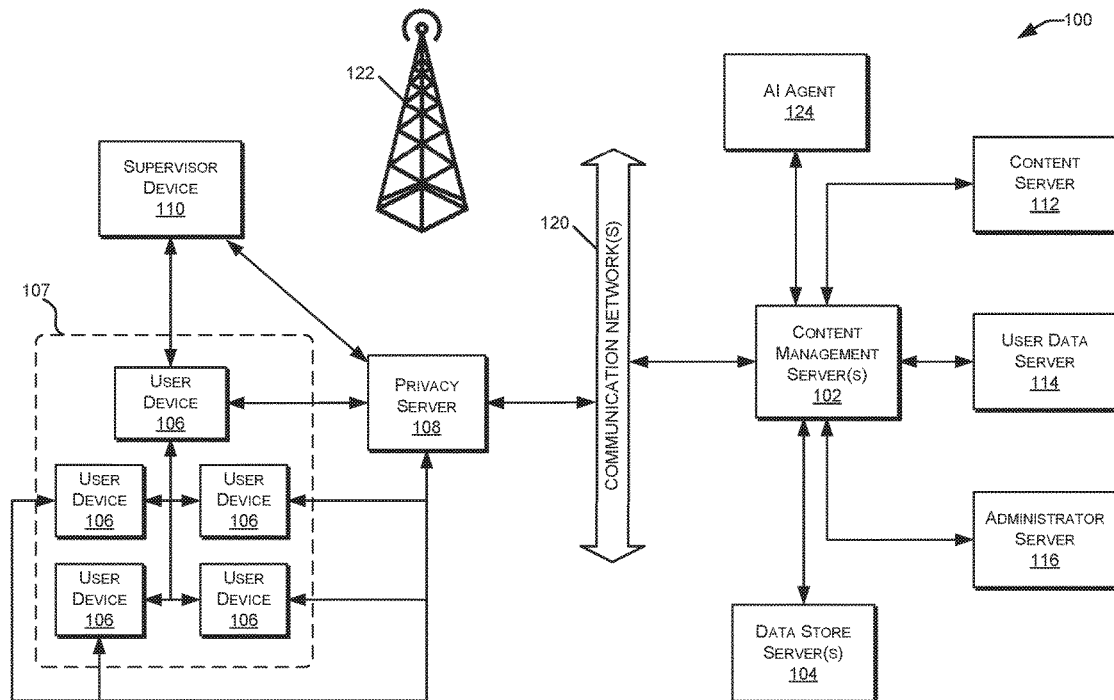




US 20170293841A1

(19) **United States**(12) **Patent Application Publication**
McAllister et al.(10) **Pub. No.: US 2017/0293841 A1**(43) **Pub. Date: Oct. 12, 2017**(54) **METHOD AND SYSTEM FOR AUTOMATED
CONTENT SELECTION AND DELIVERY**(71) Applicant: **Pearson Education, Inc.**, New York,
NY (US)(72) Inventors: **Angie McAllister**, New York, NY (US);
Brian Moriarty, Fort Lee, NJ (US);
Greg McFall, Poolesville, MD (US)(21) Appl. No.: **15/482,576**(22) Filed: **Apr. 7, 2017****Related U.S. Application Data**(60) Provisional application No. 62/320,213, filed on Apr.
8, 2016.**Publication Classification**(51) **Int. Cl.**
G06N 5/02 (2006.01)
G06N 5/00 (2006.01)
G06N 3/02 (2006.01)
G06N 99/00 (2006.01)
G06F 17/30 (2006.01)(52) **G06Q 30/02** (2006.01)
H04L 29/08 (2006.01)
U.S. Cl.
CPC **G06N 5/02** (2013.01); **G06Q 30/02**
(2013.01); **G06N 5/003** (2013.01); **H04L**
67/22 (2013.01); **G06Q 30/0269** (2013.01);
G06N 99/005 (2013.01); **G06F 17/30702**
(2013.01); **G06N 3/02** (2013.01); **G06Q**
30/0224 (2013.01)(57) **ABSTRACT**

Systems and methods of artificial intelligence based recommendation are disclosed herein. The system can include: a user device including: a network interface; and an I/O subsystem. The system can include an artificial intelligence engine that can provide a remediation dialogue. The system can include a content management server that can: receive a user identification identifying a user from the user device; retrieve user information from a memory; identify and deliver a question to the user device based on the retrieved user information; receive a response to the delivered question from the user device; determine that the received response is incorrect; trigger the launch of the artificial intelligence engine; receive an indication of completion of the dialogue; and provide a second question after receipt of the indication of completion of the dialogue.



