular communication specification, and

the a particular communication decision tree is further defined such that, when implemented and in response to detecting that the given trajectory associated with particular profile data has extended to another particular decision action represented by the second switch visual element, triggers:

using the machine-learning technique and other learned parameter data to process at least part of the particular profile data to produce another machine-learning result indicating which one of the third and fourth particular communication specifications is to be applied for another content transmission; and

transmitting other content to the user device in accordance with the one of the third and fourth particular communication specifications as indicated in the other machine-learning result.

8. The method of claim 1, wherein the interface is configured such that each individual visual element of the set of positionable visual elements can be dragged and dropped to a location within the canvas.

9. The method of claim 1, wherein using the machine-learning technique and learned parameter data to process the particular profile data includes applying a regression model to the particular profile data.

10. A computer-program product tangibly embodied in a non-transitory machine-readable storage medium, including instructions configured to cause one or more data processors to perform operations including:

 availing an interface for defining communication decision trees through which trajectories corresponding to communication exchanges are extended, the interface including:

 a set of positionable visual elements that includes:

 a set of action-defining visual elements, each action-defining visual element of the set of action-defining visual elements representing a particular action that is to be performed when a given trajectory has extended to the action-defining visual element, the set of action-defining visual elements including:

 a switch visual element that represents a decision action to identify a communication specification using a machine-learning technique; and

 a set of communication visual elements, each of the set of communication visual elements representing a particular communication specification; and

 a connection visual element configured to directionally connect multiple positioned visual elements, each positioned visual element of the multiple positioned visual elements corresponding to an action-defining visual element of the set of action-defining visual elements, wherein the directional connection indicates an order in which particular actions represented by the multiple positioned visual elements are to occur; and

 a canvas configured to accept positioning of one or more of the set of positionable visual elements;

 detecting an update to the canvas, wherein the updated canvas includes:

 the switch visual element positioned at a first position;

a first communication visual element of the set of communication visual elements positioned at a second position, the first communication visual element representing a first particular communication specification;

5 a second communication visual element of the set of communication visual elements positioned at a third position, the second communication visual element representing a second particular communication specification; and

a set of connection visual elements, each of the set of connection visual elements being an instance of the connection visual element, a first connection of the set of connection visual elements connecting the switch visual element to the first
10 communication visual element, a second connection of the set of connection visual elements connecting the switch visual element to the second communication visual element, the set of connection visual elements indicating that potential results of execution of the machine-learning technique at the switch visual element include a first result that triggers a communication transmission having the first particular
15 communication specification and a second result that triggers a communication transmission having the second particular communication specification, and

defining a particular communication decision tree based on the updated canvas that, when implemented and in response to detecting that a given trajectory associated with particular profile data has extended to a particular decision action represented by the switch
20 visual element, triggers:

using the machine-learning technique and learned parameter data to process the particular profile data to produce a machine-learning result indicating which one of the first and second particular communication specifications is to be applied for a content transmission; and

25 transmitting content to a user device associated with the trajectory, the transmission being in accordance with the one of the first and second particular communication specifications as indicated in the machine-learning result.

11. The computer-program product of claim 10, wherein the interface or another interface further includes one or more input components configured to receive an
30 identification of one or more profile attributes to be processed by the machine-learning

technique, wherein the particular profile data selectively includes the one or more profile attributes.

12. The computer-program product of claim 10, the operations further including:

5 receiving an identification of a target output that is indicative of a trajectory characteristic;

configuring a hyperparameter of the machine-learning technique based on the identification of the target output; and

10 identifying a set of parameters for the machine-learning technique, each of the set of parameters configured to be learned using trajectory data and the hyperparameter, wherein the machine-learning technique used to process the particular profile data is configured based on the set of parameters.

13. The computer-program product of claim 10, the operations further including:

15 receiving an indication of a bias towards or against a specific communication attribute; and

implementing an adjustment to the machine-learning technique that corresponds to the bias, wherein the machine-learning result is produced using the machine-learning technique having had the adjustment implemented.

20 14. The computer-program product of claim 10, the operations further including:

accessing a set of trajectories, each trajectory of at least some of the set of trajectories identifying a path along at least part of a different communication decision tree; and

25 performing a training task using the particular communication decision tree and the set of trajectories to generate the learned parameter data.

15. The computer-program product of claim 10, wherein each of the first particular communication specification and the second particular communication specification identifies a different type of communication channel, and wherein the content is transmitted via one of the different types of communication channel as indicated by the machine-learning result.

16. The computer-program product of claim 10, wherein:

the updated canvas further includes:

a second switch visual element positioned at a third position;

a third communication visual element positioned at a fourth position, the

5 third communication visual element representing a third particular communication specification; and

a fourth communication visual element positioned a fifth position, the fourth communication visual element representing a fourth particular communication specification;

10 the set of connection visual elements further includes:

a third connection connecting the second switch visual element to the third communication visual element and a fourth connection connecting the second switch visual element to the fourth communication visual element, the third and fourth communication visual elements indicating that potential results of execution of the machine-learning technique at the second switch visual element include a third result that triggers a communication transmission having the third particular communication specification and fourth result that triggers a communication transmission having the fourth particular communication specification, and

20 the a particular communication decision tree is further defined such that, when implemented and in response to detecting that the given trajectory associated with particular profile data has extended to another particular decision action represented by the second switch visual element, triggers:

25 using the machine-learning technique and other learned parameter data to process at least part of the particular profile data to produce another machine-learning result indicating which one of the third and fourth particular communication specifications is to be applied for another content transmission; and

transmitting other content to the user device in accordance with the one of the third and fourth particular communication specifications as indicated in the other machine-learning result.

17. The computer-program product of claim 9, wherein the interface is configured such that each individual visual element of the set of positionable visual elements can be dragged and dropped to a location within the canvas.

18. The computer-program product of claim 9, wherein using the machine-learning technique and learned parameter data to process the particular profile data includes applying a regression model to the particular profile data.

19. A system comprising:
one or more data processors; and
a non-transitory computer readable storage medium containing instructions which
when executed on the one or more data processors, cause the one or more data processors to perform operations including:

 availing an interface for defining communication decision trees through which trajectories corresponding to communication exchanges are extended, the interface including:

 a set of positionable visual elements that includes:

 a set of action-defining visual elements, each action-defining visual element of the set of action-defining visual elements representing a particular action that is to be performed when a given trajectory has extended to the action-defining visual element, the set of action-defining visual elements including:

 a switch visual element that represents a decision action to identify a communication specification using a machine-learning technique; and

 a set of communication visual elements, each of the set of communication visual elements representing a particular communication specification; and

 a connection visual element configured to directionally connect multiple positioned visual elements, each positioned visual element of the multiple positioned visual elements corresponding to an

action-defining visual element of the set of action-defining visual elements, wherein the directional connection indicates an order in which particular actions represented by the multiple positioned visual elements are to occur; and

5 a canvas configured to accept positioning of one or more of the set of positionable visual elements;

detecting an update to the canvas, wherein the updated canvas includes:

the switch visual element positioned at a first position;

10 a first communication visual element of the set of communication visual elements positioned at a second position, the first communication visual element representing a first particular communication specification;

15 a second communication visual element of the set of communication visual elements positioned at a third position, the second communication visual element representing a second particular communication specification; and

20 a set of connection visual elements, each of the set of connection visual elements being an instance of the connection visual element, a first connection of the set of connection visual elements connecting the switch visual element to the first communication visual element, a second connection of the set of connection visual elements connecting the switch visual element to the second communication visual element, the set of connection visual elements indicating that potential results of execution of the machine-learning technique at the switch visual element include a first result that triggers a communication transmission having the first particular communication specification and a second result that
25 triggers a communication transmission having the second particular communication specification, and

30 defining a particular communication decision tree based on the updated canvas that, when implemented and in response to detecting that a given trajectory associated with particular profile data has extended to a particular decision action represented by the switch visual element, triggers:

using the machine-learning technique and learned parameter data to process the particular profile data to produce a machine-learning result indicating which one of the first and second particular communication specifications is to be applied for a content transmission; and

5 transmitting content to a user device associated with the trajectory, the transmission being in accordance with the one of the first and second particular communication specifications as indicated in the machine-learning result.

20. The system of claim 19, the operations further including:

10 receiving an identification of a target output that is indicative of a trajectory characteristic;

configuring a hyperparameter of the machine-learning technique based on the identification of the target output; and

15 identifying a set of parameters for the machine-learning technique, each of the set of parameters configured to be learned using trajectory data and the hyperparameter, wherein the machine-learning technique used to process the particular profile data is configured based on the set of parameters.