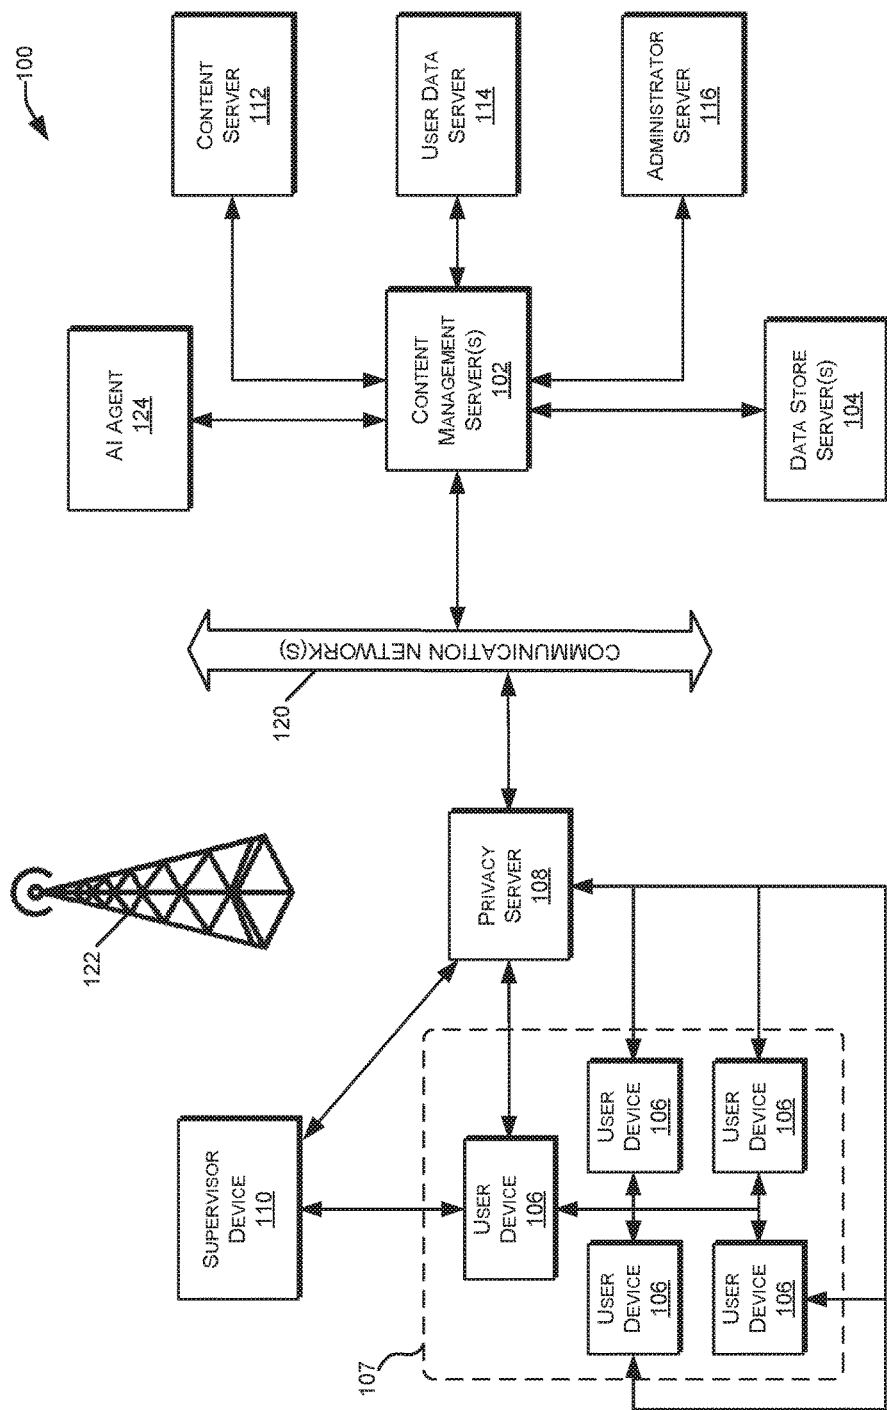


FIG. 1

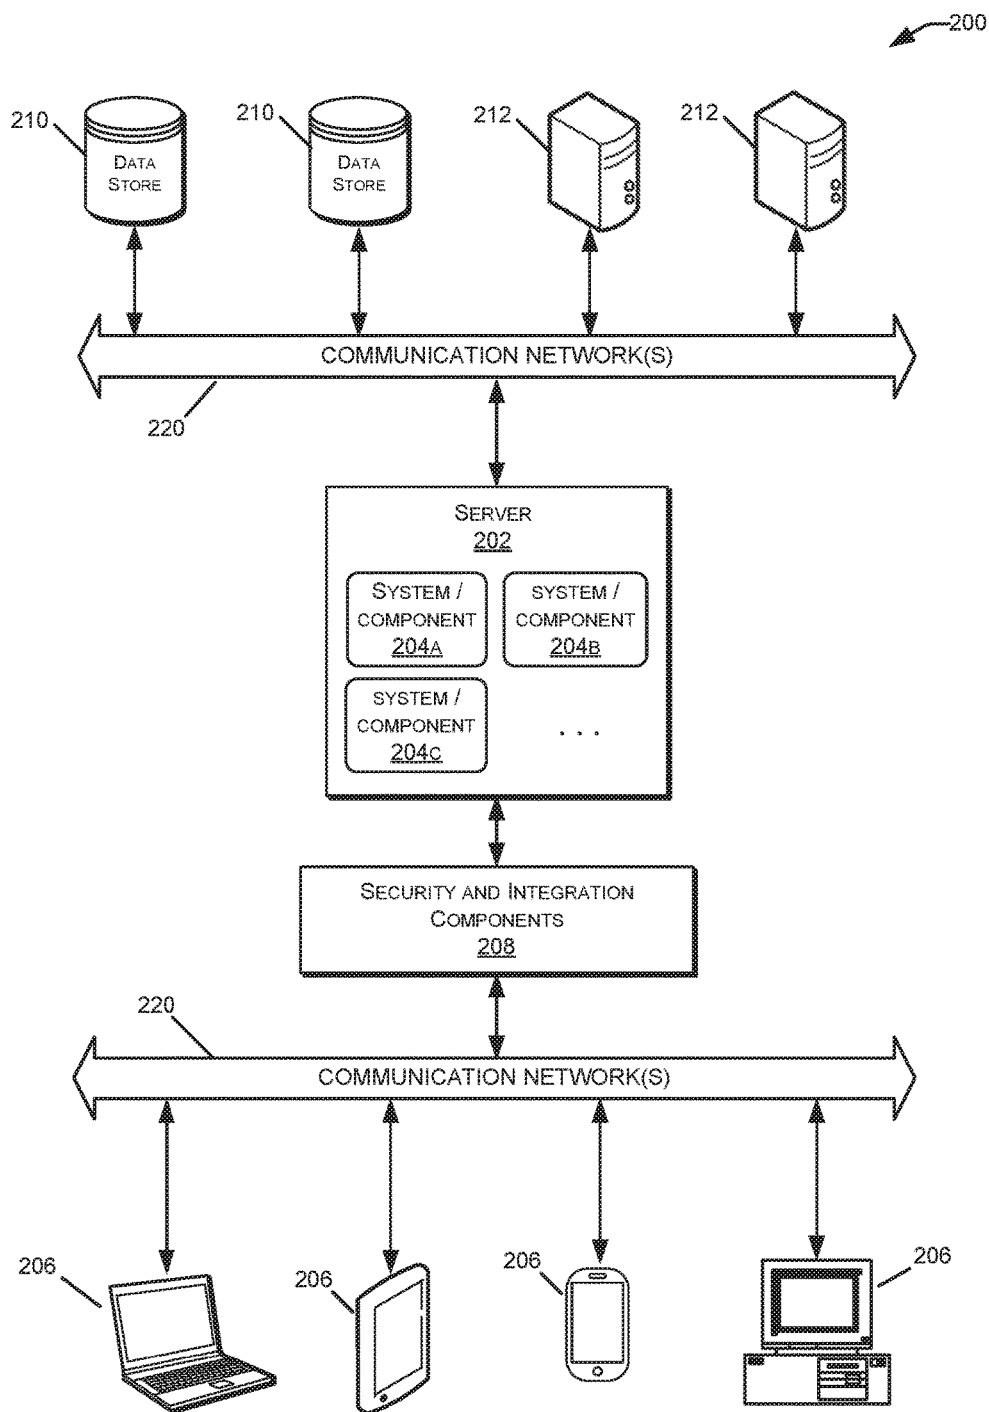


FIG. 2

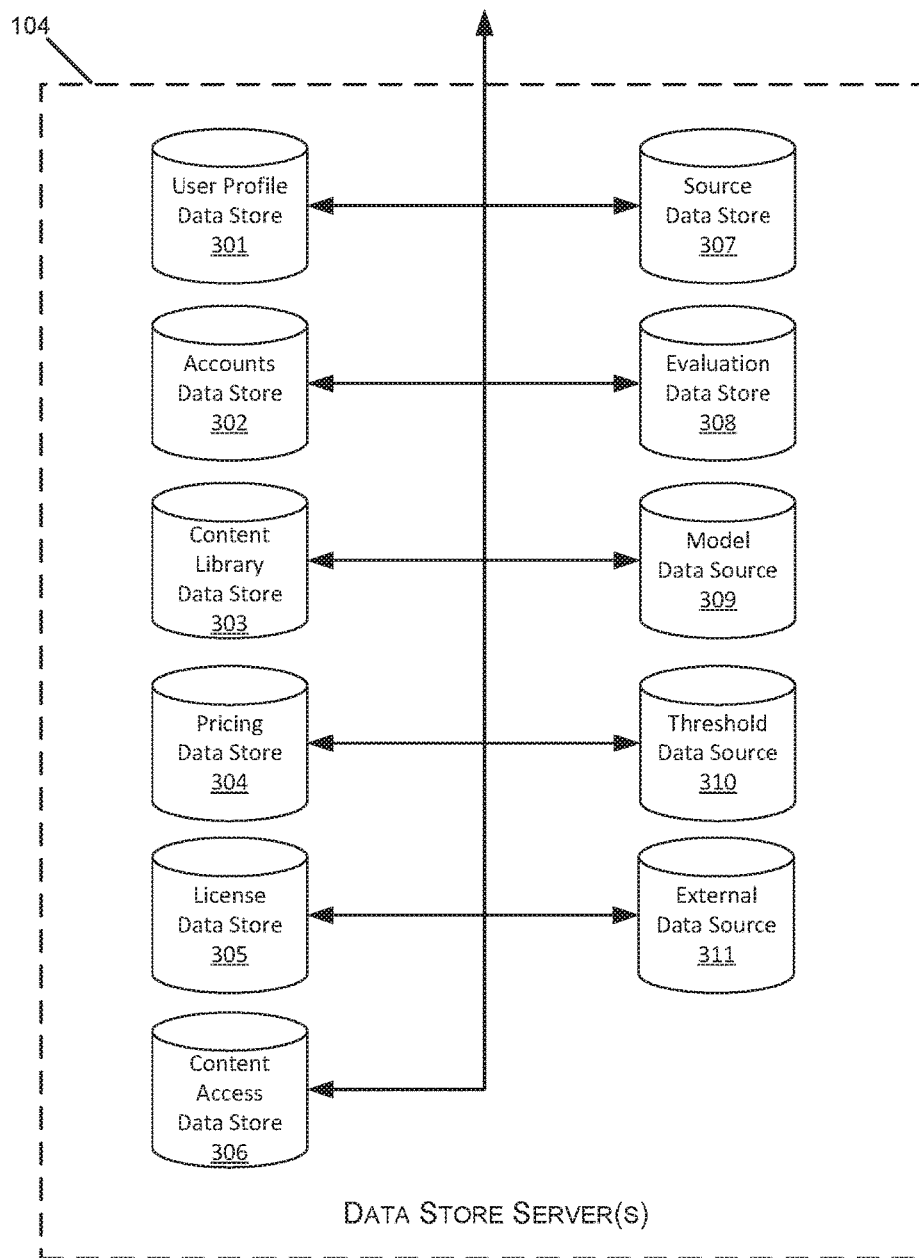


FIG. 3

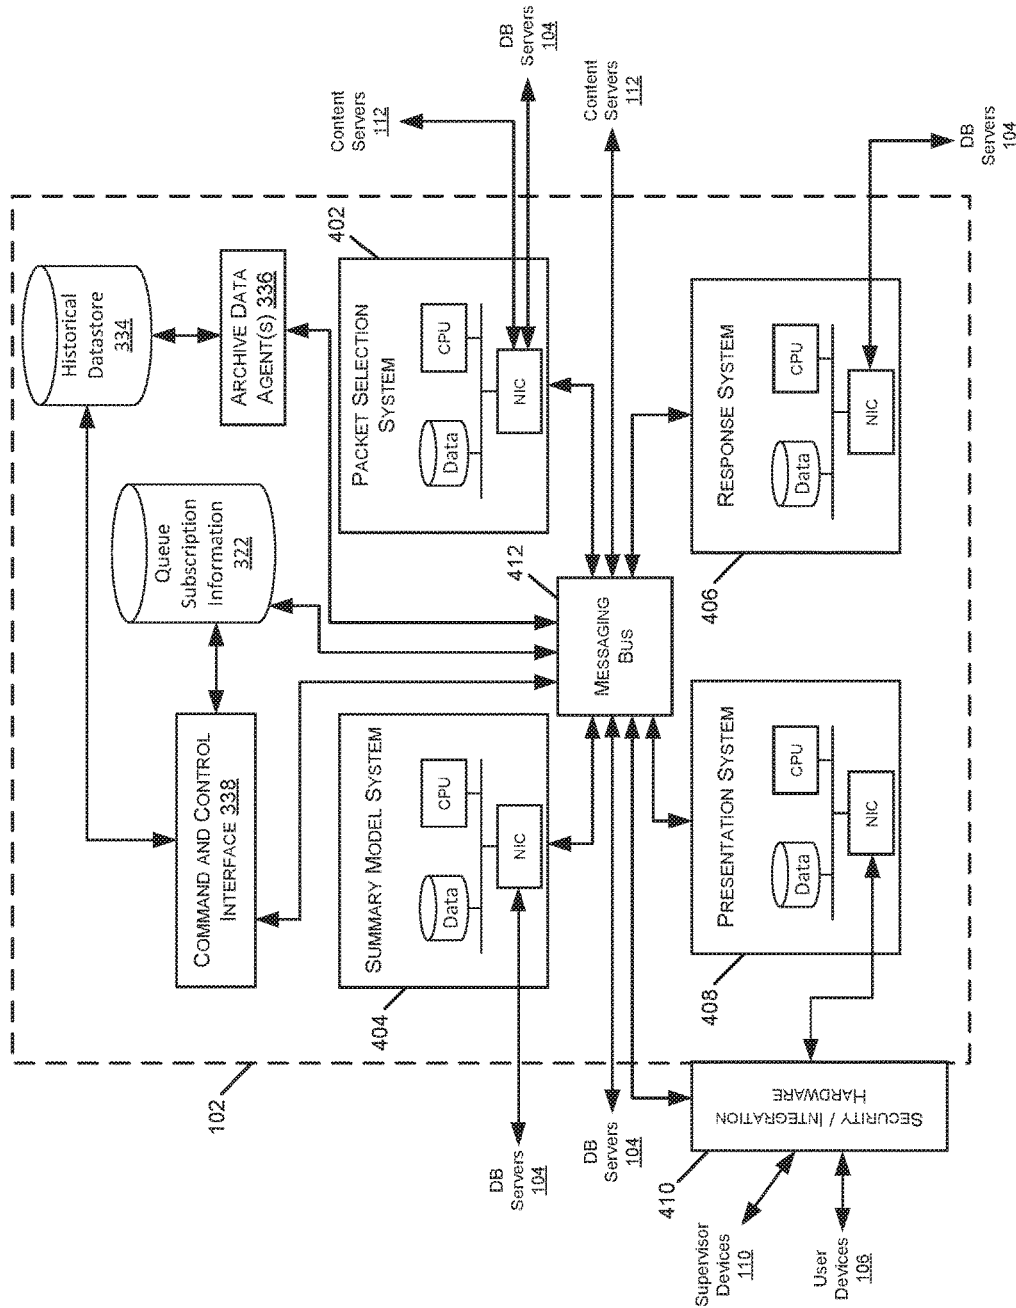


FIG. 4

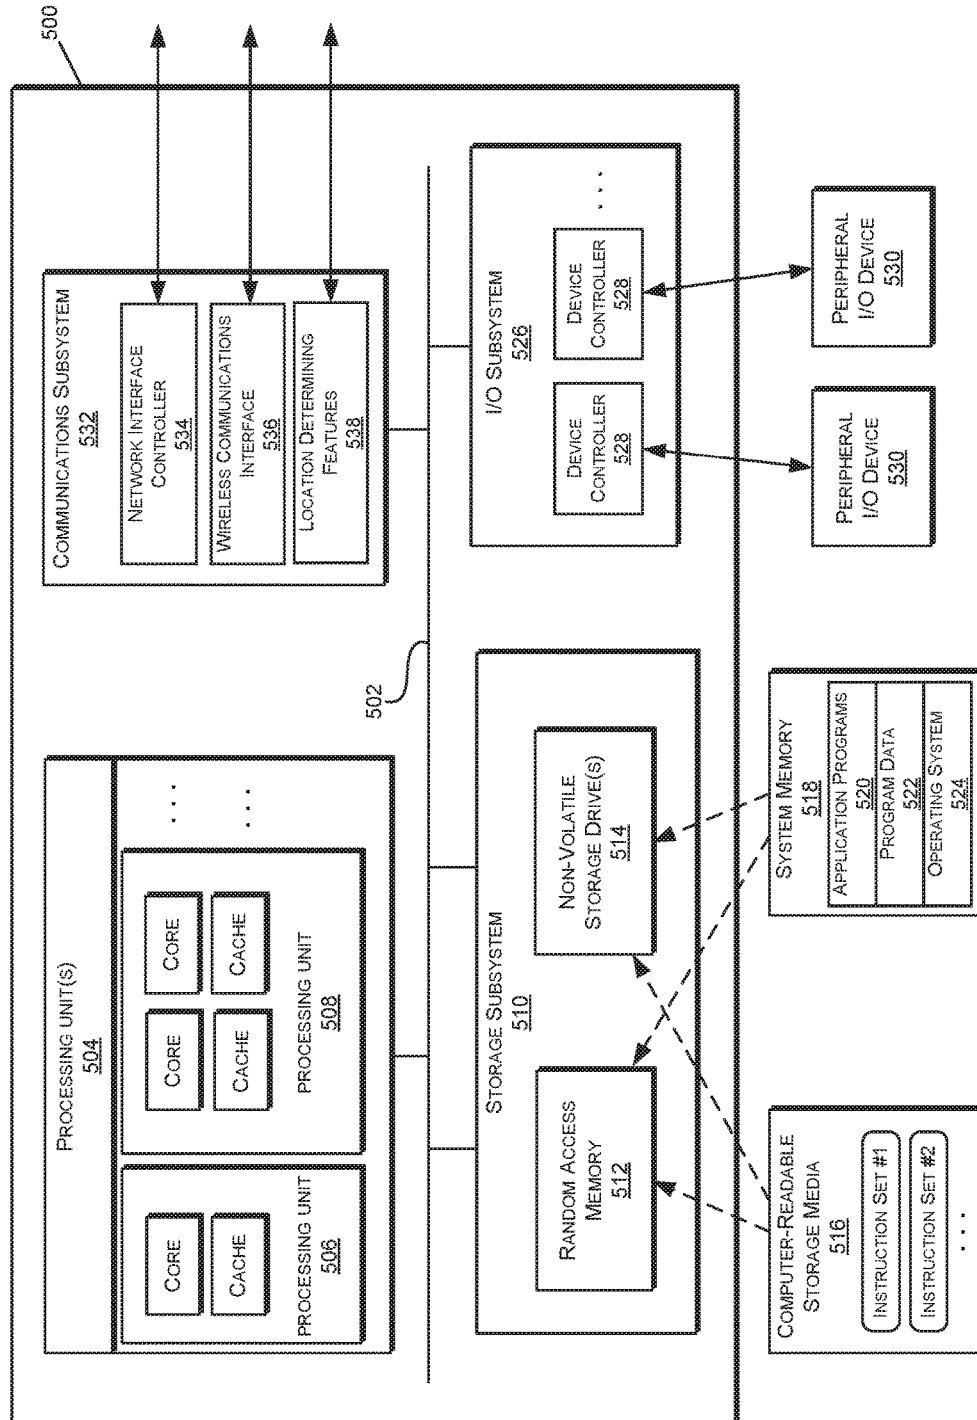


FIG. 5

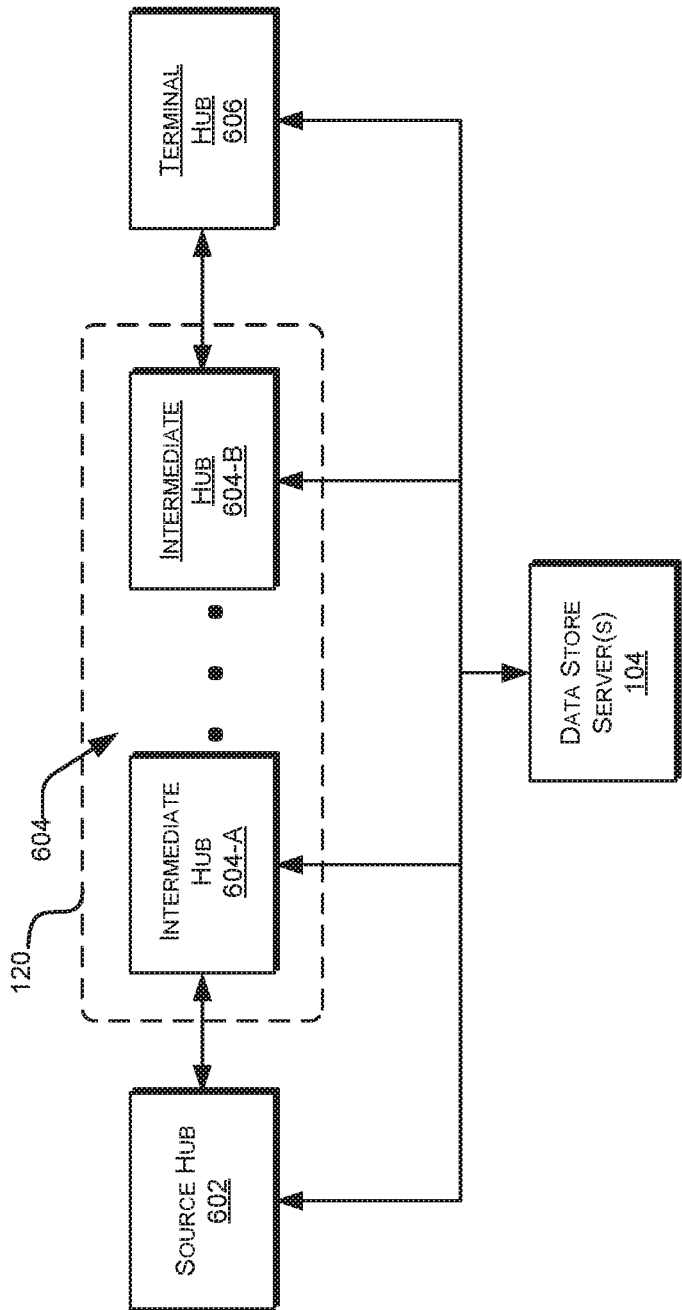


FIG. 6

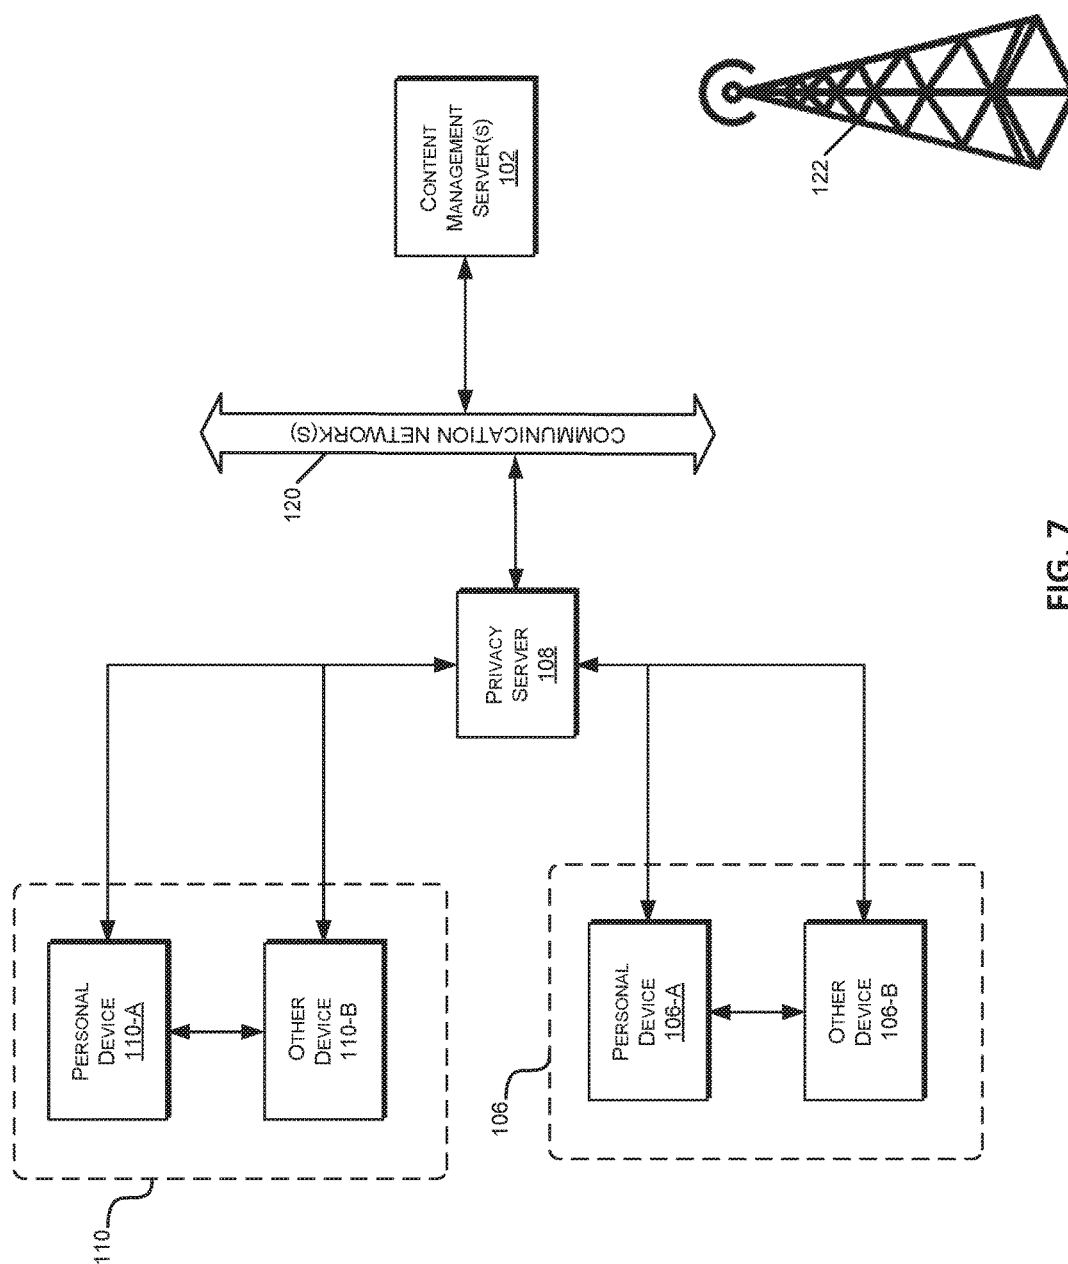