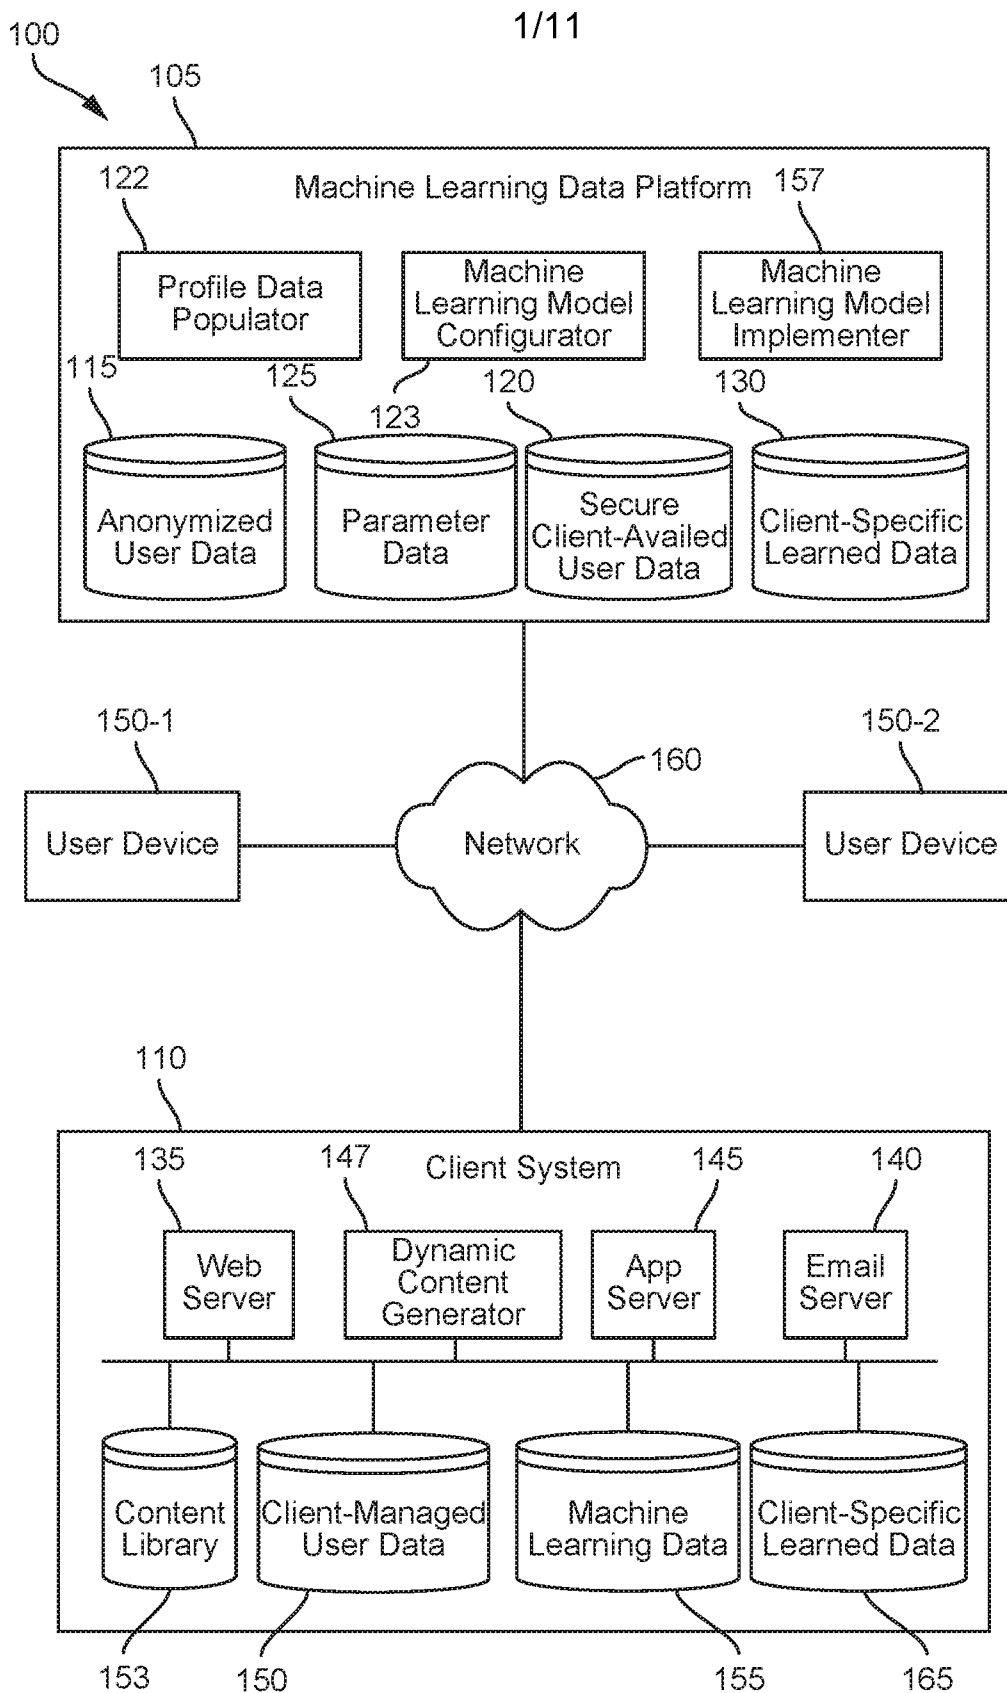


FIG. 1

2/11

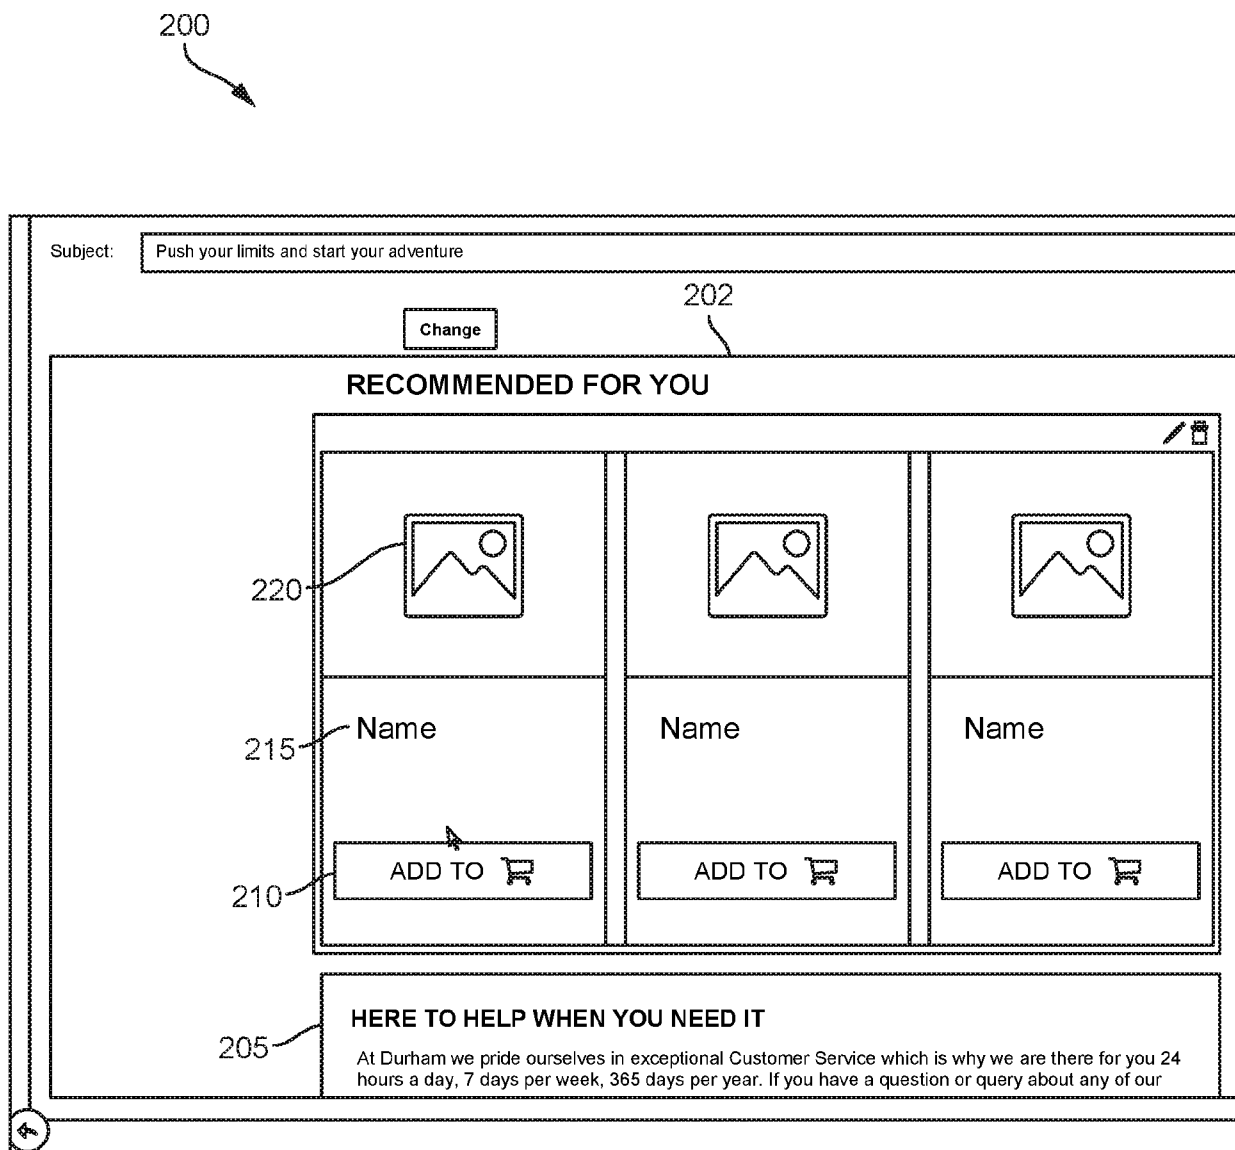


FIG. 2

3/11

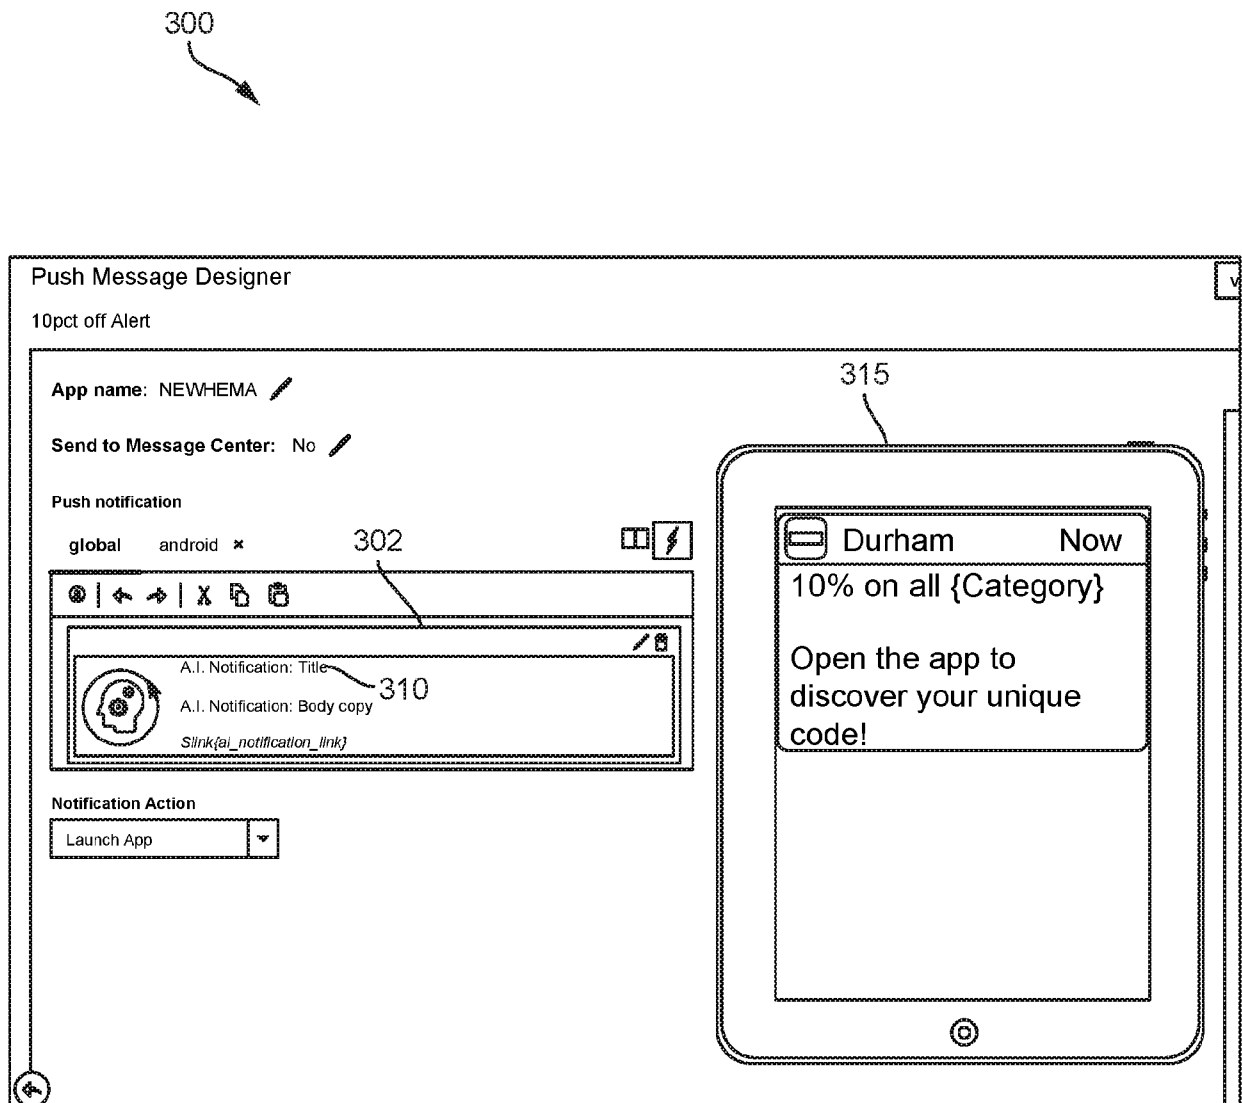


FIG. 3

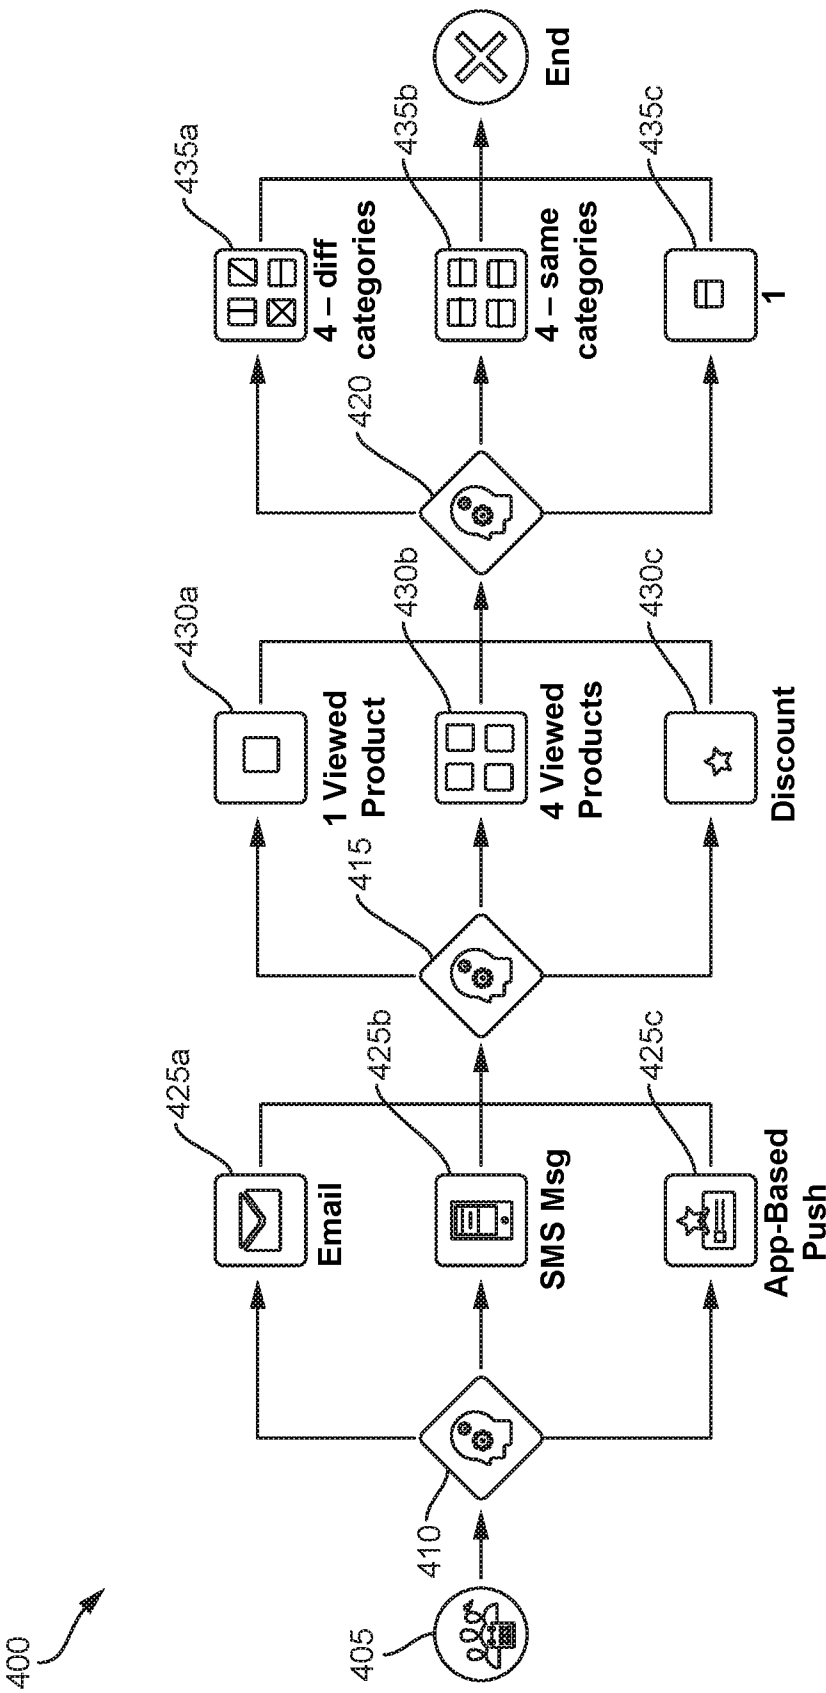


FIG. 4

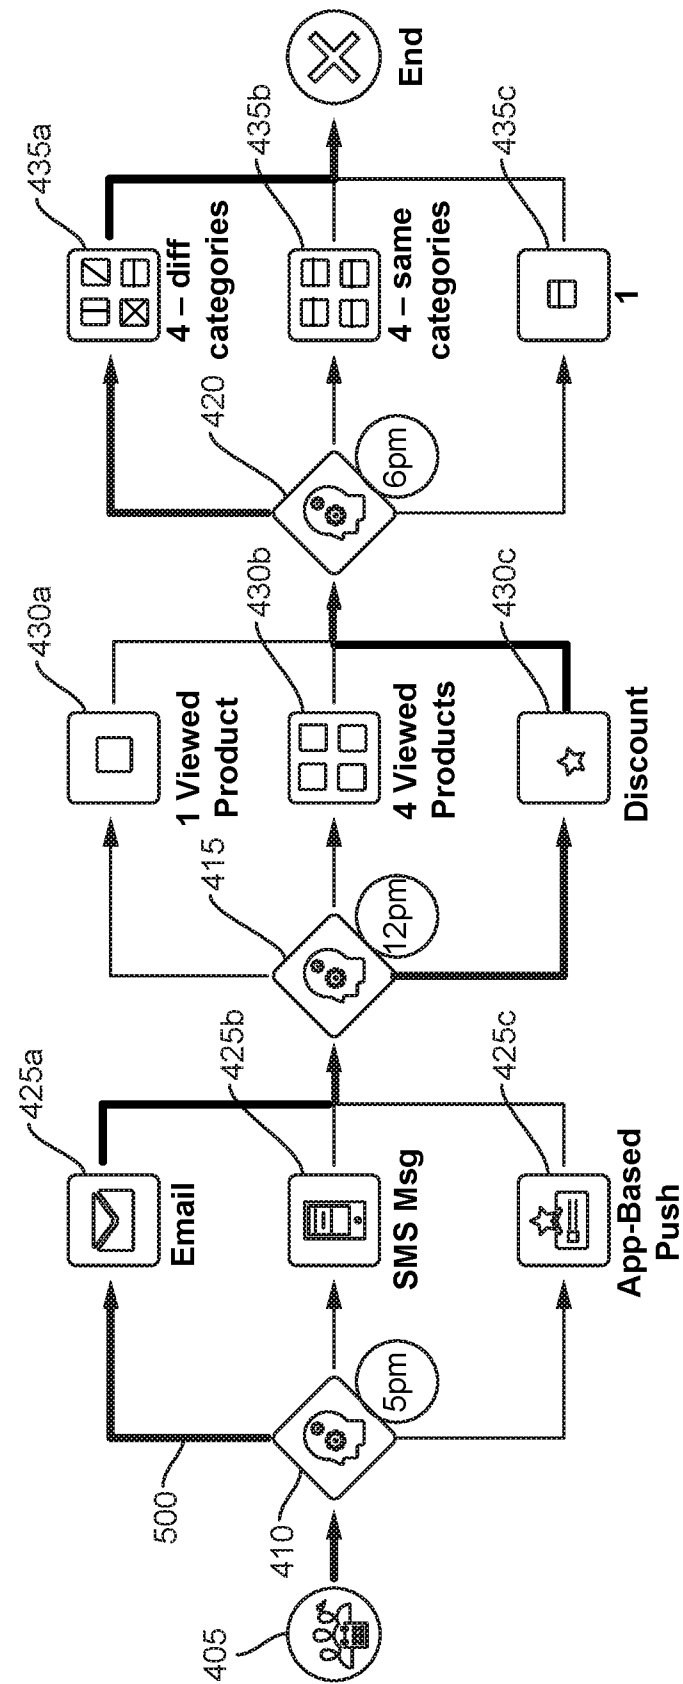


FIG. 5

6/11

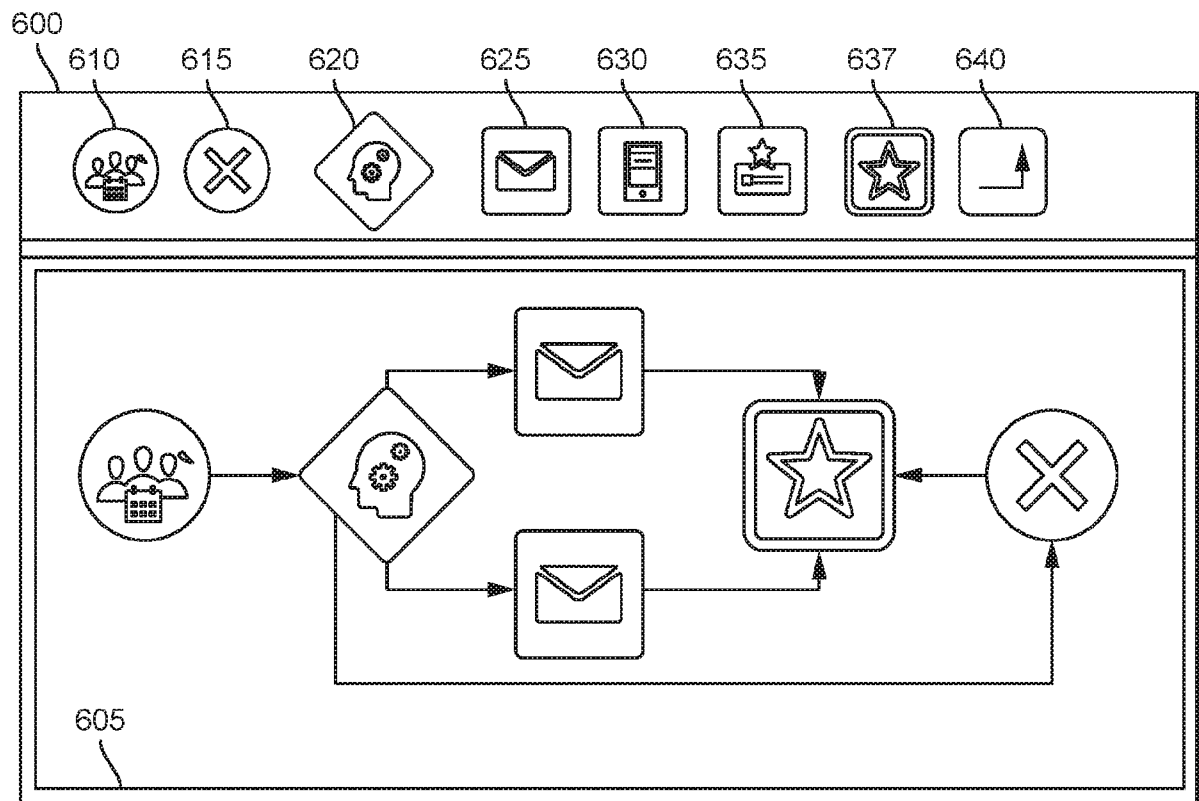


FIG. 6

7/11

700


Adaptive intelligence switch

×

Stage label

New Content Notification

Description

Select metric

Conversion

▼

Paths

Default Path

Email - Content #1

Constrain

Boost

Other Path 1

Email - Content #2

Constrain

Boost

Other Path 2

No Action

Constrain

Boost

Default Communication Time

After 1 day

▼

Cancel

Done

FIG. 7

8/11

800


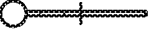



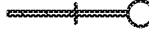












 Durham 30L Adventurer Backpack Durham 1 Category Daily Recommendations 0 Adjust Recommendations  <input type="checkbox"/> Never Offer	 Durham Curved Helmet Mount Durham 1 Category Daily Recommendations 2780 Adjust Recommendations 	 Durham Eagle 3 Action Camera Durham 1 Category Daily Recommendations 4093 Adjust Recommendations 
 Durham Eagle 5 Durham 1 Category Daily Recommendations 3724 Adjust Recommendations 	 Durham Eagle 7 Black Action Camera Durham 1 Category Daily Recommendations 3045 Adjust Recommendations 	 Durham OSC 4000 Durham 1 Category Daily Recommendations 4002 Adjust Recommendations 
 Durham OSC 4300 Durham 1 Category Daily Recommendations 0 Adjust Recommendations 	 Durham OSC 4700 Durham 1 Category Daily Recommendations 0 Adjust Recommendations 	 Durham ROC 3200 Durham 1 Category Daily Recommendations 0 Adjust Recommendations 