FIG. 7

650

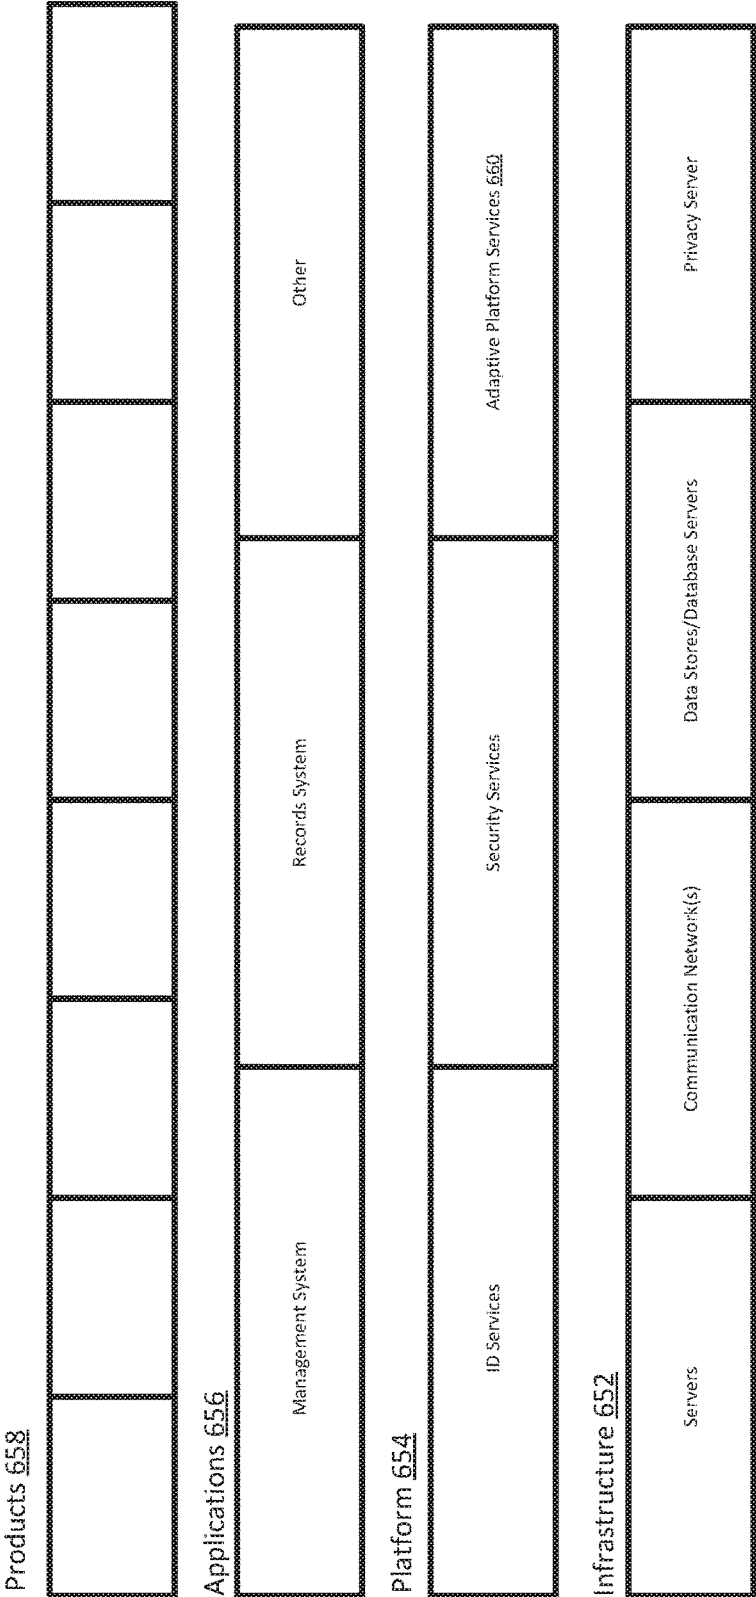


FIG. 8

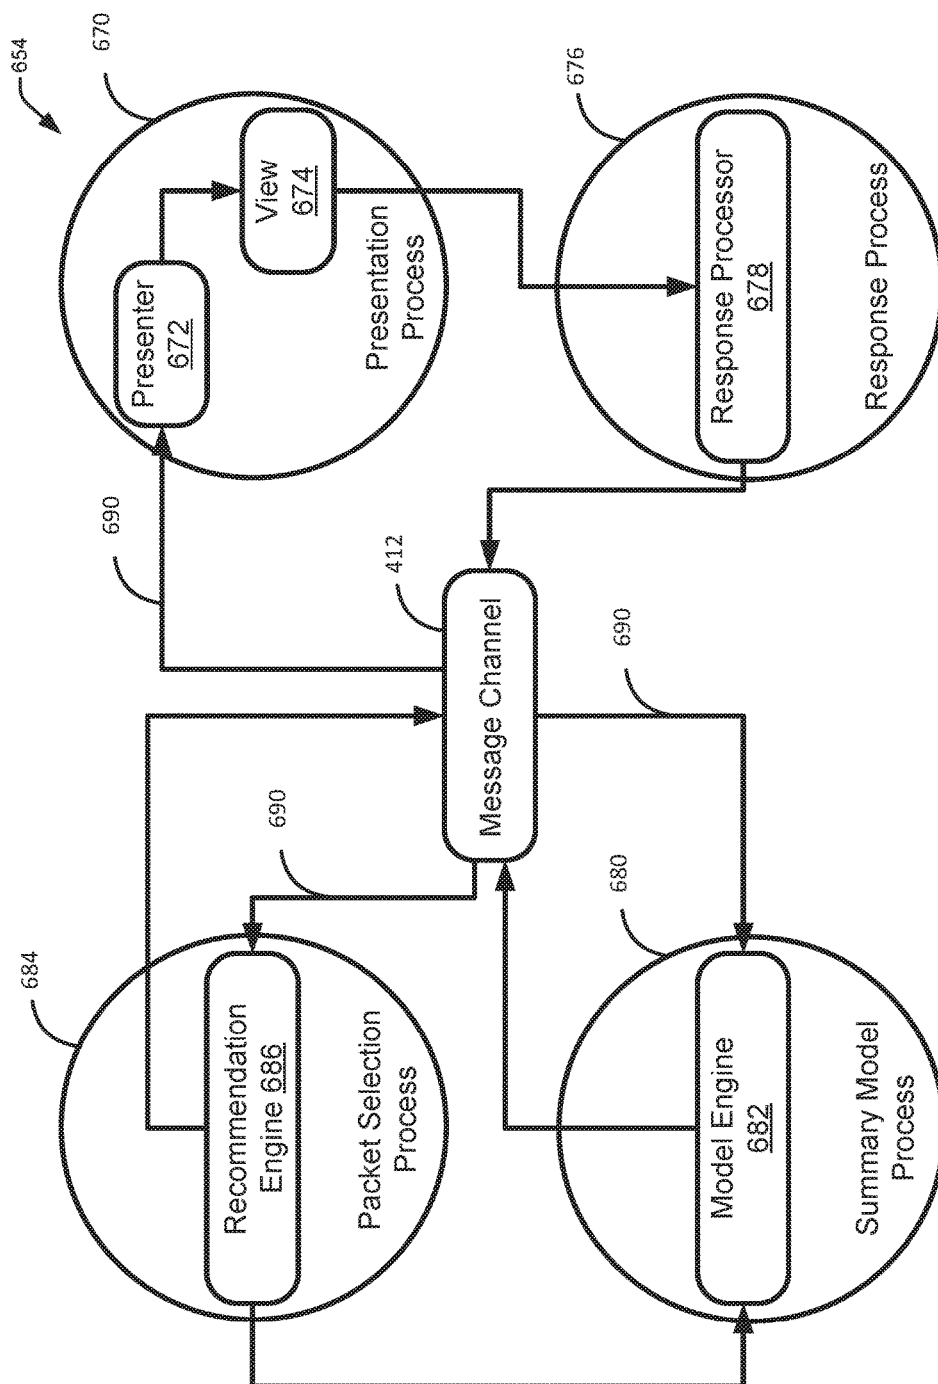


FIG. 9

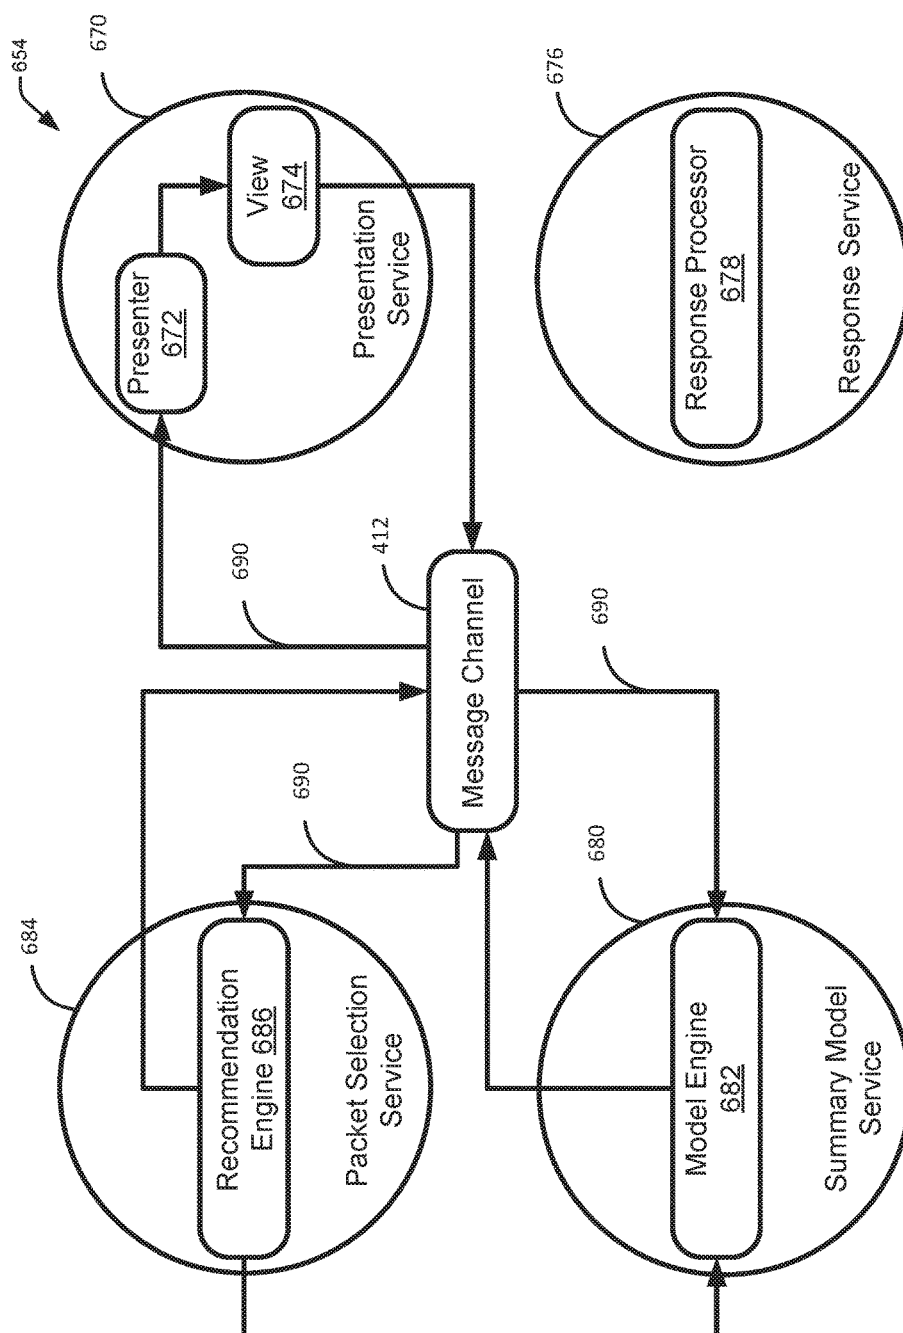


FIG. 10

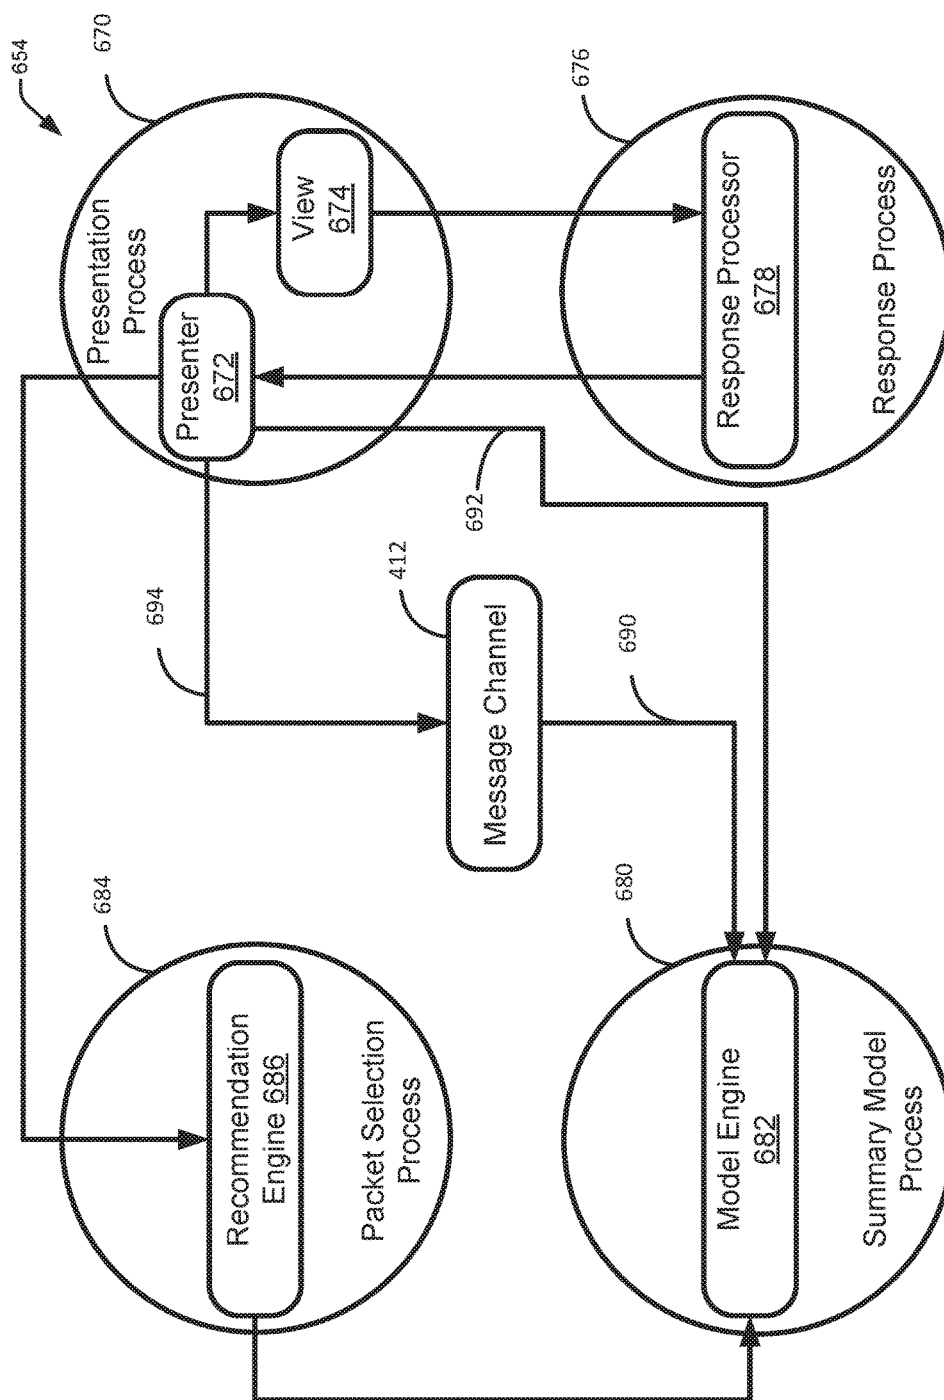


FIG. 11

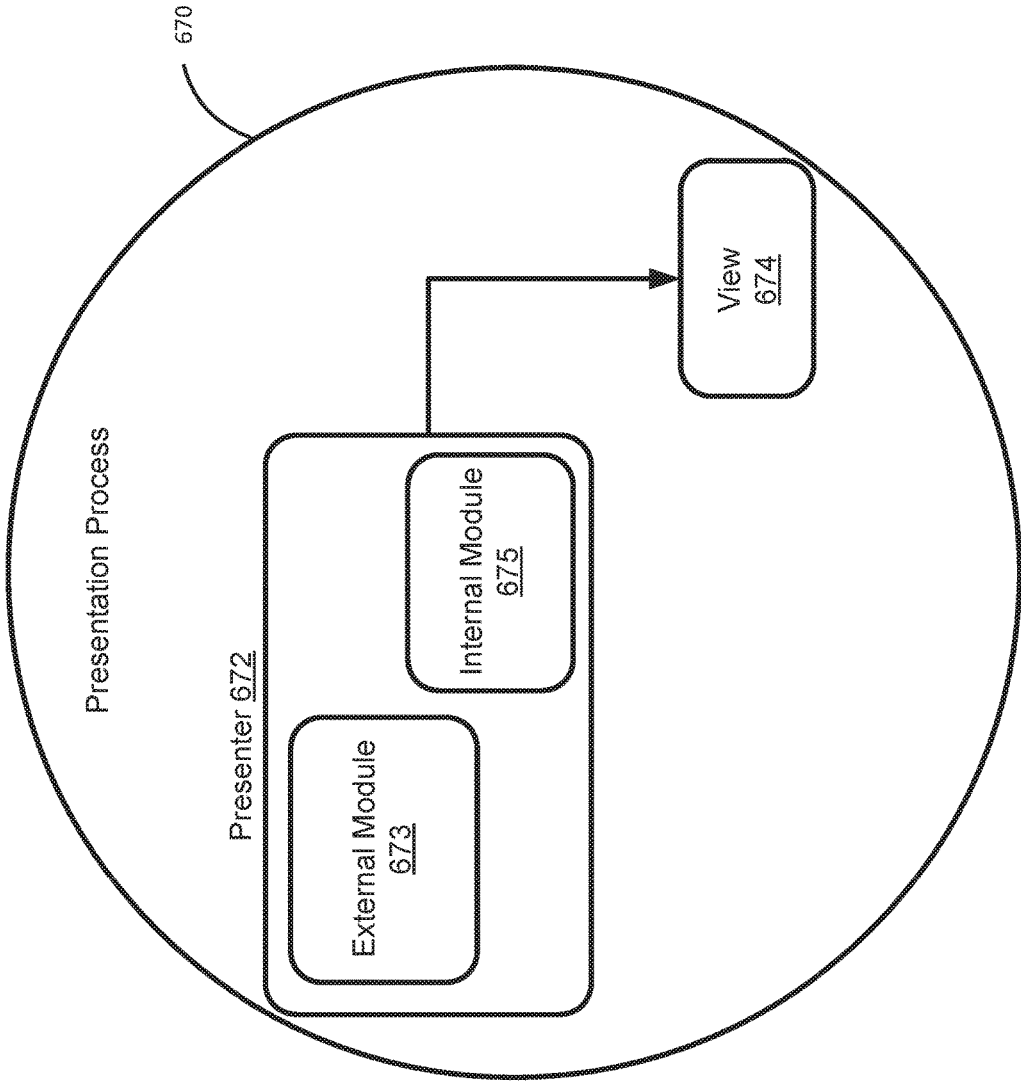


FIG. 12

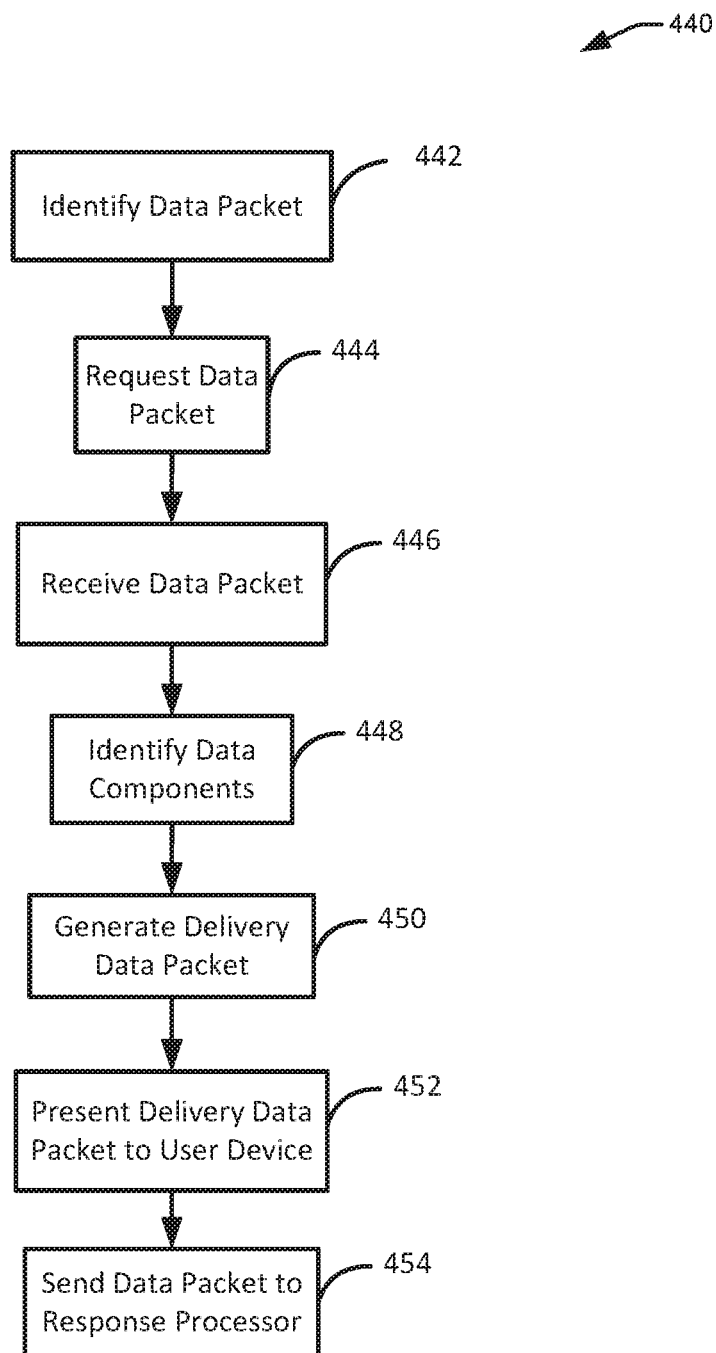
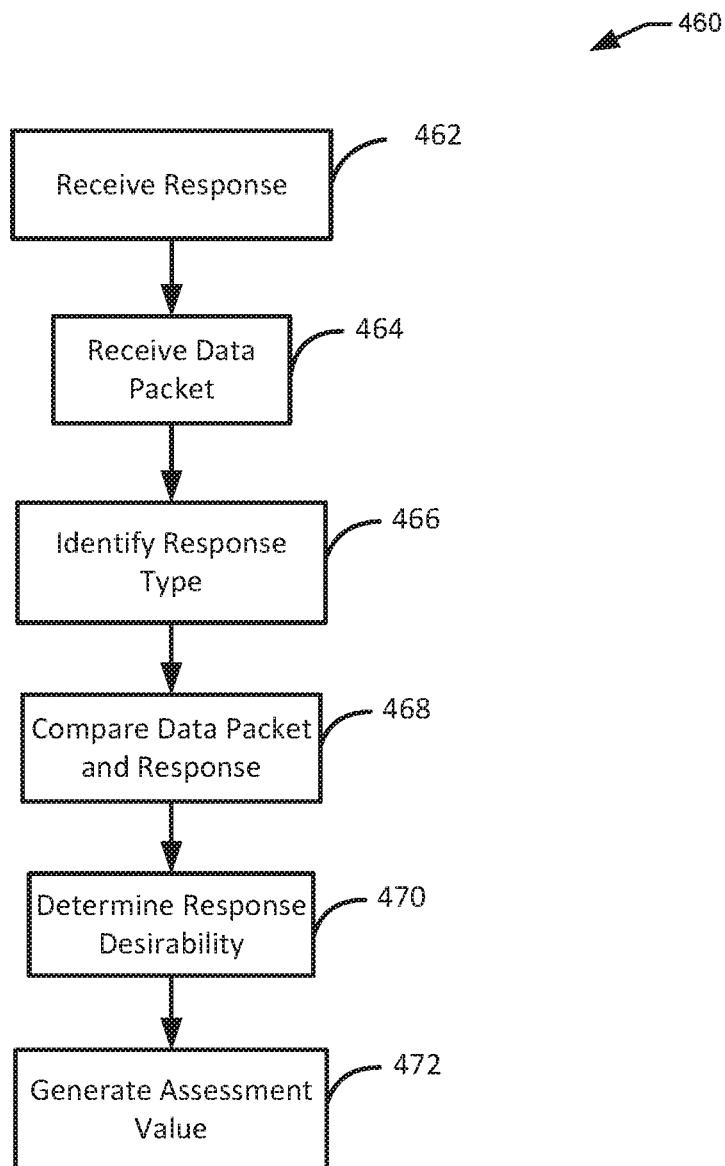