FIG. 8

9/11

900

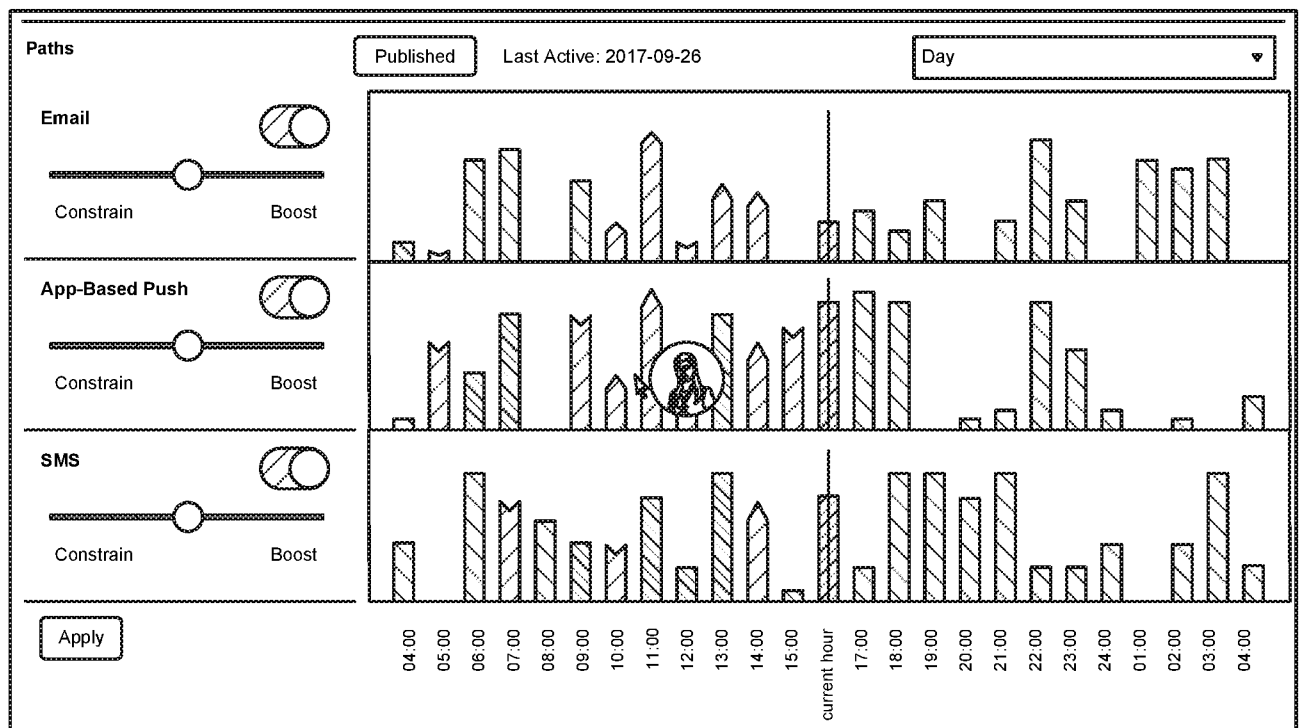



FIG. 9

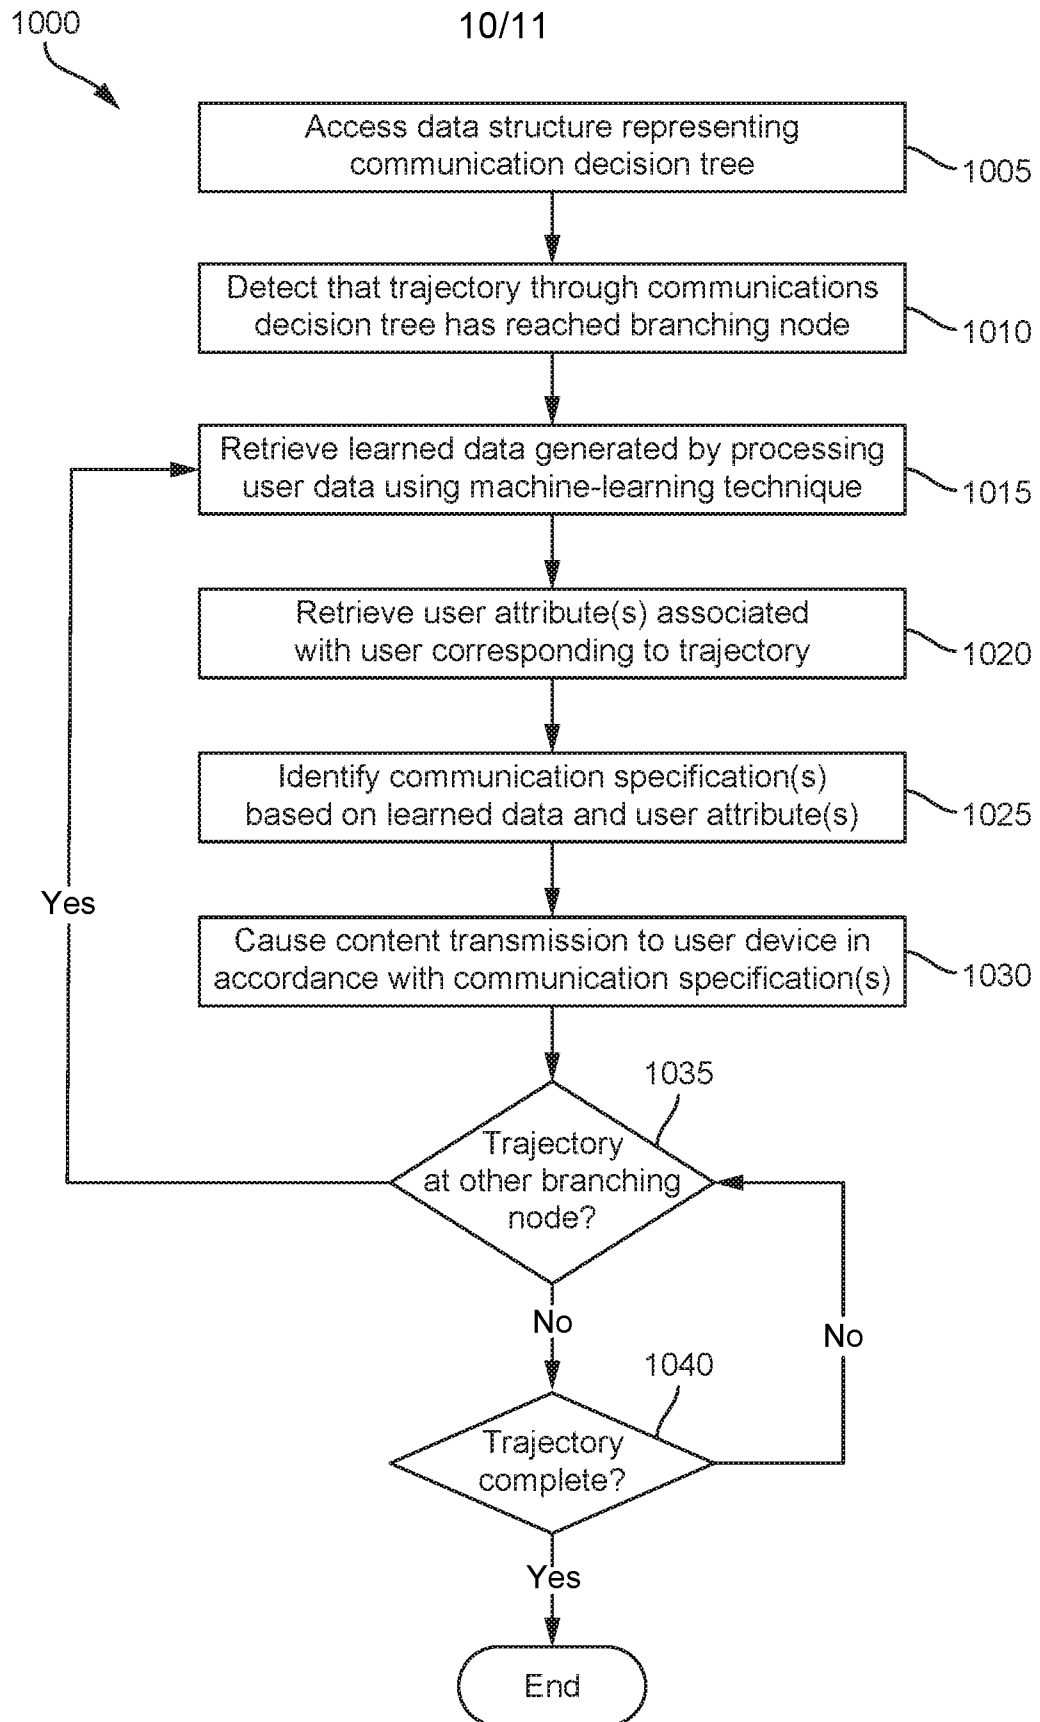
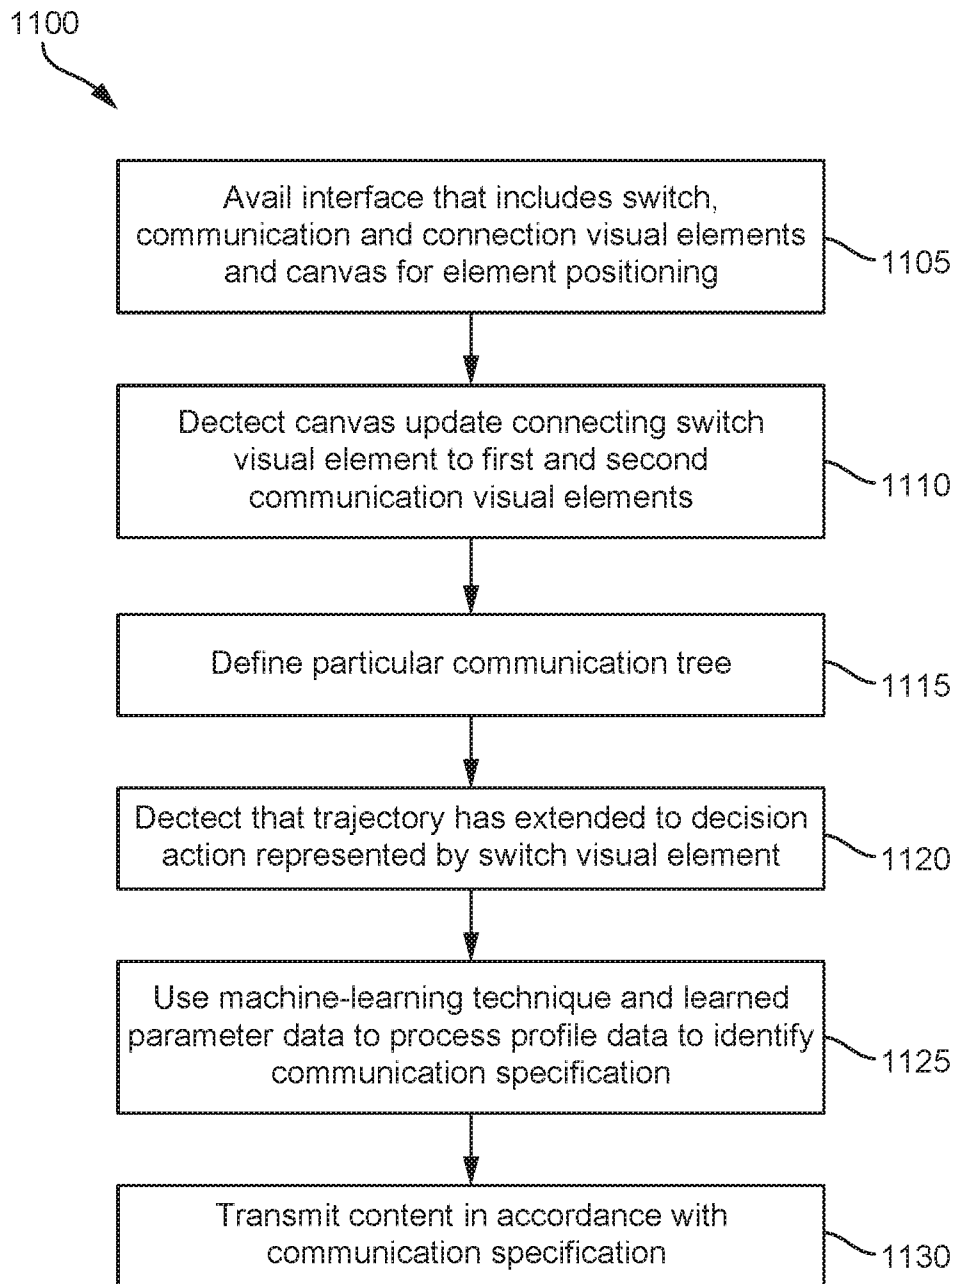


FIG. 10

11/11

**FIG. 11**

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US2018/051422

A. CLASSIFICATION OF SUBJECT MATTER

G06F 3/0484 (2013.01) G06N 20/00 (2019.01)

According to International Patent Classification (IPC)

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

GO6N, G06F

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

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

decision tree, 决策树, machine learning, regression, 机器学习, 回归模型, attribute, characteristics, profile, feature, aspect, facet, trait, 属性, 特征, 模型, communication, mode, type, medium, protocol, specifications, 通讯, 通信, 协议, 规范, content delivery, targeting, transmission, UI, interface, canvas, 用户界面, 图形界面 and related terms.

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2014/0358828 A1 (PHILLIPPS K. D. & WELLMAN R. W.) 4 December 2014 Paragraphs [0044]-[0048], [0051]-[0052], [0055]-[0059], [0106], [0110], [0133], [0194], Figure 5	1-20
X	US 2017/0061472 A1 (MATHIS C. M.) 2 March 2017 Paragraphs [0026], [0033]-[0034], [0038]-[0040], [0044]-[0046], [0052]	1-20
X	US 2015/0006294 A1 (IRMAK U. ET AL.) 1 January 2015 Paragraphs [0020], [0029], [0034], [0038]-[0039], [0046], [0059], [0062], [0064], [0078]	1-20
X	CN 105631707 A (BEIJING QIHOO TECHNOLOGY CO LTD & QIZHI SOFTWARE (BEIJING) CO LTD) 1 June 2016 Paragraphs [0117], [0129]-[0132], [0143]-[0151], [0173]-[0175] of the original non-English language document (a machine translation is enclosed only for your reference)	1-20

☒ 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

28/12/2018 (day/month/year)

Date of mailing of the international search report

04/01/2019 (day/month/year)

Name and mailing address of the ISA/SG

Intellectual Property Office of Singapore
51 Bras Basah Road
#01-01 Manulife Centre
Singapore 189554

Email: pct@ipos.gov.sg

Authorized officer

Lai Tan Lei Tanny (Dr)

IPOS Customer Service Tel. No.: (+65) 6339 8616

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US2018/051422**C (Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT**

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2009/0058859 A1 (CRAWFORD S. L.) 5 March 2009 Whole document	-
A	US 7392162 B1 (SRINIVASAN V. ET AL.) 24 June 2008 Whole document	-

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2018/051422

Note: This Annex lists known patent family members relating to the patent documents cited in this International Search Report. This Authority is in no way liable for these particulars which are merely given for the purpose of information.

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2014/0358828 A1	04/12/2014	WO 2014/194002 A1 EP 3005138 A1	04/12/2014 13/04/2016
US 2017/0061472 A1	02/03/2017	NONE	
US 2015/0006294 A1	01/01/2015	NONE	
CN 105631707 A	01/06/2016	NONE	
US 2009/0058859 A1	05/03/2009	NONE	
US 7392162 B1	24/06/2008	NONE	