FIG. 13

**FIG. 14**

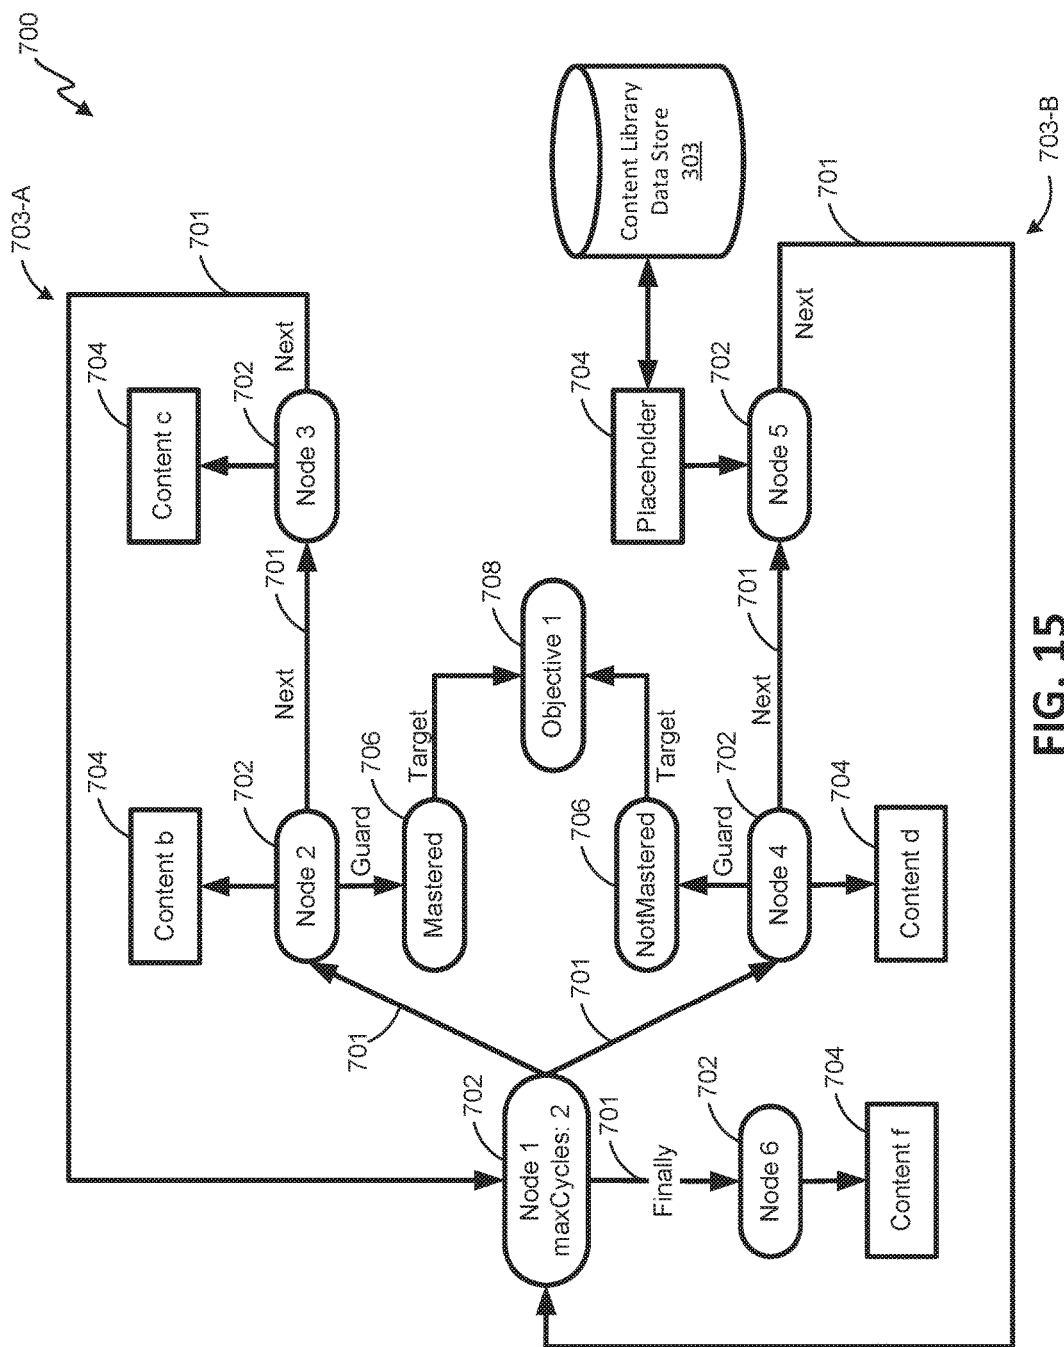


FIG. 15

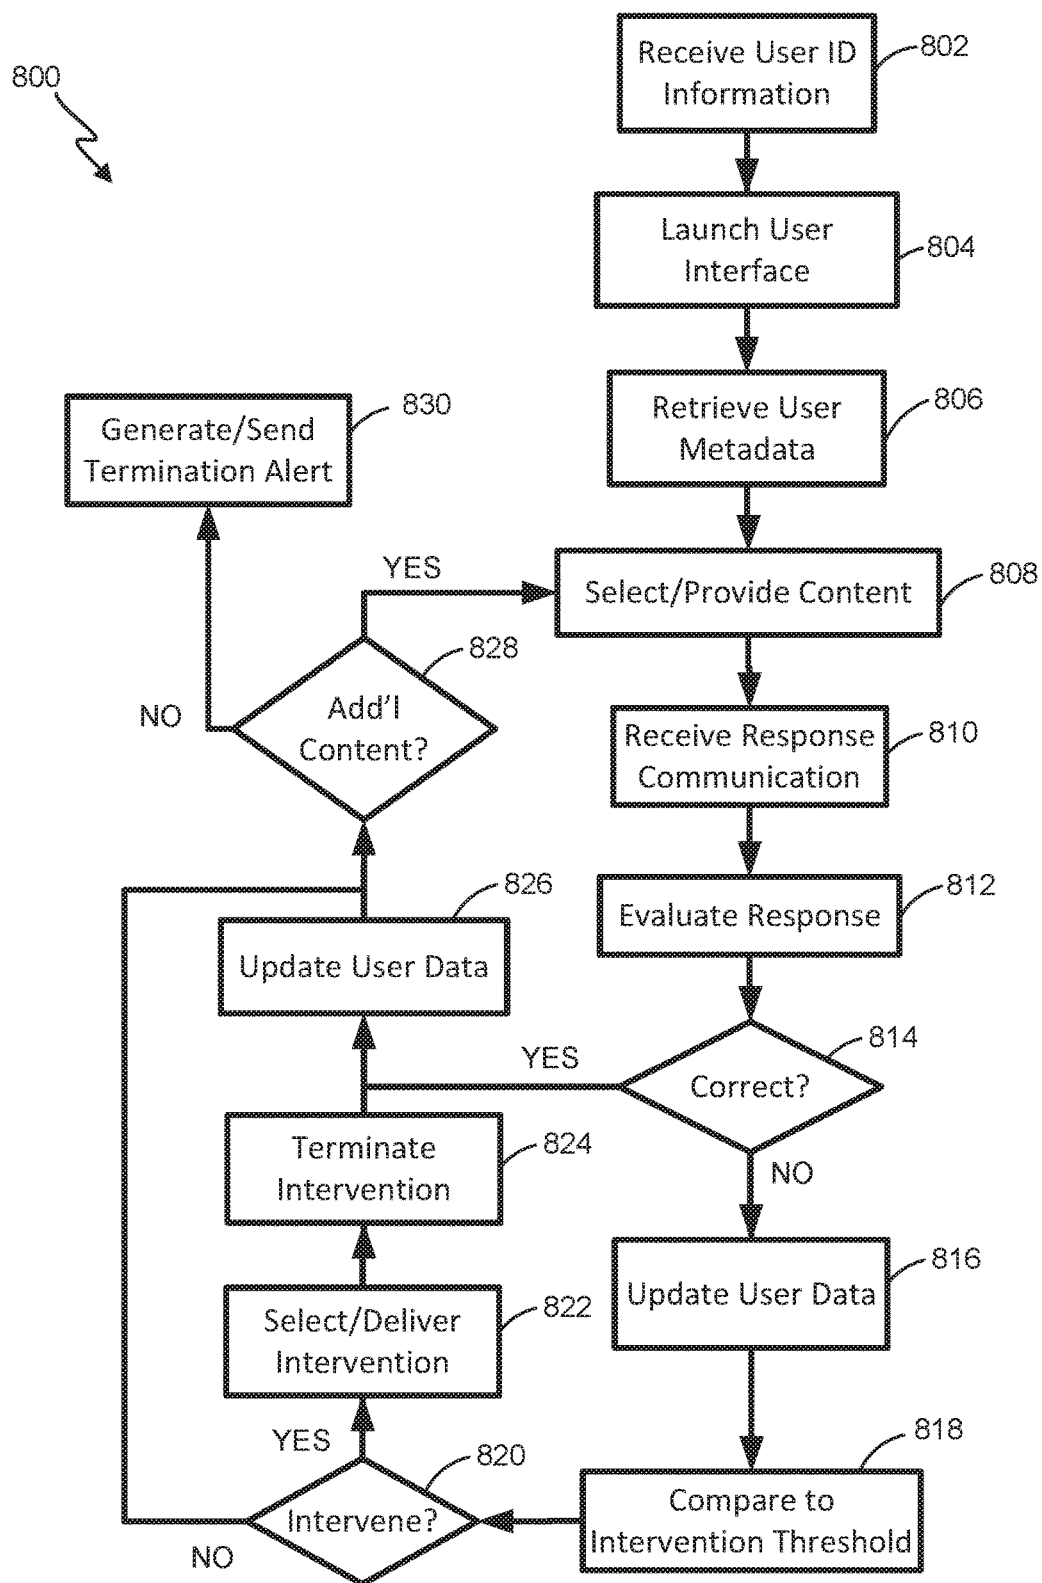


FIG. 16

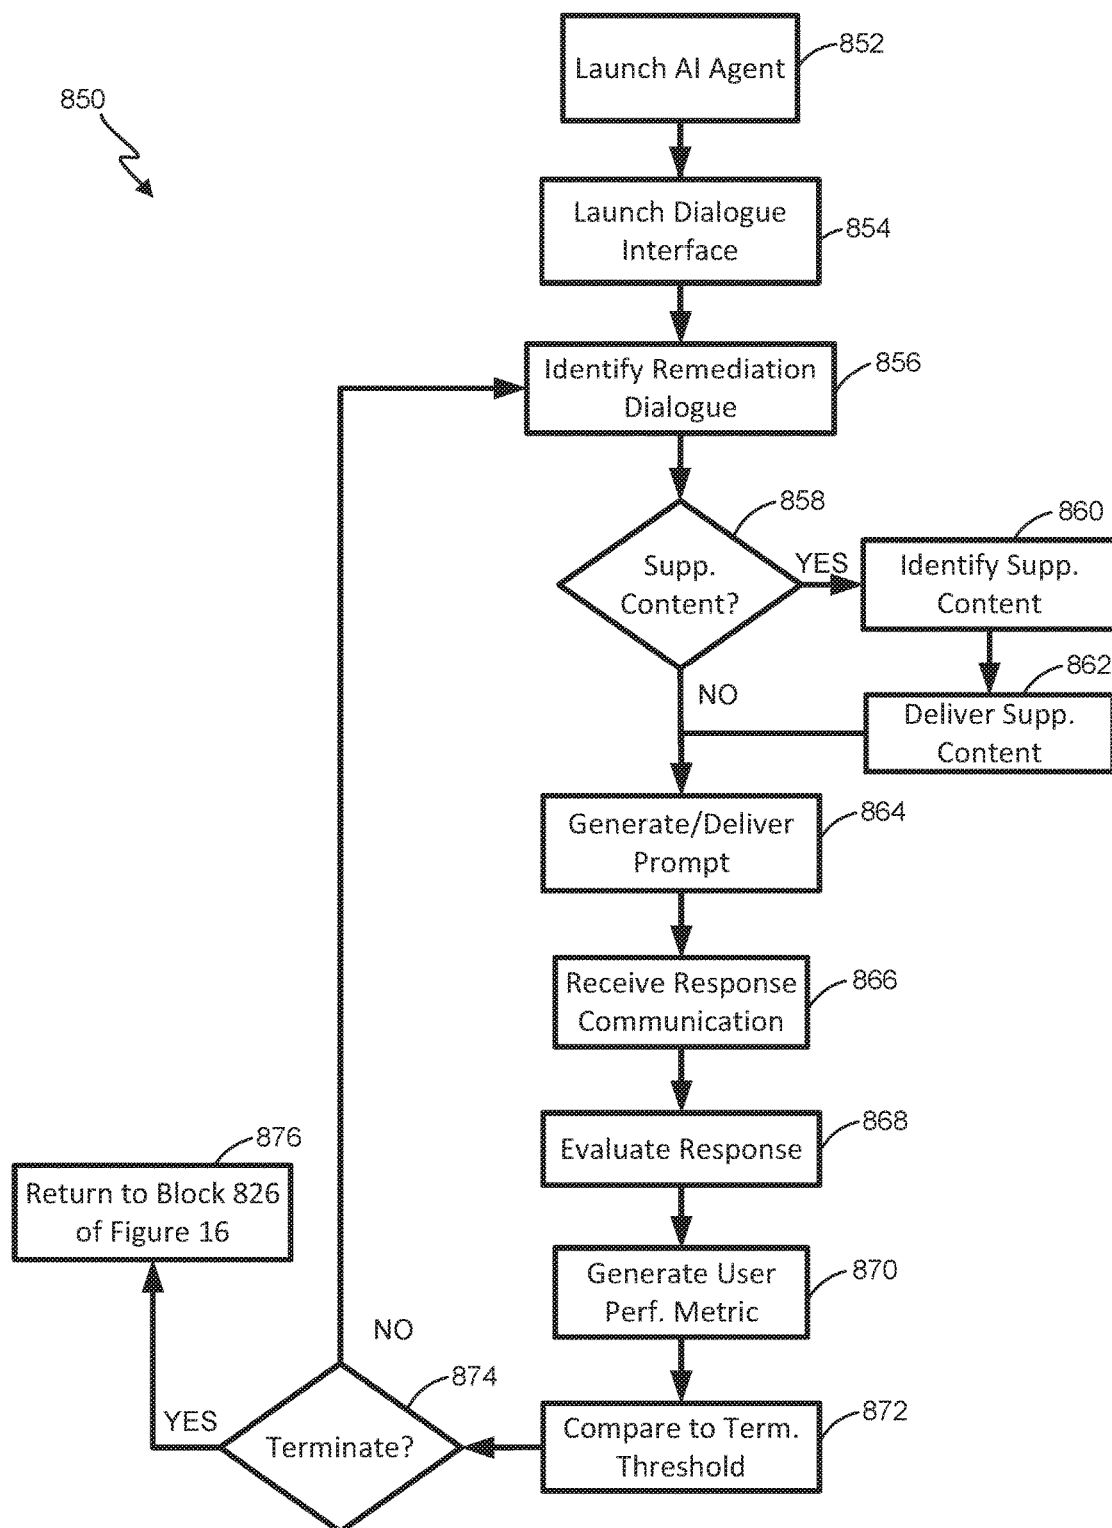
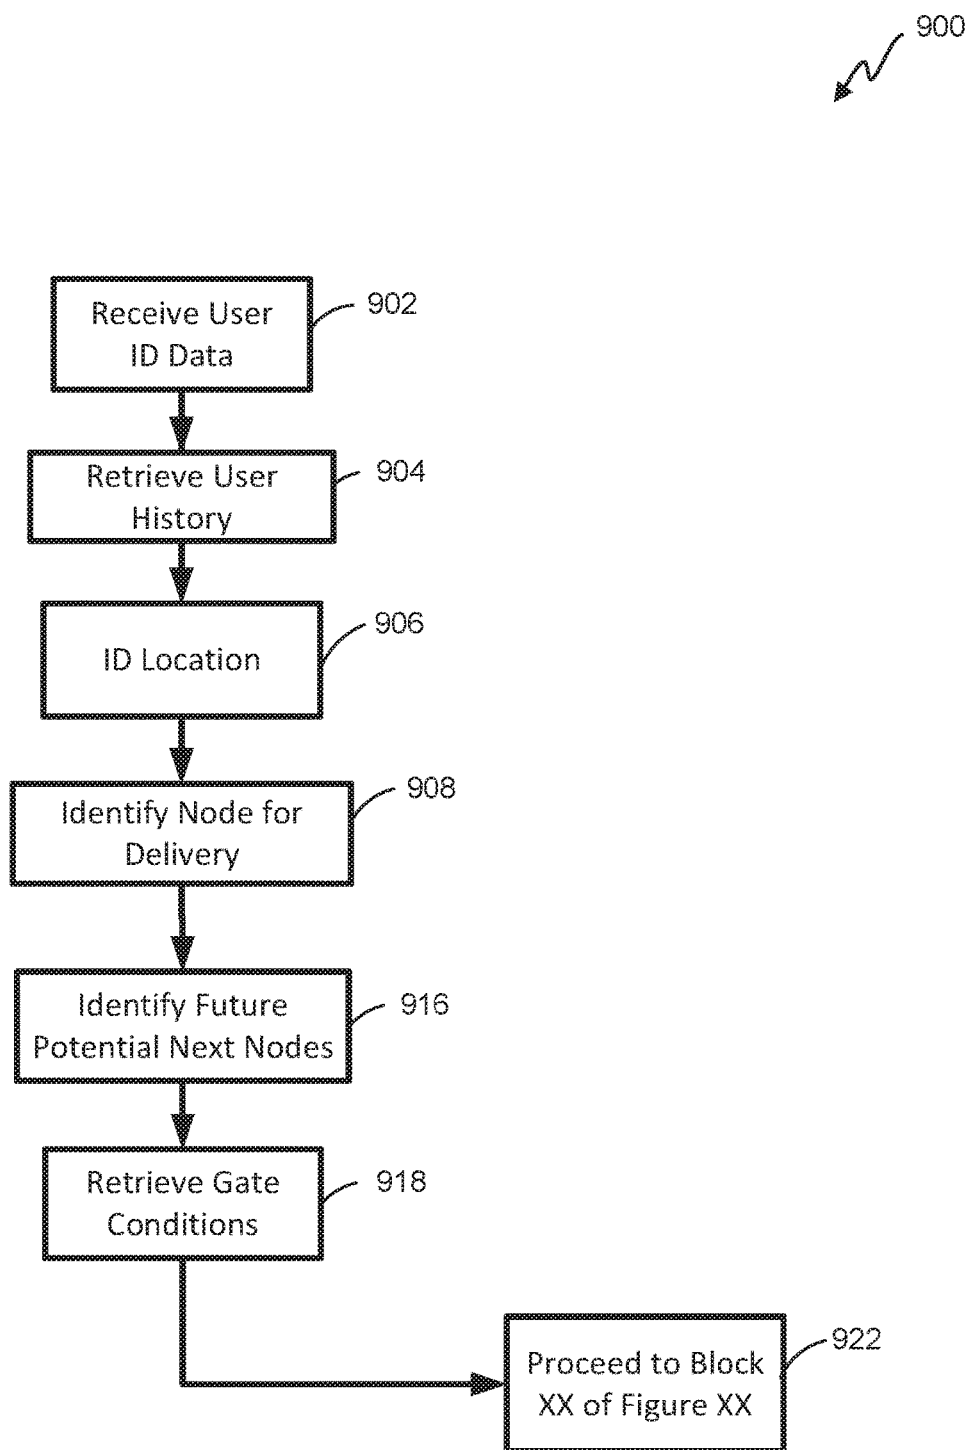


FIG. 17

**FIG. 18**

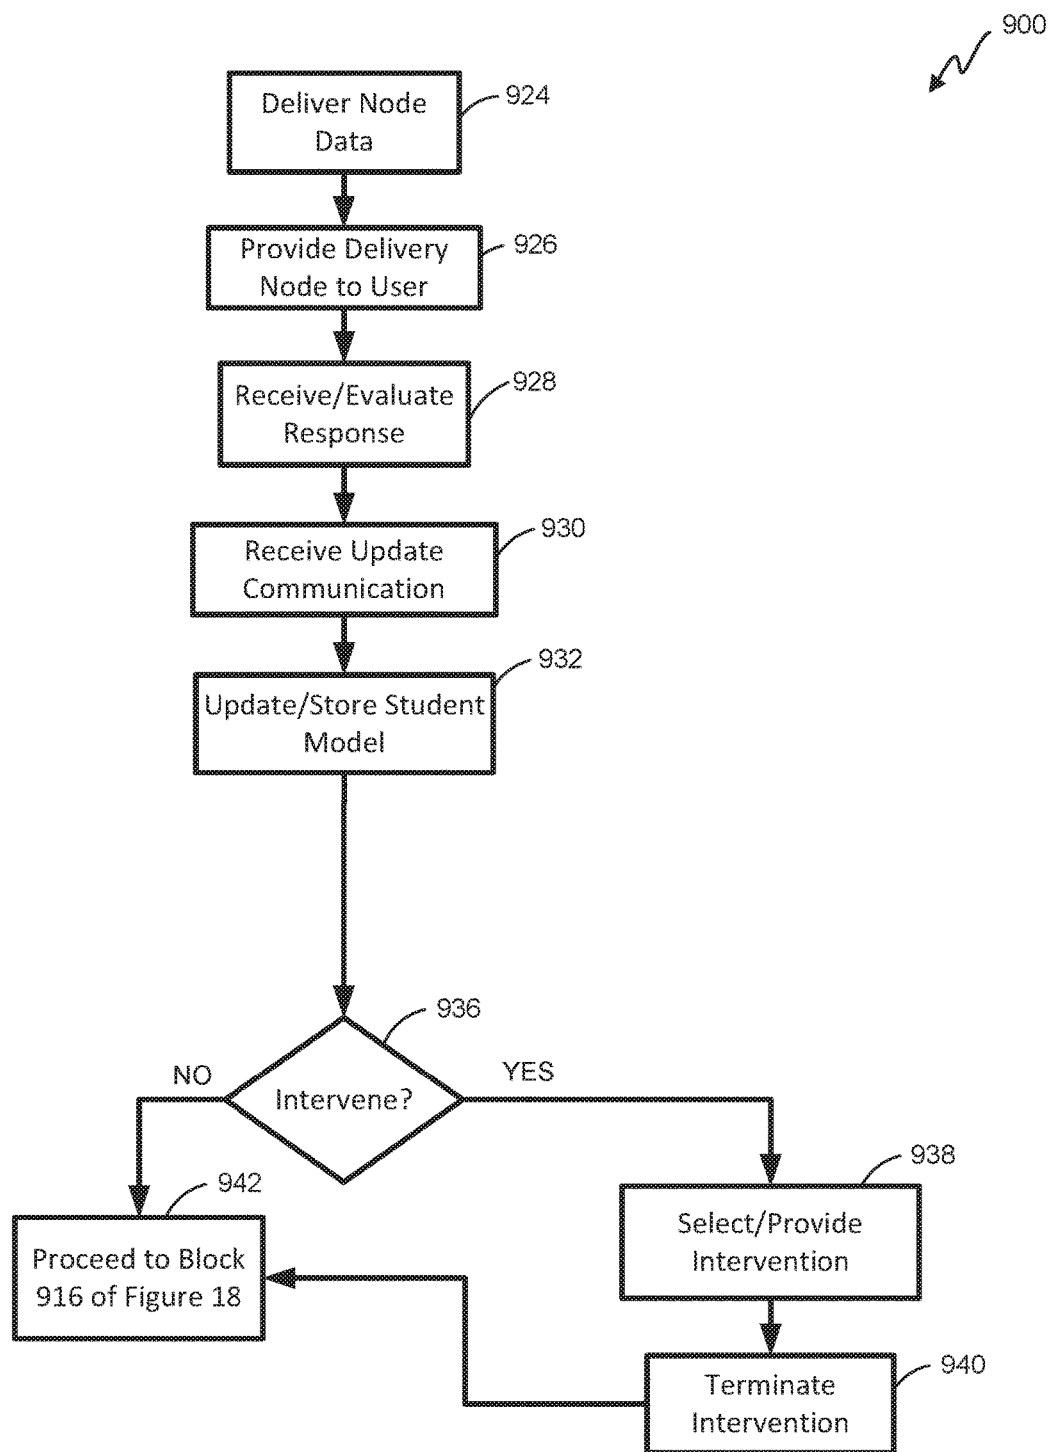


FIG. 19

METHOD AND SYSTEM FOR AUTOMATED CONTENT SELECTION AND DELIVERY

CROSS-REFERENCES TO RELATED APPLICATIONS

[0001] This application claims the benefit of U.S. Provisional Application No. 62/320,213, filed on Apr. 8, 2016, and entitled “ADAPTIVE PATHWAYS AND COGNITIVE TUTORING”, the entirety of each of which is hereby incorporated by reference herein.

BACKGROUND

[0002] A computer network or data network is a telecommunications network which allows computers to exchange data. In computer networks, networked computing devices exchange data with each other along network links (data connections). The connections between nodes are established using either cable media or wireless media. The best-known computer network is the Internet.

[0003] Network computer devices that originate, route and terminate the data are called network nodes. Nodes can include hosts such as personal computers, phones, servers as well as networking hardware. Two such devices can be said to be networked together when one device is able to exchange information with the other device, whether or not they have a direct connection to each other.

[0004] Computer networks differ in the transmission media used to carry their signals, the communications protocols to organize network traffic, the network's size, topology and organizational intent. In most cases, communications protocols are layered on other more specific or more general communications protocols, except for the physical layer that directly deals with the transmission media.

BRIEF SUMMARY

[0005] In one aspect, the present disclosure relates an artificial intelligence based recommendation system. The system includes a user device having: a network interface that can exchange data via the communication network; and an I/O subsystem that can convert electrical signals to user interpretable outputs via a user interface. The system includes an artificial intelligence engine that can receive inputs from at least the user device and that can: launch a dialogue based on the received inputs; identify remediation dialogue based on received user inputs; and terminate the dialogue based on at least one halting criterion. In some embodiments, terminating the dialogue can be based on a user performance metric indicative of a user skill level. In some embodiments, the dialogue is determined for termination when the user performance metric identifies a sufficient user skill level. The system includes a content management server. The content management server can: receive a user identification identifying a user from the user device; receive data indicative of a user pattern of interaction; generate a trigger value characterizing the user pattern of interaction; trigger the launch of the artificial intelligence engine; receive an indication of completion of the dialogue; and provide second content after receipt of the indication of completion of the dialogue.

[0006] In some embodiments, the content management server can be configured to: retrieve user information from a memory, which user information identifies progress by the user through a content progression; identify and deliver a

question to the user device based on the retrieved user information; receive a response to the delivered question from the user device; and determine that the received response is incorrect. In some embodiments, the content management server can receive an indication of an updated user skill level with the indication of completion of the dialogue. In some embodiments, the content management server can select the second content based on the user skill level. In some embodiments, the remediation dialogue includes multiple levels of interrogation and response.

[0007] In some embodiments, at least some of the levels of interrogation include a question and delivered content. In some embodiments, the delivered content corresponds to the user skill level as identified in a user model at the time of selection of the delivered. In some embodiments, the delivered content corresponds to a request received from the user as part of the remediation dialogue.

[0008] In some embodiments, the artificial intelligence engine can reevaluate the user skill level subsequent to each level of interrogation and response. In some embodiments, the reevaluation of the user skill level includes application of a natural language processing algorithm to the received response. In some embodiments, the natural language processing algorithm includes at least one of: a speech recognition algorithm; and natural language understanding algorithm. In some embodiments, launching the dialogue includes launching a user interface on the user device, wherein the user interface can display dialogue content and receive user responses to the displayed dialogue content.

[0009] One aspect of the present disclosure relates to a method for artificial intelligence based content recommendation and provisioning. The method includes: receiving at a content management server a user identification from a user device, which user identification identifies a user; receiving data indicative of a user pattern of interaction; generating a trigger value characterizing the user pattern of interaction; triggering the launch of an artificial intelligence engine; launching a dialogue based on the received inputs with the artificial intelligence engine; identifying remediation dialogue with the artificial intelligence engine based on received user inputs; terminating the dialogue with the artificial intelligence engine; receiving an indication of completion of the dialogue with the content management server; and delivering second content to the user device after receipt of the indication of completion of the dialogue.

[0010] In some embodiments, the dialogue is terminated based on at least one halting criterion. In some embodiments, the at least one halting criterion includes a termination threshold, and terminating the dialogue can include generating a user performance metric indicative of at least one of: a user skill level; a correct response; and an incorrect response. In some embodiments terminating the dialogue can include comparing the user performance metric to the termination threshold.

[0011] In some embodiments the method includes receiving an indication of an updated user skill level with the indication of completion of the dialogue at the content management server. In some embodiments the method includes, selecting the second content with the content management server, which second question is selected based on the user skill level.

[0012] In some embodiments, the remediation dialogue includes multiple levels of interrogation and response. In some embodiments, at least some of the levels of interro-

gation include a question and delivered content. In some embodiments, the delivered content corresponds to the user skill level as identified in a user model at the time of selection of the delivered content. In some embodiments, the delivered content corresponds to a request received from the user as part of the remediation dialogue.

[0013] In some embodiments the method includes reevaluating the user skill level subsequent to each level of interrogation and response with the artificial intelligence engine. In some embodiments, reevaluating the user skill level includes application of a natural language processing algorithm to the received response. In some embodiments, the natural language processing algorithm includes at least one of: a speech recognition algorithm; and natural language understanding algorithm. In some embodiments, launching the dialogue includes launching a user interface on the user device, which user interface can display dialogue content and receive user responses to the displayed dialogue content.

[0014] In some embodiments, the method includes: retrieving with the content management server user information from a memory, which user information identifies progress by the user through a content progression; identifying and delivering a question to the user device based on the retrieved user information; receiving a response to the delivered question from the user device; and determining that the received response is incorrect. In some embodiments, the second content comprises a second question. In some embodiments, the launch of the artificial intelligence engine is triggered via at least one triggering condition. In some embodiments, the triggering condition delineates between patterns of user interaction.

[0015] One aspect of the present disclosure relates to a recommendation system. The recommendation system includes: memory including: a user profile database including information relating to a plurality of users, which information relating to the plurality of users comprises unique user history data associated with each of the users in the plurality of users; and a content library database including a plurality of nodes linked in a plurality of sequential relationships in a content network. In some embodiments, some of the plurality of nodes is associated with guard conditions. In some embodiments, each guard condition identifies at least one of: an entrance condition controlling entrance into the node associated with the guard condition; and an exit condition controlling exit from the node associated with the guard condition. The system can include a user device including: a network interface that can exchange data via the communication network; and an I/O subsystem that can convert electrical signals to user interpretable outputs via a user interface. The system can include one or more servers that can: receive user identification information, which user identification information identifies a user; determine a location of the user identified by the received user identification information in the content network; select a next node based on the location of the user in the content network and on at least some of the guard conditions; deliver next node content to the user device via a network such as, for example, wireless network, which next node content is adaptively selected from a plurality of selectable content options when the next node is a placeholder node, and which next content is directly linked to the next node when the next node is not a placeholder node.

[0016] In some embodiments, at least some of the guard conditions regulate entry into a learning sequence including

some of the plurality of nodes. In some embodiments, at least some of the guard conditions regulate exit from a learning sequence including some of the plurality of nodes. In some embodiments, the system includes an artificial intelligence engine that can receive inputs from at least the user device. In some embodiments, at least one of the selectable content options of the placeholder node comprises a dialogue with the artificial intelligence engine. In some embodiments, the artificial intelligence engine can: launch the dialogue; deliver interrogation and content based on the received user responses; and determine termination of the dialogue based on a user performance metric indicative of a user skill level, which dialogue is determined for termination when the user performance metric identifies a sufficient user skill level.

[0017] In some embodiments, the remediation dialogue includes multiple levels of interrogation and response. In some embodiments, the artificial intelligence engine can reevaluate the user skill level subsequent to each level of interrogation and response. In some embodiments, the reevaluation of the user skill level includes application of a natural language processing algorithm to the received response, which natural language processing algorithm comprises a speech recognition algorithm and natural language understanding algorithm.

[0018] In some embodiments, the one or several servers can determine a user skill level based on user responses received relating to the delivered next node content. In some embodiments the system includes a supervisor device in communicating connection with the one or several servers via the wireless network, which one or several servers can send an alert to the supervisor device when the determined user skill level drops below a predetermined threshold value, which alert includes computer code that can direct the supervisor device to display the alert upon receipt.

[0019] One aspect of the present disclosure relates to a method of content selection and delivery. The method includes: receiving user identification information at one or several servers from a user device via a communication network such as, for example, a wireless network, which user identification information identifies a user; determining with the one or several servers a location of the user identified by the received user identification information in the content network; selecting with the one or several servers a next node based on the location of the user in the content network and on at least some of the guard conditions; and delivering with the one or several servers next node content to the user device via the communication network such as, for example, the wireless network, which next node content is adaptively selected from a plurality of selectable content options when the next node is a placeholder node, and which next content is directly linked to the next node when the next node is not a placeholder node.

[0020] In some embodiments, at least some of the guard conditions regulate entry into a learning sequence including some of the plurality of nodes. In some embodiments, at least some of the guard conditions regulate exit from a learning sequence including some of the plurality of nodes. In some embodiments, the next content can be an intervention. In some embodiments, the intervention can be at least one of: targeted remediation content; a scaffolding hint; question generation; answer specific feedback; a text; a text summary; and a follow-up question.

[0021] In some embodiments, the next content directly linked to the next node can include a dialogue with an artificial intelligence engine, which artificial intelligence engine can receive inputs from at least the user device. In some embodiments, at least one of the selectable content options of the placeholder node can be a dialogue with an artificial intelligence engine, which artificial intelligence engine can receive inputs from at least the user device.

[0022] In some embodiments the method includes: launching the dialogue with the artificial intelligence engine; delivering interrogation and content based on the received user responses from the artificial intelligence engine to the user device; and determining termination of the dialogue with the artificial intelligence engine based on at least one halting criteria. In some embodiments, the halting criteria can include a termination threshold. In some embodiments, terminating the dialogue includes: generating a user performance metric indicative of at least one of: a skill level; a correct response; and an incorrect response; and comparing the user performance metric to the termination threshold. In some embodiments, the dialogue is terminated when the user performance metric identifies a sufficient user skill level.

[0023] In some embodiments, the remediation dialogue includes multiple levels of interrogation and response. In some embodiments the method includes reevaluating the user skill level with the artificial intelligence engine subsequent to each level of interrogation and response. In some embodiments, the reevaluation of the user skill level includes application by the one or several servers of a natural language processing algorithm to the received response, which natural language processing algorithm comprises a speech recognition algorithm and natural language understanding algorithm. In some embodiments the method includes determining a user skill level with the one or several servers based on user responses received relating to the delivered next node content. In some embodiments the method includes sending an alert from the one or several servers to a supervisor device when the determined user skill level drops below a predetermined threshold value, which alert includes computer code directing the supervisor device to display the alert upon receipt.

[0024] Further areas of applicability of the present disclosure will become apparent from the detailed description provided hereinafter. It should be understood that the detailed description and specific examples, while indicating various embodiments, are intended for purposes of illustration only and are not intended to necessarily limit the scope of the disclosure.

BRIEF DESCRIPTION OF THE DRAWINGS

[0025] FIG. 1 is a block diagram illustrating an example of a content distribution network.

[0026] FIG. 2 is a block diagram illustrating a computer server and computing environment within a content distribution network.

[0027] FIG. 3 is a block diagram illustrating an embodiment of one or more data store servers within a content distribution network.

[0028] FIG. 4 is a block diagram illustrating an embodiment of one or more content management servers within a content distribution network.

[0029] FIG. 5 is a block diagram illustrating the physical and logical components of a special-purpose computer device within a content distribution network.

[0030] FIG. 6 is a block diagram illustrating one embodiment of the communication network.

[0031] FIG. 7 is a block diagram illustrating one embodiment of user device and supervisor device communication.

[0032] FIG. 8 is a schematic illustration of one embodiment of a computing stack.

[0033] FIG. 9 is a schematic illustration of one embodiment of communication and processing flow of modules within the content distribution network.

[0034] FIG. 10 is a schematic illustration of another embodiment of communication and processing flow of modules within the content distribution network.

[0035] FIG. 11 is a schematic illustration of another embodiment of communication and processing flow of modules within the content distribution network.

[0036] FIG. 12 is a schematic illustration of one embodiment of the presentation process.

[0037] FIG. 13 is a flowchart illustrating one embodiment of a process for data management.

[0038] FIG. 14 is a flowchart illustrating one embodiment of a process for evaluating a response.

[0039] FIG. 15 is a schematic illustration of one embodiment of the content network.

[0040] FIG. 16 is a flowchart illustrating one embodiment of a process for artificial intelligence based content recommendation and provisioning.

[0041] FIG. 17 is a flowchart illustrating one embodiment of a process for selecting and delivering an intervention.

[0042] FIG. 18 is a flowchart illustrating a first portion of one embodiment of a process for a method of content selection and delivery.

[0043] FIG. 19 is a flowchart illustrating a second portion of one embodiment of a process for a method of content selection and delivery.

[0044] In the appended figures, similar components and/or features may have the same reference label. Further, various components of the same type may be distinguished by following the reference label by a dash and a second label that distinguishes among the similar components. If only the first reference label is used in the specification, the description is applicable to any one of the similar components having the same first reference label irrespective of the second reference label.

DETAILED DESCRIPTION

[0045] The ensuing description provides illustrative embodiment(s) only and is not intended to limit the scope, applicability or configuration of the disclosure. Rather, the ensuing description of the illustrative embodiment(s) will provide those skilled in the art with an enabling description for implementing a preferred exemplary embodiment. It is understood that various changes can be made in the function and arrangement of elements without departing from the spirit and scope as set forth in the appended claims.

[0046] This application is related to U.S. application Ser. No. 15/236,103, [Attorney Docket No. 056838-1014017-0365.01.0.US], filed on Aug. 12, 2016, and entitled SYSTEMS AND METHODS FOR DATA PACKET META-DATA STABILIZATION; U.S. application Ser. No. 15/236,275, [Attorney Docket No. 056838-1014019-0365.04.0.

US], filed on Aug. 12, 2016, and entitled SYSTEM AND METHOD FOR AUTOMATIC CONTENT AGGREGATION

[0047] GENERATION; and U.S. application Ser. No. 15/236,238, [Attorney Docket No. 056838-1014018-0365. 03.0.US], filed on Aug. 12, 2016, and entitled SYSTEM AND METHOD FOR CONTENT PROVISIONING WITH DUAL RECOMMENDATION ENGINES, the entirety of each of which is hereby incorporated by reference herein.

[0048] With reference now to FIG. 1, a block diagram is shown illustrating various components of a content distribution network (CDN) 100 which implements and supports certain embodiments and features described herein. In some embodiments, the content distribution network 100 can comprise one or several physical components and/or one or several virtual components such as, for example, one or several cloud computing components. In some embodiments, the content distribution network 100 can comprise a mixture of physical and cloud computing components.

[0049] Content distribution network 100 may include one or more content management servers 102. As discussed below in more detail, content management servers 102 may be any desired type of server including, for example, a rack server, a tower server, a miniature server, a blade server, a mini rack server, a mobile server, an ultra-dense server, a super server, or the like, and may include various hardware components, for example, a motherboard, a processing unit, memory systems, hard drives, network interfaces, power supplies, etc. Content management server 102 may include one or more server farms, clusters, or any other appropriate arrangement and/or combination of computer servers. Content management server 102 may act according to stored instructions located in a memory subsystem of the server 102, and may run an operating system, including any commercially available server operating system and/or any other operating systems discussed herein.

[0050] The content distribution network 100 may include one or more data store servers 104, such as database servers and file-based storage systems. The database servers 104 can access data that can be stored on a variety of hardware components. These hardware components can include, for example, components forming tier 0 storage, components forming tier 1 storage, components forming tier 2 storage, and/or any other tier of storage. In some embodiments, tier 0 storage refers to storage that is the fastest tier of storage in the database server 104, and particularly, the tier 0 storage is the fastest storage that is not RAM or cache memory. In some embodiments, the tier 0 memory can be embodied in solid state memory such as, for example, a solid-state drive (SSD) and/or flash memory.

[0051] In some embodiments, the tier 1 storage refers to storage that is one or several higher performing systems in the memory management system, and that is relatively slower than tier 0 memory, and relatively faster than other tiers of memory. The tier 1 memory can be one or several hard disks that can be, for example, high-performance hard disks. These hard disks can be one or both of physically or communicatively connected such as, for example, by one or several fiber channels. In some embodiments, the one or several disks can be arranged into a disk storage system, and specifically can be arranged into an enterprise class disk storage system. The disk storage system can include any desired level of redundancy to protect data stored therein, and in one embodiment, the disk storage system can be made

with grid architecture that creates parallelism for uniform allocation of system resources and balanced data distribution.

[0052] In some embodiments, the tier 2 storage refers to storage that includes one or several relatively lower performing systems in the memory management system, as compared to the tier 1 and tier 2 storages. Thus, tier 2 memory is relatively slower than tier 1 and tier 0 memories. Tier 2 memory can include one or several SATA-drives or one or several NL-SATA drives.

[0053] In some embodiments, the one or several hardware and/or software components of the database server 104 can be arranged into one or several storage area networks (SAN), which one or several storage area networks can be one or several dedicated networks that provide access to data storage, and particularly that provide access to consolidated, block level data storage. A SAN typically has its own network of storage devices that are generally not accessible through the local area network (LAN) by other devices. The SAN allows access to these devices in a manner such that these devices appear to be locally attached to the user device.

[0054] Data stores 104 may comprise stored data relevant to the functions of the content distribution network 100. Illustrative examples of data stores 104 that may be maintained in certain embodiments of the content distribution network 100 are described below in reference to FIG. 3. In some embodiments, multiple data stores may reside on a single server 104, either using the same storage components of server 104 or using different physical storage components to assure data security and integrity between data stores. In other embodiments, each data store may have a separate dedicated data store server 104.

[0055] Content distribution network 100 also may include one or more user devices 106 and/or supervisor devices 110. User devices 106 and supervisor devices 110 may display content received via the content distribution network 100, and may support various types of user interactions with the content. User devices 106 and supervisor devices 110 may include mobile devices such as smartphones, tablet computers, personal digital assistants, and wearable computing devices. Such mobile devices may run a variety of mobile operating systems, and may be enabled for Internet, e-mail, short message service (SMS), Bluetooth®, mobile radio-frequency identification (M-RFID), and/or other communication protocols. Other user devices 106 and supervisor devices 110 may be general purpose personal computers or special-purpose computing devices including, by way of example, personal computers, laptop computers, workstation computers, projection devices, and interactive room display systems. Additionally, user devices 106 and supervisor devices 110 may be any other electronic devices, such as thin-client computers, Internet-enabled gaming systems, business or home appliances, and/or personal messaging devices, capable of communicating over network(s) 120.

[0056] In different contexts of content distribution networks 100, user devices 106 and supervisor devices 110 may correspond to different types of specialized devices, for example, student devices and teacher devices in an educational network, employee devices and presentation devices in a company network, different gaming devices in a gaming network, etc. In some embodiments, user devices 106 and supervisor devices 110 may operate in the same physical location 107, such as a classroom or conference room. In

such cases, the devices may contain components that support direct communications with other nearby devices, such as a wireless transceivers and wireless communications interfaces, Ethernet sockets or other Local Area Network (LAN) interfaces, etc. In other implementations, the user devices **106** and supervisor devices **110** need not be used at the same location **107**, but may be used in remote geographic locations in which each user device **106** and supervisor device **110** may use security features and/or specialized hardware (e.g., hardware-accelerated SSL and HTTPS, WS-Security, firewalls, etc.) to communicate with the content management server **102** and/or other remotely located user devices **106**. Additionally, different user devices **106** and supervisor devices **110** may be assigned different designated roles, such as presenter devices, teacher devices, administrator devices, or the like, and in such cases the different devices may be provided with additional hardware and/or software components to provide content and support user capabilities not available to the other devices.

[0057] The content distribution network **100** also may include a privacy server **108** that maintains private user information at the privacy server **108** while using applications or services hosted on other servers. For example, the privacy server **108** may be used to maintain private data of a user within one jurisdiction even though the user is accessing an application hosted on a server (e.g., the content management server **102**) located outside the jurisdiction. In such cases, the privacy server **108** may intercept communications between a user device **106** or supervisor device **110** and other devices that include private user information. The privacy server **108** may create a token or identifier that does not disclose the private information and may use the token or identifier when communicating with the other servers and systems, instead of using the user's private information.

[0058] As illustrated in FIG. 1, the content management server **102** may be in communication with one or more additional servers, such as a content server **112**, a user data server **114**, and/or an administrator server **116**. Each of these servers may include some or all of the same physical and logical components as the content management server(s) **102**, and in some cases, the hardware and software components of these servers **112-116** may be incorporated into the content management server(s) **102**, rather than being implemented as separate computer servers.

[0059] Content server **112** may include hardware and software components to generate, store, and maintain the content resources for distribution to user devices **106** and other devices in the network **100**. For example, in content distribution networks **100** used for professional training and educational purposes, content server **112** may include data stores of training materials, presentations, plans, syllabi, reviews, evaluations, interactive programs and simulations, course models, course outlines, and various training interfaces that correspond to different materials and/or different types of user devices **106**. In content distribution networks **100** used for media distribution, interactive gaming, and the like, a content server **112** may include media content files such as music, movies, television programming, games, and advertisements.

[0060] User data server **114** may include hardware and software components that store and process data for multiple users relating to each user's activities and usage of the content distribution network **100**. For example, the content management server **102** may record and track each user's

system usage, including his or her user device **106**, content resources accessed, and interactions with other user devices **106**. This data may be stored and processed by the user data server **114**, to support user tracking and analysis features. For instance, in the professional training and educational contexts, the user data server **114** may store and analyze each user's training materials viewed, presentations attended, courses completed, interactions, evaluation results, and the like. The user data server **114** may also include a repository for user-generated material, such as evaluations and tests completed by users, and documents and assignments prepared by users. In the context of media distribution and interactive gaming, the user data server **114** may store and process resource access data for multiple users (e.g., content titles accessed, access times, data usage amounts, gaming histories, user devices and device types, etc.).

[0061] Administrator server **116** may include hardware and software components to initiate various administrative functions at the content management server **102** and other components within the content distribution network **100**. For example, the administrator server **116** may monitor device status and performance for the various servers, data stores, and/or user devices **106** in the content distribution network **100**. When necessary, the administrator server **116** may add or remove devices from the network **100**, and perform device maintenance such as providing software updates to the devices in the network **100**. Various administrative tools on the administrator server **116** may allow authorized users to set user access permissions to various content resources, monitor resource usage by users and devices **106**, and perform analyses and generate reports on specific network users and/or devices (e.g., resource usage tracking reports, training evaluations, etc.).

[0062] The content distribution network **100** may include one or more communication networks **120**. Although only a single network **120** is identified in FIG. 1, the content distribution network **100** may include any number of different communication networks between any of the computer servers and devices shown in FIG. 1 and/or other devices described herein. Communication networks **120** may enable communication between the various computing devices, servers, and other components of the content distribution network **100**. As discussed below, various implementations of content distribution networks **100** may employ different types of networks **120**, for example, computer networks, telecommunications networks, wireless networks, and/or any combination of these and/or other networks.

[0063] The content distribution network **100** may include one or several navigation systems or features including, for example, the Global Positioning System ("GPS"), GALILEO, or the like, or location systems or features including, for example, one or several transceivers that can determine location of the one or several components of the content distribution network **100** via, for example, triangulation. All of these are depicted as navigation system **122**.

[0064] In some embodiments, navigation system **122** can include one or several features that can communicate with one or several components of the content distribution network **100** including, for example, with one or several of the user devices **106** and/or with one or several of the supervisor devices **110**. In some embodiments, this communication can include the transmission of a signal from the navigation system **122** which signal is received by one or several

components of the content distribution network **100** and can be used to determine the location of the one or several components of the content distribution network **100**.

[0065] The content distribution network **100** may include one or several artificial intelligence agents **124**, also referred to herein as a AI agent **124**, intelligence agent **123**, and IA agent **124**. The AI agent **124** can be an autonomous entity that can receive inputs via one or several sensors or from one or several user devices and can provide responses to those received inputs. In some embodiments, the AI agent **124** can comprise a question answering computing system, all or portions of an intelligent tutoring system, pedagogical agent, or the like. In some embodiments, the functioning of the AI agent **124** can be predicated upon artificial intelligence, machine learning, or the like. The AI agent **124** can reside in the server **102** and/or another component of the content distribution network **100**, or can reside on a separate server or on separate computing resources. The AI agent **124** can be configured to receive inputs from the user device **106**. The AI agent **124** can further launch a dialogue based on the received inputs and identify a remediation dialogue based on the received user inputs. In some embodiments, the AI agent **124** can determine termination of the dialogue and/or the remediation dialogue based on a user performance metric indicative of a user skill level. In some embodiments, the dialogue can be determined for termination when the user performance metric identifies a sufficient user skill level.

[0066] With reference to FIG. 2, an illustrative distributed computing environment **200** is shown including a computer server **202**, four client computing devices **206**, and other components that may implement certain embodiments and features described herein. In some embodiments, the server **202** may correspond to the content management server **102** discussed above in FIG. 1, and the client computing devices **206** may correspond to the user devices **106**. However, the computing environment **200** illustrated in FIG. 2 may correspond to any other combination of devices and servers configured to implement a client-server model or other distributed computing architecture.

[0067] Client devices **206** may be configured to receive and execute client applications over one or more networks **220**. Such client applications may be web browser-based applications and/or standalone software applications, such as mobile device applications. Server **202** may be communicatively coupled with the client devices **206** via one or more communication networks **220**. Client devices **206** may receive client applications from server **202** or from other application providers (e.g., public or private application stores). Server **202** may be configured to run one or more server software applications or services, for example, web-based or cloud-based services, to support content distribution and interaction with client devices **206**. Users operating client devices **206** may in turn utilize one or more client applications (e.g., virtual client applications) to interact with server **202** to utilize the services provided by these components.

[0068] Various different subsystems and/or components **204** may be implemented on server **202**. Users operating the client devices **206** may initiate one or more client applications to use services provided by these subsystems and components. The subsystems and components within the server **202** and client devices **206** may be implemented in hardware, firmware, software, or combinations thereof. Various different system configurations are possible in different

distributed computing systems **200** and content distribution networks **100**. The embodiment shown in FIG. 2 is thus one example of a distributed computing system and is not intended to be limiting.

[0069] Although exemplary computing environment **200** is shown with four client computing devices **206**, any number of client computing devices may be supported. Other devices, such as specialized sensor devices, etc., may interact with client devices **206** and/or server **202**.

[0070] As shown in FIG. 2, various security and integration components **208** may be used to send and manage communications between the server **202** and user devices **206** over one or more communication networks **220**. The security and integration components **208** may include separate servers, such as web servers and/or authentication servers, and/or specialized networking components, such as firewalls, routers, gateways, load balancers, and the like. In some cases, the security and integration components **208** may correspond to a set of dedicated hardware and/or software operating at the same physical location and under the control of same entities as server **202**. For example, components **208** may include one or more dedicated web servers and network hardware in a datacenter or a cloud infrastructure. In other examples, the security and integration components **208** may correspond to separate hardware and software components which may be operated at a separate physical location and/or by a separate entity.

[0071] Security and integration components **208** may implement various security features for data transmission and storage, such as authenticating users and restricting access to unknown or unauthorized users. In various implementations, security and integration components **208** may provide, for example, a file-based integration scheme or a service-based integration scheme for transmitting data between the various devices in the content distribution network **100**. Security and integration components **208** also may use secure data transmission protocols and/or encryption for data transfers, for example, File Transfer Protocol (FTP), Secure File Transfer Protocol (SFTP), and/or Pretty Good Privacy (PGP) encryption.

[0072] In some embodiments, one or more web services may be implemented within the security and integration components **208** and/or elsewhere within the content distribution network **100**. Such web services, including cross-domain and/or cross-platform web services, may be developed for enterprise use in accordance with various web service standards, such as RESTful web services (i.e., services based on the Representation State Transfer (REST) architectural style and constraints), and/or web services designed in accordance with the Web Service Interoperability (WS-I) guidelines. Some web services may use the Secure Sockets Layer (SSL) or Transport Layer Security (TLS) protocol to provide secure connections between the server **202** and user devices **206**. SSL or TLS may use HTTP or HTTPS to provide authentication and confidentiality. In other examples, web services may be implemented using REST over HTTPS with the OAuth open standard for authentication, or using the WS-Security standard which provides for secure SOAP messages using XML encryption. In other examples, the security and integration components **208** may include specialized hardware for providing secure web services. For example, security and integration components **208** may include secure network appliances having built-in features such as hardware-accelerated SSL and

HTTPS, WS-Security, and firewalls. Such specialized hardware may be installed and configured in front of any web servers, so that any external devices may communicate directly with the specialized hardware.

[0073] Communication network(s) **220** may be any type of network familiar to those skilled in the art that can support data communications using any of a variety of commercially-available protocols, including without limitation, TCP/IP (transmission control protocol/Internet protocol), SNA (systems network architecture), IPX (Internet packet exchange), Secure Sockets Layer (SSL) or Transport Layer Security (TLS) protocols, Hyper Text Transfer Protocol (HTTP) and Secure Hyper Text Transfer Protocol (HTTPS), Bluetooth®, Near Field Communication (NFC), and the like. Merely by way of example, network(s) **220** may be local area networks (LAN), such as one based on Ethernet, Token-Ring and/or the like. Network(s) **220** also may be wide-area networks, such as the Internet. Networks **220** may include telecommunication networks such as public switched telephone networks (PSTNs), or virtual networks such as an intranet or an extranet. Infrared and wireless networks (e.g., using the Institute of Electrical and Electronics (IEEE) 802.11 protocol suite or other wireless protocols) also may be included in networks **220**.

[0074] Computing environment **200** also may include one or more data stores **210** and/or back-end servers **212**. In certain examples, the data stores **210** may correspond to data store server(s) **104** discussed above in FIG. 1, and back-end servers **212** may correspond to the various back-end servers **112-116**. Data stores **210** and servers **212** may reside in the same datacenter or may operate at a remote location from server **202**. In some cases, one or more data stores **210** may reside on a non-transitory storage medium within the server **202**. Other data stores **210** and back-end servers **212** may be remote from server **202** and configured to communicate with server **202** via one or more networks **220**. In certain embodiments, data stores **210** and back-end servers **212** may reside in a storage-area network (SAN), or may use storage-as-a-service (STaaS) architectural model.

[0075] With reference to FIG. 3, an illustrative set of data stores and/or data store servers is shown, corresponding to the data store servers **104** of the content distribution network **100** discussed above in FIG. 1. One or more individual data stores **301-311** may reside in storage on a single computer server **104** (or a single server farm or cluster) under the control of a single entity, or may reside on separate servers operated by different entities and/or at remote locations. In some embodiments, data stores **301-311** may be accessed by the content management server **102** and/or other devices and servers within the network **100** (e.g., user devices **106**, supervisor devices **110**, administrator servers **116**, etc.). Access to one or more of the data stores **301-311** may be limited or denied based on the processes, user credentials, and/or devices attempting to interact with the data store.

[0076] The paragraphs below describe examples of specific data stores that may be implemented within some embodiments of a content distribution network **100**. It should be understood that the below descriptions of data stores **301-311**, including their functionality and types of data stored therein, are illustrative and non-limiting. Data stores server architecture, design, and the execution of specific data stores **301-311** may depend on the context, size, and functional requirements of a content distribution network **100**. For example, in content distribution systems **100**

used for professional training and educational purposes, separate databases or file-based storage systems may be implemented in data store server(s) **104** to store trainee and/or student data, trainer and/or professor data, training module data and content descriptions, training results, evaluation data, and the like. In contrast, in content distribution systems **100** used for media distribution from content providers to subscribers, separate data stores may be implemented in data stores server(s) **104** to store listings of available content titles and descriptions, content title usage statistics, subscriber profiles, account data, payment data, network usage statistics, etc.

[0077] A user profile data store **301**, also referred to herein as a user profile database **301**, may include information relating to the end users within the content distribution network **100**. This information may include user characteristics such as the user names, access credentials (e.g., logins and passwords), user preferences, and information relating to any previous user interactions within the content distribution network **100** (e.g., requested content, posted content, content modules completed, training scores or evaluations, other associated users, etc.). In some embodiments, this information can relate to one or several individual end users such as, for example, one or several students, teachers, administrators, or the like, and in some embodiments, this information can relate to one or several institutional end users such as, for example, one or several schools, groups of schools such as one or several school districts, one or several colleges, one or several universities, one or several training providers, or the like. In some embodiments, this information can identify one or several user memberships in one or several groups such as, for example, a student's membership in a university, school, program, grade, course, class, or the like.

[0078] The user profile database **301** can include information relating to a user's status, location, or the like. This information can identify, for example, a device a user is using, the location of that device, or the like. In some embodiments, this information can be generated based on any location detection technology including, for example, a navigation system **122**, or the like.

[0079] Information relating to the user's status can identify, for example, logged-in status information that can indicate whether the user is presently logged-in to the content distribution network **100** and/or whether the log-in is active. In some embodiments, the information relating to the user's status can identify whether the user is currently accessing content and/or participating in an activity from the content distribution network **100**.

[0080] In some embodiments, information relating to the user's status can identify, for example, one or several attributes of the user's interaction with the content distribution network **100**, and/or content distributed by the content distribution network **100**. This can include data identifying the user's interactions with the content distribution network **100**, the content consumed by the user through the content distribution network **100**, or the like. In some embodiments, this can include data identifying the type of information accessed through the content distribution network **100** and/or the type of activity performed by the user via the content distribution network **100**, the lapsed time since the last time the user accessed content and/or participated in an activity from the content distribution network **100**, or the like. In some embodiments, this information can relate to a content

program comprising an aggregate of data, content, and/or activities, and can identify, for example, progress through the content program, or through the aggregate of data, content, and/or activities forming the content program. In some embodiments, this information can track, for example, the amount of time since participation in and/or completion of one or several types of activities, the amount of time since communication with one or several supervisors and/or supervisor devices **110**, or the like.

[0081] In some embodiments in which the one or several end users are individuals, and specifically are students, the user profile database **301** can further include information relating to these students' academic and/or educational history. This information can identify one or several courses of study that the student has initiated, completed, and/or partially completed, as well as grades received in those courses of study. In some embodiments, the student's academic and/or educational history can further include information identifying student performance on one or several tests, quizzes, and/or assignments. In some embodiments, this information can be stored in a tier of memory that is not the fastest memory in the content delivery network **100**.

[0082] The user profile database **301** can include information relating to one or several student learning preferences. In some embodiments, for example, the user, also referred to herein as the student or the student-user may have one or several preferred learning styles, one or several most effective learning styles, and/or the like. In some embodiments, the student's learning style can be any learning style describing how the student best learns or how the student prefers to learn. In one embodiment, these learning styles can include, for example, identification of the student as an auditory learner, as a visual learner, and/or as a tactile learner. In some embodiments, the data identifying one or several student learning styles can include data identifying a learning style based on the student's educational history such as, for example, identifying a student as an auditory learner when the student has received significantly higher grades and/or scores on assignments and/or in courses favorable to auditory learners. In some embodiments, this information can be stored in a tier of memory that is not the fastest memory in the content delivery network **100**.

[0083] In some embodiments, the user profile data store **301** can further include information identifying one or several user skill levels. In some embodiments, these one or several user skill levels can identify a skill level determined based on past performance by the user interacting with the content delivery network **100**, and in some embodiments, these one or several user skill levels can identify a predicted skill level determined based on past performance by the user interacting with the content delivery network **100** and one or several predictive models.

[0084] The user profile database **301** can further include information relating to one or several teachers and/or instructors who are responsible for organizing, presenting, and/or managing the presentation of information to the student. In some embodiments, user profile database **301** can include information identifying courses and/or subjects that have been taught by the teacher, data identifying courses and/or subjects currently taught by the teacher, and/or data identifying courses and/or subjects that will be taught by the teacher. In some embodiments, this can include information relating to one or several teaching styles of one or several teachers. In some embodiments, the user profile database

301 can further include information indicating past evaluations and/or evaluation reports received by the teacher. In some embodiments, the user profile database **301** can further include information relating to improvement suggestions received by the teacher, training received by the teacher, continuing education received by the teacher, and/or the like. In some embodiments, this information can be stored in a tier of memory that is not the fastest memory in the content delivery network **100**.

[0085] An accounts data store **302**, also referred to herein as an accounts database **302**, may generate and store account data for different users in various roles within the content distribution network **100**. For example, accounts may be created in an accounts data store **302** for individual end users, supervisors, administrator users, and entities such as companies or educational institutions. Account data may include account types, current account status, account characteristics, and any parameters, limits, restrictions associated with the accounts.

[0086] A content library data store **303**, also referred to herein as a content library database **303**, may include information describing the individual content items (or content resources or data packets) available via the content distribution network **100**. In some embodiments, these data packets in the content library database **303** can be linked to form an object network. In some embodiments, these data packets can be linked in the object network according to one or several prerequisite relationships that can, for example, identify the relative hierarchy and/or difficulty of the data objects. In some embodiments, this hierarchy of data objects can be generated by the content distribution network **100** according to user experience with the object network, and in some embodiments, this hierarchy of data objects can be generated based on one or several existing and/or external hierarchies such as, for example, a syllabus, a table of contents, or the like. In some embodiments, for example, the object network can correspond to a syllabus such that content for the syllabus is embodied in the object network.

[0087] In some embodiments, the content library data store **303** can comprise a syllabus, a schedule, or the like. In some embodiments, the syllabus or schedule can identify one or several tasks and/or events relevant to the user. In some embodiments, for example, when the user is a member of a group such as a section or a class, these tasks and/or events relevant to the user can identify one or several assignments, quizzes, exams, or the like.

[0088] In some embodiments, the library data store **303** may include metadata, properties, and other characteristics associated with the content resources stored in the content server **112**. Such data may identify one or more aspects or content attributes of the associated content resources, for example, subject matter, access level, or skill level of the content resources, license attributes of the content resources (e.g., any limitations and/or restrictions on the licensable use and/or distribution of the content resource), price attributes of the content resources (e.g., a price and/or price structure for determining a payment amount for use or distribution of the content resource), rating attributes for the content resources (e.g., data indicating the evaluation or effectiveness of the content resource), and the like. In some embodiments, the library data store **303** may be configured to allow updating of content metadata or properties, and to allow the addition and/or removal of information relating to the content resources. For example, content relationships may be

implemented as graph structures, which may be stored in the library data store **303** or in an additional store for use by selection algorithms along with the other metadata.

[0089] In some embodiments, the content library data store **303** can contain information used in evaluating responses received from users. In some embodiments, for example, a user can receive content from the content distribution network **100** and can, subsequent to receiving that content, provide a response to the received content. In some embodiments, for example, the received content can comprise one or several questions, prompts, or the like, and the response to the received content can comprise an answer to those one or several questions, prompts, or the like. In some embodiments, information, referred to herein as “comparative data,” from the content library data store **303** can be used to determine whether the responses are the correct and/or desired responses.

[0090] In some embodiments, the content library database **303** and/or the user profile database **301** can comprise an aggregation network, also referred to herein as a content network or content aggregation network. The aggregation network can comprise a plurality of content aggregations that can be linked together by, for example: creation by common user; relation to a common subject, topic, skill, or the like; creation from a common set of source material such as source data packets; or the like. In some embodiments, the content aggregation can comprise a grouping of content comprising the presentation portion that can be provided to the user in the form of, for example, a flash card and an extraction portion that can comprise the desired response to the presentation portion such as for example, an answer to a flash card. In some embodiments, one or several content aggregations can be generated by the content distribution network **100** and can be related to one or several data packets that can be, for example, organized in object network. In some embodiments, the one or several content aggregations can be each created from content stored in one or several of the data packets.

[0091] In some embodiments, the content aggregations located in the content library database **303** and/or the user profile database **301** can be associated with a user-creator of those content aggregations. In some embodiments, access to content aggregations can vary based on, for example, whether a user created the content aggregations. In some embodiments, the content library database **303** and/or the user profile database **301** can comprise a database of content aggregations associated with a specific user, and in some embodiments, the content library database **303** and/or the user profile database **301** can comprise a plurality of databases of content aggregations that are each associated with a specific user. In some embodiments, these databases of content aggregations can include content aggregations created by their specific user and, in some embodiments, these databases of content aggregations can further include content aggregations selected for inclusion by their specific user and/or a supervisor of that specific user. In some embodiments, these content aggregations can be arranged and/or linked in a hierarchical relationship similar to the data packets in the object network and/or linked to the object network in the object network or the tasks or skills associated with the data packets in the object network or the syllabus or schedule.

[0092] In some embodiments, the content aggregation network, and the content aggregations forming the content

aggregation network can be organized according to the object network and/or the hierarchical relationships embodied in the object network. In some embodiments, the content aggregation network, and/or the content aggregations forming the content aggregation network can be organized according to one or several tasks identified in the syllabus, schedule or the like.

[0093] A pricing data store **304** may include pricing information and/or pricing structures for determining payment amounts for providing access to the content distribution network **100** and/or the individual content resources within the network **100**. In some cases, pricing may be determined based on a user's access to the content distribution network **100**, for example, a time-based subscription fee, or pricing based on network usage. In other cases, pricing may be tied to specific content resources. Certain content resources may have associated pricing information, whereas other pricing determinations may be based on the resources accessed, the profiles and/or accounts of the user, and the desired level of access (e.g., duration of access, network speed, etc.). Additionally, the pricing data store **304** may include information relating to compilation pricing for groups of content resources, such as group prices and/or price structures for groupings of resources.

[0094] A license data store **305** may include information relating to licenses and/or licensing of the content resources within the content distribution network **100**. For example, the license data store **305** may identify licenses and licensing terms for individual content resources and/or compilations of content resources in the content server **112**, the rights holders for the content resources, and/or common or large-scale right holder information such as contact information for rights holders of content not included in the content server **112**.

[0095] A content access data store **306** may include access rights and security information for the content distribution network **100** and specific content resources. For example, the content access data store **306** may include login information (e.g., user identifiers, logins, passwords, etc.) that can be verified during user login attempts to the network **100**. The content access data store **306** also may be used to store assigned user roles and/or user levels of access. For example, a user's access level may correspond to the sets of content resources and/or the client or server applications that the user is permitted to access. Certain users may be permitted or denied access to certain applications and resources based on their subscription level, training program, course/grade level, etc. Certain users may have supervisory access over one or more end users, allowing the supervisor to access all or portions of the end user's content, activities, evaluations, etc. Additionally, certain users may have administrative access over some users and/or some applications in the content management network **100**, allowing such users to add and remove user accounts, modify user access permissions, perform maintenance updates on software and servers, etc.

[0096] A source data store **307** may include information relating to the source of the content resources available via the content distribution network. For example, a source data store **307** may identify the authors and originating devices of content resources, previous pieces of data and/or groups of data originating from the same authors or originating devices, and the like.

[0097] An evaluation data store **308** may include information used to direct the evaluation of users and content resources in the content management network **100**. In some embodiments, the evaluation data store **308** may contain, for example, the analysis criteria and the analysis guidelines for evaluating users (e.g., trainees/students, gaming users, media content consumers, etc.) and/or for evaluating the content resources in the network **100**. The evaluation data store **308** also may include information relating to evaluation processing tasks, for example, the identification of users and user devices **106** that have received certain content resources or accessed certain applications, the status of evaluations or evaluation histories for content resources, users, or applications, and the like. Evaluation criteria may be stored in the evaluation data store **308** including data and/or instructions in the form of one or several electronic rubrics or scoring guides for use in the evaluation of the content, users, or applications. The evaluation data store **308** also may include past evaluations and/or evaluation analyses for users, content, and applications, including relative rankings, characterizations, explanations, and the like.

[0098] A model data store **309**, also referred to herein as a model database **309**, can store information relating to one or several predictive models. In some embodiments, these can include one or several evidence models, risk models, skill models, or the like. In some embodiments, an evidence model can be a mathematically-based statistical model. The evidence model can be based on, for example, Item Response Theory (IRT), Bayesian Network (Bayes net), Performance Factor Analysis (PFA), or the like. The evidence model can, in some embodiments, be customizable to a user and/or to one or several content items. Specifically, one or several inputs relating to the user and/or to one or several content items can be inserted into the evidence model. These inputs can include, for example, one or several measures of user skill level, one or several measures of content item difficulty and/or skill level, or the like. The customized evidence model can then be used to predict the likelihood of the user providing desired or undesired responses to one or several of the content items.

[0099] In some embodiments, the risk models can include one or several models that can be used to calculate one or several model function values. In some embodiments, these one or several model function values can be used to calculate a risk probability, which risk probability can characterize the risk of a user such as a student-user failing to achieve a desired outcome such as, for example, failing to correctly respond to one or several data packets, failure to achieve a desired level of completion of a program, for example in a pre-defined time period, failure to achieve a desired learning outcome, or the like. In some embodiments, the risk probability can identify the risk of the student-user failing to complete 60% of the program.

[0100] In some embodiments, these models can include a plurality of model functions including, for example, a first model function, a second model function, a third model function, and a fourth model function. In some embodiments, some or all of the model functions can be associated with a portion of the program such as, for example, a completion stage and/or completion status of the program. In one embodiment, for example, the first model function can be associated with a first completion status, the second model function can be associated with a second completion status, the third model function can be associated with a

third completion status, and the fourth model function can be associated with a fourth completion status. In some embodiments, these completion statuses can be selected such that some or all of these completion statuses are less than the desired level of completion of the program. Specifically, in some embodiments, these completion statuses can be selected to all be at less than 60% completion of the program, and more specifically, in some embodiments, the first completion status can be at 20% completion of the program, the second completion status can be at 30% completion of the program, the third completion status can be at 40% completion of the program, and the fourth completion status can be at 50% completion of the program. Similarly, any desired number of model functions can be associated with any desired number of completion statuses.

[0101] In some embodiments, a model function can be selected from the plurality of model functions based on a student-user's progress through a program. In some embodiments, the student-user's progress can be compared to one or several status trigger thresholds, each of which status trigger thresholds can be associated with one or more of the model functions. If one of the status triggers is triggered by the student-user's progress, the corresponding one or several model functions can be selected.

[0102] The model functions can comprise a variety of types of models and/or functions. In some embodiments, each of the model functions outputs a function value that can be used in calculating a risk probability. This function value can be calculated by performing one or several mathematical operations on one or several values indicative of one or several user attributes and/or user parameters, also referred to herein as program status parameters. In some embodiments, each of the model functions can use the same program status parameters, and in some embodiments, the model functions can use different program status parameters. In some embodiments, the model functions use different program status parameters when at least one of the model functions uses at least one program status parameter that is not used by others of the model functions.

[0103] In some embodiments, a skill model can comprise a statistical model identifying a predictive skill level of one or several students. In some embodiments, this model can identify a single skill level of a student and/or a range of possible skill levels of a student. In some embodiments, this statistical model can identify a skill level of a student-user and an error value or error range associated with that skill level. In some embodiments, the error value can be associated with a confidence interval determined based on a confidence level. Thus, in some embodiments, as the number of student interactions with the content distribution network increases, the confidence level can increase and the error value can decrease such that the range identified by the error value about the predicted skill level is smaller.

[0104] A threshold database **310**, also referred to herein as a threshold database, can store one or several threshold values. These one or several threshold values can delineate between states or conditions. In one exemplary embodiment, for example, a threshold value can delineate between an acceptable user performance and an unacceptable user performance, between content appropriate for a user and content that is inappropriate for a user, between risk levels, or the like.

[0105] In addition to the illustrative data stores described above, data store server(s) **104** (e.g., database servers, file-

based storage servers, etc.) may include one or more external data aggregators **311**. External data aggregators **311** may include third-party data sources accessible to the content management network **100**, but not maintained by the content management network **100**. External data aggregators **311** may include any electronic information source relating to the users, content resources, or applications of the content distribution network **100**. For example, external data aggregators **311** may be third-party data stores containing demographic data, education-related data, consumer sales data, health-related data, and the like. Illustrative external data aggregators **311** may include, for example, social networking web servers, public records data stores, learning management systems, educational institution servers, business servers, consumer sales data stores, medical record data stores, etc. Data retrieved from various external data aggregators **311** may be used to verify and update user account information, suggest user content, and perform user and content evaluations.

[**0106**] With reference now to FIG. 4, a block diagram is shown illustrating an embodiment of one or more content management servers **102** within a content distribution network **100**. In such an embodiment, content management server **102** performs internal data gathering and processing of streamed content along with external data gathering and processing. Other embodiments could have either all external or all internal data gathering. This embodiment allows reporting timely information that might be of interest to the reporting party or other parties. In this embodiment, the content management server **102** can monitor gathered information from several sources to allow it to make timely business and/or processing decisions based upon that information. For example, reports of user actions and/or responses, as well as the status and/or results of one or several processing tasks could be gathered and reported to the content management server **102** from a number of sources.

[**0107**] Internally, the content management server **102** gathers information from one or more internal components **402-408**. The internal components **402-408** gather and/or process information relating to such things as: content provided to users; content consumed by users; responses provided by users; user skill levels; content difficulty levels; next content for providing to users; etc. The internal components **402-408** can report the gathered and/or generated information in real-time, near real-time or along another time line. To account for any delay in reporting information, a time stamp or staleness indicator can inform others of how timely the information was sampled. The content management server **102** can opt to allow third parties to use internally or externally gathered information that is aggregated within the server **102** by subscription to the content distribution network **100**.

[**0108**] A command and control (CC) interface **338** configures the gathered input information to an output of data streams, also referred to herein as content streams. APIs for accepting gathered information and providing data streams are provided to third parties external to the server **102** who want to subscribe to data streams. The server **102** or a third party can design as yet undefined APIs using the CC interface **338**. The server **102** can also define authorization and authentication parameters using the CC interface **338** such as authentication, authorization, login, and/or data encryption. CC information is passed to the internal com-

ponents **402-408** and/or other components of the content distribution network **100** through a channel separate from the gathered information or data stream in this embodiment, but other embodiments could embed CC information in these communication channels. The CC information allows throttling information reporting frequency, specifying formats for information and data streams, deactivation of one or several internal components **402-408** and/or other components of the content distribution network **100**, updating authentication and authorization, etc.

[**0109**] The various data streams that are available can be researched and explored through the CC interface **338**. Those data stream selections for a particular subscriber, which can be one or several of the internal components **402-408** and/or other components of the content distribution network **100**, are stored in the queue subscription information database **322**. The server **102** and/or the CC interface **338** then routes selected data streams to processing subscribers that have selected delivery of a given data stream. Additionally, the server **102** also supports historical queries of the various data streams that are stored in an historical data store **334** as gathered by an archive data agent **336**. Through the CC interface **238** various data streams can be selected for archiving into the historical data store **334**.

[**0110**] Components of the content distribution network **100** outside of the server **102** can also gather information that is reported to the server **102** in real-time, near real-time or along another time line. There is a defined API between those components and the server **102**. Each type of information or variable collected by server **102** falls within a defined API or multiple APIs. In some cases, the CC interface **338** is used to define additional variables to modify an API that might be of use to processing subscribers. The additional variables can be passed to all processing subscribers or just a subset. For example, a component of the content distribution network **100** outside of the server **102** may report a user response but define an identifier of that user as a private variable that would not be passed to processing subscribers lacking access to that user and/or authorization to receive that user data. Processing subscribers having access to that user and/or authorization to receive that user data would receive the subscriber identifier along with response reported that component. Encryption and/or unique addressing of data streams or sub-streams can be used to hide the private variables within the messaging queues.

[**0111**] The user devices **106** and/or supervisor devices **110** communicate with the server **102** through security and/or integration hardware **410**. The communication with security and/or integration hardware **410** can be encrypted or not. For example, a socket using a TCP connection could be used. In addition to TCP, other transport layer protocols like SCTP and UDP could be used in some embodiments to intake the gathered information. A protocol such as SSL could be used to protect the information over the TCP connection. Authentication and authorization can be performed to any user devices **106** and/or supervisor device interfacing to the server **102**. The security and/or integration hardware **410** receives the information from one or several of the user devices **106** and/or the supervisor devices **110** by providing the API and any encryption, authorization, and/or authentication. In some cases, the security and/or integration hardware **410** reformats or rearranges this received information.

[0112] The messaging bus 412, also referred to herein as a messaging queue or a messaging channel, can receive information from the internal components of the server 102 and/or components of the content distribution network 100 outside of the server 102 and distribute the gathered information as a data stream to any processing subscribers that have requested the data stream from the messaging queue 412. Specifically, in some embodiments, the messaging bus 412 can receive and output information from at least one of the packet selection system, the presentation system, the response system, and the summary model system. In some embodiments, this information can be output according to a “push” model, and in some embodiments, this information can be output according to a “pull” model.

[0113] As indicated in FIG. 4, processing subscribers are indicated by a connector to the messaging bus 412, the connector having an arrow head pointing away from the messaging bus 412. Only data streams within the messaging queue 412 that a particular processing subscriber has subscribed to may be read by that processing subscriber if received at all. Gathered information sent to the messaging queue 412 is processed and returned in a data stream in a fraction of a second by the messaging queue 412. Various multicasting and routing techniques can be used to distribute a data stream from the messaging queue 412 that a number of processing subscribers have requested. Protocols such as Multicast or multiple Unicast could be used to distribute streams within the messaging queue 412. Additionally, transport layer protocols like TCP, SCTP and UDP could be used in various embodiments.

[0114] Through the CC interface 338, an external or internal processing subscriber can be assigned one or more data streams within the messaging queue 412. A data stream is a particular type of message in a particular category. For example, a data stream can comprise all of the data reported to the messaging bus 412 by a designated set of components. One or more processing subscribers could subscribe and receive the data stream to process the information and make a decision and/or feed the output from the processing as gathered information fed back into the messaging queue 412. Through the CC interface 338, a developer can search the available data streams or specify a new data stream and its API. The new data stream might be determined by processing a number of existing data streams with a processing subscriber.

[0115] The CDN 110 has internal processing subscribers 402-408 that process assigned data streams to perform functions within the server 102. Internal processing subscribers 402-408 could perform functions such as providing content to a user, receiving a response from a user, determining the correctness of the received response, updating one or several models based on the correctness of the response, recommending new content for providing to one or several users, or the like. The internal processing subscribers 402-408 can decide filtering and weighting of records from the data stream. To the extent that decisions are made based upon analysis of the data stream, each data record is time stamped to reflect when the information was gathered such that additional credibility could be given to more recent results, for example. Other embodiments may filter out records in the data stream that are from an unreliable source or stale. For example, a particular con-

tributor of information may prove to have less than optimal gathered information and that could be weighted very low or removed altogether.

[0116] Internal processing subscribers 402-408 may additionally process one or more data streams to provide different information to feed back into the messaging queue 412 to be part of a different data stream. For example, hundreds of user devices 106 could provide responses that are put into a data stream on the messaging queue 412. An internal processing subscriber 402-408 could receive the data stream and process it to determine the difficulty of one or several data packets provided to one or several users, and supply this information back onto the messaging queue 412 for possible use by other internal and external processing subscribers.

[0117] As mentioned above, the CC interface 338 allows the CDN 110 to query historical messaging queue 412 information. An archive data agent 336 listens to the messaging queue 412 to store data streams in a historical database 334. The historical database 334 may store data streams for varying amounts of time and may not store all data streams. Different data streams may be stored for different amounts of time.

[0118] With regard to the components 402-48, the content management server(s) 102 may include various server hardware and software components that manage the content resources within the content distribution network 100 and provide interactive and adaptive content to users on various user devices 106. For example, content management server (s) 102 may provide instructions to and receive information from the other devices within the content distribution network 100, in order to manage and transmit content resources, user data, and server or client applications executing within the network 100.

[0119] A content management server 102 may include a packet selection system 402. The packet selection system 402 may be implemented using dedicated hardware within the content distribution network 100 (e.g., a packet selection server 402), or using designated hardware and software resources within a shared content management server 102. In some embodiments, the packet selection system 402 may adjust the selection and adaptive capabilities of content resources to match the needs and desires of the users receiving the content. For example, the packet selection system 402 may query various data stores and servers 104 to retrieve user information, such as user preferences and characteristics (e.g., from a user profile data store 301), user access restrictions to content resources (e.g., from a content access data store 306), previous user results and content evaluations (e.g., from an evaluation data store 308), and the like. Based on the retrieved information from data stores 104 and other data sources, the packet selection system 402 may modify content resources for individual users.

[0120] In some embodiments, the packet selection system 402 can include a recommendation engine, also referred to herein as an adaptive recommendation engine. In some embodiments, the recommendation engine can select one or several pieces of content, also referred to herein as data packets, for providing to a user. These data packets can be selected based on, for example, the information retrieved from the database server 104 including, for example, the user profile database 301, the content library database 303, the model database 309, or the like. In some embodiments, these one or several data packets can be adaptively selected and/or selected according to one or several selection rules.

In one embodiment, for example, the recommendation engine can retrieve information from the user profile database 301 identifying, for example, a skill level of the user. The recommendation engine can further retrieve information from the content library database 303 identifying, for example, potential data packets for providing to the user and the difficulty of those data packets and/or the skill level associated with those data packets.

[0121] The recommendation engine can identify one or several potential data packets for providing and/or one or several data packets for providing to the user based on, for example, one or several rules, models, predictions, or the like. The recommendation engine can use the skill level of the user to generate a prediction of the likelihood of one or several users providing a desired response to some or all of the potential data packets. In some embodiments, the recommendation engine can pair one or several data packets with selection criteria that may be used to determine which packet should be delivered to a student-user based on one or several received responses from that student-user. In some embodiments, one or several data packets can be eliminated from the pool of potential data packets if the prediction indicates either too high a likelihood of a desired response or too low a likelihood of a desired response. In some embodiments, the recommendation engine can then apply one or several selection criteria to the remaining potential data packets to select a data packet for providing to the user. These one or several selection criteria can be based on, for example, criteria relating to a desired estimated time for receipt of response to the data packet, one or several content parameters, one or several assignment parameters, or the like.

[0122] A content management server 102 also may include a summary model system 404. The summary model system 404 may be implemented using dedicated hardware within the content distribution network 100 (e.g., a summary model server 404), or using designated hardware and software resources within a shared content management server 102. In some embodiments, the summary model system 404 may monitor the progress of users through various types of content resources and groups, such as media compilations, courses or curriculums in training or educational contexts, interactive gaming environments, and the like. For example, the summary model system 404 may query one or more databases and/or data store servers 104 to retrieve user data such as associated content compilations or programs, content completion status, user goals, results, and the like.

[0123] A content management server 102 also may include a response system 406, which can include, in some embodiments, a response processor. The response system 406 may be implemented using dedicated hardware within the content distribution network 100 (e.g., a response server 406), or using designated hardware and software resources within a shared content management server 102. The response system 406 may be configured to receive and analyze information from user devices 106. For example, various ratings of content resources submitted by users may be compiled and analyzed, and then stored in a data store (e.g., a content library data store 303 and/or evaluation data store 308) associated with the content. In some embodiments, the response server 406 may analyze the information to determine the effectiveness or appropriateness of content resources with, for example, a subject matter, an age group, a skill level, or the like. In some embodiments, the response

system 406 may provide updates to the packet selection system 402 or the summary model system 404, with the attributes of one or more content resources or groups of resources within the network 100. The response system 406 also may receive and analyze user evaluation data from user devices 106, supervisor devices 110, and administrator servers 116, etc. For instance, response system 406 may receive, aggregate, and analyze user evaluation data for different types of users (e.g., end users, supervisors, administrators, etc.) in different contexts (e.g., media consumer ratings, trainee or student comprehension levels, teacher effectiveness levels, gamer skill levels, etc.).

[0124] In some embodiments, the response system 406 can be further configured to receive one or several responses from the user and analyze these one or several responses. In some embodiments, for example, the response system 406 can be configured to translate the one or several responses into one or several observables. As used herein, an observable is a characterization of a received response. In some embodiments, the translation of the one or several responses into one or several observables can include determining whether the one or several responses are correct responses, also referred to herein as desired responses, or are incorrect responses, also referred to herein as undesired responses. In some embodiments, the translation of the one or several responses into one or several observables can include characterizing the degree to which one or several responses are desired responses and/or undesired responses. In some embodiments, one or several values can be generated by the response system 406 to reflect user performance in responding to the one or several data packets. In some embodiments, these one or several values can comprise one or several scores for one or several responses and/or data packets.

[0125] A content management server 102 also may include a presentation system 408. The presentation system 408 may be implemented using dedicated hardware within the content distribution network 100 (e.g., a presentation server 408), or using designated hardware and software resources within a shared content management server 102. The presentation system 408 can include a presentation engine that can be, for example, a software module running on the content delivery system.

[0126] The presentation system 408, also referred to herein as the presentation module or the presentation engine, may receive content resources from the packet selection system 402 and/or from the summary model system 404, and provide the resources to user devices 106. The presentation system 408 may determine the appropriate presentation format for the content resources based on the user characteristics and preferences, and/or the device capabilities of user devices 106. If needed, the presentation system 408 may convert the content resources to the appropriate presentation format and/or compress the content before transmission. In some embodiments, the presentation system 408 may also determine the appropriate transmission media and communication protocols for transmission of the content resources.

[0127] In some embodiments, the presentation system 408 may include specialized security and integration hardware 410, along with corresponding software components to implement the appropriate security features, content transmission and storage, to provide the supported network and client access models, and to support the performance and scalability requirements of the network 100. The security

and integration layer **410** may include some or all of the security and integration components **208** discussed above in FIG. 2, and may control the transmission of content resources and other data, as well as the receipt of requests and content interactions, to and from the user devices **106**, supervisor devices **110**, administrative servers **116**, and other devices in the network **100**.

[10128] With reference now to FIG. 5, a block diagram of an illustrative computer system is shown. The system **500** may correspond to any of the computing devices or servers of the content distribution network **100** described above, or any other computing devices described herein, and specifically can include, for example, one or several of the user devices **106**, the supervisor device **110**, and/or any of the servers **102**, **104**, **108**, **112**, **114**, **116**. In this example, computer system **500** includes processing units **504** that communicate with a number of peripheral subsystems via a bus subsystem **502**. These peripheral subsystems include, for example, a storage subsystem **510**, an I/O subsystem **526**, and a communications subsystem **532**.

[10129] Bus subsystem **502** provides a mechanism for letting the various components and subsystems of computer system **500** communicate with each other as intended. Although bus subsystem **502** is shown schematically as a single bus, alternative embodiments of the bus subsystem may utilize multiple buses. Bus subsystem **502** may be any of several types of bus structures including a memory bus or memory controller, a peripheral bus, and a local bus using any of a variety of bus architectures. Such architectures may include, for example, an Industry Standard Architecture (ISA) bus, Micro Channel Architecture (MCA) bus, Enhanced ISA (EISA) bus, Video Electronics Standards Association (VESA) local bus, and Peripheral Component Interconnect (PCI) bus, which can be implemented as a Mezzanine bus manufactured to the IEEE P1386.1 standard.

[10130] Processing unit **504**, which may be implemented as one or more integrated circuits (e.g., a conventional micro-processor or microcontroller), controls the operation of computer system **500**. One or more processors, including single core and/or multicore processors, may be included in processing unit **504**. As shown in the figure, processing unit **504** may be implemented as one or more independent processing units **506** and/or **508** with single or multicore processors and processor caches included in each processing unit. In other embodiments, processing unit **504** may also be implemented as a quad-core processing unit or larger multicore designs (e.g., hexa-core processors, octo-core processors, ten-core processors, or greater).

[10131] Processing unit **504** may execute a variety of software processes embodied in program code, and may maintain multiple concurrently executing programs or processes. At any given time, some or all of the program code to be executed can be resident in processor(s) **504** and/or in storage subsystem **510**. In some embodiments, computer system **500** may include one or more specialized processors, such as digital signal processors (DSPs), outboard processors, graphics processors, application-specific processors, and/or the like.

[10132] I/O subsystem **526** may include device controllers **528** for one or more user interface input devices and/or user interface output devices **530**. User interface input and output devices **530** may be integral with the computer system **500** (e.g., integrated audio/video systems, and/or touchscreen displays), or may be separate peripheral devices which are

attachable/detachable from the computer system **500**. The I/O subsystem **526** may provide one or several outputs to a user by converting one or several electrical signals to the user in perceptible and/or interpretable form, and may receive one or several inputs from the user by generating one or several electrical signals based on one or several user-caused interactions with the I/O subsystem such as the depressing of a key or button, the moving of a mouse, the interaction with a touchscreen or trackpad, the interaction of a sound wave with a microphone, or the like.

[10133] Input devices **530** may include a keyboard, pointing devices such as a mouse or trackball, a touchpad or touch screen incorporated into a display, a scroll wheel, a click wheel, a dial, a button, a switch, a keypad, audio input devices with voice command recognition systems, microphones, and other types of input devices. Input devices **530** may also include three dimensional (3D) mice, joysticks or pointing sticks, gamepads and graphic tablets, and audio/visual devices such as speakers, digital cameras, digital camcorders, portable media players, webcams, image scanners, fingerprint scanners, barcode reader 3D scanners, 3D printers, laser rangefinders, and eye gaze tracking devices. Additional input devices **530** may include, for example, motion sensing and/or gesture recognition devices that enable users to control and interact with an input device through a natural user interface using gestures and spoken commands, eye gesture recognition devices that detect eye activity from users and transform the eye gestures as input into an input device, voice recognition sensing devices that enable users to interact with voice recognition systems through voice commands, medical imaging input devices, MIDI keyboards, digital musical instruments, and the like.

[10134] Output devices **530** may include one or more display subsystems, indicator lights, or non-visual displays such as audio output devices, etc. Display subsystems may include, for example, cathode ray tube (CRT) displays, flat-panel devices, such as those using a liquid crystal display (LCD) or plasma display, light-emitting diode (LED) displays, projection devices, touch screens, and the like. In general, use of the term “output device” is intended to include all possible types of devices and mechanisms for outputting information from computer system **500** to a user or other computer. For example, output devices **530** may include, without limitation, a variety of display devices that visually convey text, graphics and audio/video information such as monitors, printers, speakers, headphones, automotive navigation systems, plotters, voice output devices, and modems.

[10135] Computer system **500** may comprise one or more storage subsystems **510**, comprising hardware and software components used for storing data and program instructions, such as system memory **518** and computer-readable storage media **516**. The system memory **518** and/or computer-readable storage media **516** may store program instructions that are loadable and executable on processing units **504**, as well as data generated during the execution of these programs.

[10136] Depending on the configuration and type of computer system **500**, system memory **318** may be stored in volatile memory (such as random access memory (RAM) **512**) and/or in non-volatile storage drives **514** (such as read-only memory (ROM), flash memory, etc.). The RAM **512** may contain data and/or program modules that are immediately accessible to and/or presently being operated

and executed by processing units **504**. In some implementations, system memory **518** may include multiple different types of memory, such as static random access memory (SRAM) or dynamic random access memory (DRAM). In some implementations, a basic input/output system (BIOS), containing the basic routines that help to transfer information between elements within computer system **500**, such as during start-up, may typically be stored in the non-volatile storage drives **514**. By way of example, and not limitation, system memory **518** may include application programs **520**, such as client applications, Web browsers, mid-tier applications, server applications, etc., program data **522**, and an operating system **524**.

[0137] Storage subsystem **510** also may provide one or more tangible computer-readable storage media **516** for storing the basic programming and data constructs that provide the functionality of some embodiments. Software (programs, code modules, instructions) that, when executed by a processor, provide the functionality described herein may be stored in storage subsystem **510**. These software modules or instructions may be executed by processing units **504**. Storage subsystem **510** may also provide a repository for storing data used in accordance with the present invention.

[0138] Storage subsystem **300** may also include a computer-readable storage media reader that can further be connected to computer-readable storage media **516**. Together and, optionally, in combination with system memory **518**, computer-readable storage media **516** may comprehensively represent remote, local, fixed, and/or removable storage devices plus storage media for temporarily and/or more permanently containing, storing, transmitting, and retrieving computer-readable information.

[0139] Computer-readable storage media **516** containing program code, or portions of program code, may include any appropriate media known or used in the art, including storage media and communication media, such as but not limited to, volatile and non-volatile, removable and non-removable media implemented in any method or technology for storage and/or transmission of information. This can include tangible computer-readable storage media such as RAM, ROM, electronically erasable programmable ROM (EEPROM), flash memory or other memory technology, CD-ROM, digital versatile disk (DVD), or other optical storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or other tangible computer readable media. This can also include nontangible computer-readable media, such as data signals, data transmissions, or any other medium which can be used to transmit the desired information and which can be accessed by computer system **500**.

[0140] By way of example, computer-readable storage media **516** may include a hard disk drive that reads from or writes to non-removable, nonvolatile magnetic media, a magnetic disk drive that reads from or writes to a removable, nonvolatile magnetic disk, and an optical disk drive that reads from or writes to a removable, nonvolatile optical disk such as a CD ROM, DVD, and Blu-Ray® disk, or other optical media. Computer-readable storage media **516** may include, but is not limited to, Zip® drives, flash memory cards, universal serial bus (USB) flash drives, secure digital (SD) cards, DVD disks, digital video tape, and the like. Computer-readable storage media **516** may also include, solid-state drives (SSD) based on non-volatile memory such

as flash-memory-based SSDs, enterprise flash drives, solid state ROM, and the like, SSDs based on volatile memory such as solid state RAM, dynamic RAM, static RAM, DRAM-based SSDs, magnetoresistive RAM (MRAM) SSDs, and hybrid SSDs that use a combination of DRAM and flash memory-based SSDs. The disk drives and their associated computer-readable media may provide non-volatile storage of computer-readable instructions, data structures, program modules, and other data for computer system **500**.

[0141] Communications subsystem **532** may provide a communication interface from computer system **500** and external computing devices via one or more communication networks, including local area networks (LANs), wide area networks (WANs) (e.g., the Internet), and various wireless telecommunications networks. As illustrated in FIG. 5, the communications subsystem **532** may include, for example, one or more network interface controllers (NICs) **534**, such as Ethernet cards, Asynchronous Transfer Mode NICs, Token Ring NICs, and the like, as well as one or more wireless communications interfaces **536**, such as wireless network interface controllers (WNICs), wireless network adapters, and the like. As illustrated in FIG. 5, the communications subsystem **532** may include, for example, one or more location determining features **538** such as one or several navigation system features and/or receivers, and the like. Additionally and/or alternatively, the communications subsystem **532** may include one or more modems (telephone, satellite, cable, ISDN), synchronous or asynchronous digital subscriber line (DSL) units, FireWire® interfaces, USB® interfaces, and the like. Communications subsystem **536** also may include radio frequency (RF) transceiver components for accessing wireless voice and/or data networks (e.g., using cellular telephone technology, advanced data network technology, such as 3G, 4G or EDGE (enhanced data rates for global evolution), WiFi (IEEE 802.11 family standards, or other mobile communication technologies, or any combination thereof), global positioning system (GPS) receiver components, and/or other components.

[0142] The various physical components of the communications subsystem **532** may be detachable components coupled to the computer system **500** via a computer network, a FireWire® bus, or the like, and/or may be physically integrated onto a motherboard of the computer system **500**. Communications subsystem **532** also may be implemented in whole or in part by software.

[0143] In some embodiments, communications subsystem **532** may also receive input communication in the form of structured and/or unstructured data feeds, event streams, event updates, and the like, on behalf of one or more users who may use or access computer system **500**. For example, communications subsystem **532** may be configured to receive data feeds in real-time from users of social networks and/or other communication services, web feeds such as Rich Site Summary (RSS) feeds, and/or real-time updates from one or more third party information sources (e.g., data aggregators **311**). Additionally, communications subsystem **532** may be configured to receive data in the form of continuous data streams, which may include event streams of real-time events and/or event updates (e.g., sensor data applications, financial tickers, network performance measuring tools, clickstream analysis tools, automobile traffic monitoring, etc.). Communications subsystem **532** may output such structured and/or unstructured data feeds, event

streams, event updates, and the like to one or more data stores **104** that may be in communication with one or more streaming data source computers coupled to computer system **500**.

[0144] Due to the ever-changing nature of computers and networks, the description of computer system **500** depicted in the figure is intended only as a specific example. Many other configurations having more or fewer components than the system depicted in the figure are possible. For example, customized hardware might also be used and/or particular elements might be implemented in hardware, firmware, software, or a combination. Further, connection to other computing devices, such as network input/output devices, may be employed. Based on the disclosure and teachings provided herein, a person of ordinary skill in the art will appreciate other ways and/or methods to implement the various embodiments.

[0145] With reference now to FIG. 6, a block diagram illustrating one embodiment of the communication network is shown. Specifically, FIG. 6 depicts one hardware configuration in which messages are exchanged between a source hub **602** via the communication network **120** that can include one or several intermediate hubs **604**. In some embodiments, the source hub **602** can be any one or several components of the content distribution network generating and initiating the sending of a message, and the terminal hub **606** can be any one or several components of the content distribution network receiving and not re-sending the message. In some embodiments, for example, the source hub **602** can be one or several of the user device **106**, the supervisor device **110**, and/or the server **102**, and the terminal hub **606** can likewise be one or several of the user device **106**, the supervisor device **110**, and/or the server **102**. In some embodiments, the intermediate hubs **604** can include any computing device that receives the message and resends the message to a next node.

[0146] As seen in FIG. 6, in some embodiments, each of the hubs **602**, **604**, **606** can be communicatively connected with the data store **104**. In such an embodiment, some or all of the hubs **602**, **604**, **606** can send information to the data store **104** identifying a received message and/or any sent or resent message. This information can, in some embodiments, be used to determine the completeness of any sent and/or received messages and/or to verify the accuracy and completeness of any message received by the terminal hub **606**.

[0147] In some embodiments, the communication network **120** can be formed by the intermediate hubs **604**. In some embodiments, the communication network **120** can comprise a single intermediate hub **604**, and in some embodiments, the communication network **120** can comprise a plurality of intermediate hubs. In one embodiment, for example, and as depicted in FIG. 6, the communication network **120** includes a first intermediate hub **604-A** and a second intermediate hub **604-B**.

[0148] With reference now to FIG. 7, a block diagram illustrating one embodiment of user device **106** and supervisor device **110** communication is shown. In some embodiments, for example, a user may have multiple devices that can connect with the content distribution network **100** to send or receive information. In some embodiments, for example, a user may have a personal device such as a mobile device, a Smartphone, a tablet, a Smartwatch, a laptop, a PC, or the like. In some embodiments, the other device can be

any computing device in addition to the personal device. This other device can include, for example, a laptop, a PC, a Smartphone, a tablet, a Smartwatch, or the like. In some embodiments, the other device differs from the personal device in that the personal device is registered as such within the content distribution network **100** and the other device is not registered as a personal device within the content distribution network **100**.

[0149] Specifically with respect to FIG. 7, the user device **106** can include a personal user device **106-A** and one or several other user devices **106-B**. In some embodiments, one or both of the personal user device **106-A** and the one or several other user devices **106-B** can be communicatively connected to the content management server **102** and/or to the navigation system **122**. Similarly, the supervisor device **110** can include a personal supervisor device **110-A** and one or several other supervisor devices **110-B**. In some embodiments, one or both of the personal supervisor device **110-A** and the one or several other supervisor devices **110-B** can be communicatively connected to the content management server **102** and/or to the navigation system **122**.

[0150] In some embodiments, the content distribution network can send one or more alerts to one or more user devices **106** and/or one or more supervisor devices **110** via, for example, the communication network **120**. In some embodiments, the receipt of the alert can result in the launching of an application within the receiving device, and in some embodiments, the alert can include a link that, when selected, launches the application or navigates a web-browser of the device of the selector of the link to a page or portal associated with the alert.

[0151] In some embodiments, for example, the providing of this alert can include the identification of one or several user devices **106** and/or student-user accounts associated with the student-user and/or one or several supervisor devices **110** and/or supervisor-user accounts associated with the supervisor-user. After these one or several devices **106**, **110** and/or accounts have been identified, the providing of this alert can include determining an active device of the devices **106**, **110** based on determining which of the devices **106**, **110** and/or accounts are actively being used, and then providing the alert to that active device.

[0152] Specifically, if the user is actively using one of the devices **106**, **110** such as the other user device **106-B** and the other supervisor device **110-B**, and/or accounts, the alert can be provided to the user via that other device **106-B**, **110-B** and/or account that is actively being used. If the user is not actively using an other device **106-B**, **110-B** and/or account, a personal device **106-A**, **110-A** device, such as a smart phone or tablet, can be identified and the alert can be provided to this personal device **106-A**, **110-A**. In some embodiments, the alert can include code to direct the default device to provide an indicator of the received alert such as, for example, an aural, tactile, or visual indicator of receipt of the alert.

[0153] In some embodiments, the recipient device **106**, **110** of the alert can provide an indication of receipt of the alert. In some embodiments, the presentation of the alert can include the control of the I/O subsystem **526** to, for example, provide an aural, tactile, and/or visual indicator of the alert and/or of the receipt of the alert. In some embodiments, this can include controlling a screen of the supervisor device **110** to display the alert, data contained in alert and/or an indicator of the alert.

[0154] With reference now to FIG. 8, a schematic illustration of one embodiment of an application stack, and particularly of a stack 650 is shown. In some embodiments, the content distribution network 100 can comprise a portion of the stack 650 that can include an infrastructure layer 652, a platform layer 654, an applications layer 656, and a products layer 658. In some embodiments, the stack 650 can comprise some or all of the layers, hardware, and/or software to provide one or several desired functionalities and/or productions.

[0155] As depicted in FIG. 8, the infrastructure layer 652 can include one or several servers, communication networks, data stores, privacy servers, and the like. In some embodiments, the infrastructure layer can further include one or several user devices 106 and/or supervisor devices 110 connected as part of the content distribution network.

[0156] The platform layer can include one or several platform software programs, modules, and/or capabilities. These can include, for example, identification services, security services, and/or adaptive platform services 660. In some embodiments, the identification services can, for example, identify one or several users, components of the content distribution network 100, or the like. The security services can monitor the content distribution network for one or several security threats, breaches, viruses, malware, or the like. The adaptive platform services 660 can receive information from one or several components of the content distribution network 100 and can provide predictions, models, recommendations, or the like based on that received information. The functionality of the adaptive platform services 660 will be discussed in greater detail in FIGS. 9-12, below.

[0157] The applications layer 656 can include software or software modules upon or in which one or several product softwares or product software modules can operate. In some embodiments, the applications layer 656 can include, for example, a management system, record system, or the like. In some embodiments, the management system can include, for example, a Learning Management System (LMS), a Content Management System (CMS), or the like. The management system can be configured to control the delivery of one or several resources to a user and/or to receive one or several responses from the user. In some embodiments, the records system can include, for example, a virtual gradebook, a virtual counselor, or the like.

[0158] The products layer can include one or several software products and/or software module products. These software products and/or software module products can provide one or several services and/or functionalities to one or several users of the software products and/or software module products.

[0159] With reference now to FIG. 9-12, schematic illustrations of embodiments of communication and processing flow of modules within the content distribution network 100 are shown. In some embodiments, the communication and processing can be performed in portions of the platform layer 654 and/or applications layer 656. FIG. 9 depicts a first embodiment of such communications or processing that can be in the platform layer 654 and/or applications layer 656 via the message channel 412.

[0160] The platform layer 654 and/or applications layer 656 can include a plurality of modules that can be embodied in software or hardware. In some embodiments, some or all of the modules can be embodied in hardware and/or soft-

ware at a single location, and in some embodiments, some or all of these modules can be embodied in hardware and/or software at multiple locations. These modules can perform one or several processes including, for example, a presentation process 670, a response process 676, a summary model process 680, and a packet selection process 684.

[0161] The presentation process 670 can, in some embodiments, include one or several methods and/or steps to deliver content to one or several user devices 106 and/or supervisor devices 110. The presentation process 670 can be performed by a presenter module 672 and a view module 674. The presenter module 672 can be a hardware or software module of the content distribution network 100, and specifically of the server 102. In some embodiments, the presenter module 672 can include one or several portions, features, and/or functionalities that are located on the server 102 and/or one or several portions, features, and/or functionalities that are located on the user device 106. In some embodiments, the presenter module 672 can be embodied in the presentation system 408.

[0162] The presenter module 672 can control the providing of content to one or several user devices 106 and/or supervisor devices 110. Specifically, the presenter module 672 can control the generation of one or several messages to provide content to one or several desired user devices 106 and/or supervisor devices 110. The presenter module 672 can further control the providing of these one or several messages to the desired one or several desired user devices 106 and/or supervisor devices 110. Thus, in some embodiments, the presenter module 672 can control one or several features of the communications subsystem 532 to generate and send one or several electrical signals comprising content to one or several user devices 106 and/or supervisor devices 110.

[0163] In some embodiments, the presenter module 672 can control and/or manage a portion of the presentation functions of the presentation process 670, and can specifically manage an “outer loop” of presentation functions. As used herein, the outer loop refers to tasks relating to the tracking of a user’s progress through all or a portion of a group of data packets. In some embodiments, this can include the identification of one or several completed data packets or nodes and/or the non-adaptive selection of one or several next data packets or nodes according to, for example, one or several fixed rules. Such non-adaptive selection does not rely on the use of predictive models, but rather on rules identifying next data packets based on data relating to the completion of one or several previously completed data packets or assessments and/or whether one or several previously completed data packets were successfully completed.

[0164] In some embodiments, and due to the management of the outer loop of presentation functions including the non-adaptive selection of one or several next data packets, nodes, or tasks by the presenter module, the presenter module can function as a recommendation engine referred to herein as a first recommendation engine or a rules-based recommendation engine. In some embodiments, the first recommendation engine can be configured to select a next node for a user based on one or all of: the user’s current location in the content network; potential next nodes; the user’s history including the user’s previous responses; and one or several guard conditions associated with the potential

next nodes. In some embodiments, a guard condition defines one or several prerequisites for entry into, or exit from, a node.

[0165] In some embodiments, the presenter module 672 can include a portion located on the server 102 and/or a portion located on the user device 106. In some embodiments, the portion of the presenter module 672 located on the server 102 can receive data packet information and provide a subset of the received data packet information to the portion of the presenter module 672 located on the user device 106. In some embodiments, this segregation of functions and/or capabilities can prevent solution data from being located on the user device 106 and from being potentially accessible by the user of the user device 106.

[0166] In some embodiments, the portion of the presenter module 672 located on the user device 106 can be further configured to receive the subset of the data packet information from the portion of the presenter module 672 located on the server 102 and provide that subset of the data packet information to the view module 674. In some embodiments, the portion of the presenter module 672 located on the user device 106 can be further configured to receive a content request from the view module 674 and to provide that content request to the portion of the presenter module 674 located on the server 102.

[0167] The view module 674 can be a hardware or software module of some or all of the user devices 106 and/or supervisor devices 110 of the content distribution network 100. The view module 674 can receive one or several electrical signals and/or communications from the presenter module 672 and can provide the content received in those one or several electrical signals and/or communications to the user of the user device 106 and/or supervisor device 110 via, for example, the I/O subsystem 526.

[0168] In some embodiments, the view module 674 can control and/or monitor an “inner loop” of presentation functions. As used herein, the inner loop refers to tasks relating to the tracking and/or management of a user’s progress through a data packet. This can specifically relate to the tracking and/or management of a user’s progression through one or several pieces of content, questions, assessments, and/or the like of a data packet. In some embodiments, this can further include the selection of one or several next pieces of content, next questions, next assessments, and/or the like of the data packet for presentation and/or providing to the user of the user device 106.

[0169] In some embodiments, one or both of the presenter module 672 and the view module 674 can comprise one or several presentation engines. In some embodiments, these one or several presentation engines can comprise different capabilities and/or functions. In some embodiments, one of the presentation engines can be configured to track the progress of a user through a single data packet, task, content item, or the like, and in some embodiments, one of the presentation engines can track the progress of a user through a series of data packets, tasks, content items, or the like.

[0170] The response process 676 can comprise one or several methods and/or steps to evaluate a response. In some embodiments, this can include, for example, determining whether the response comprises a desired response and/or an undesired response. In some embodiments, the response process 676 can include one or several methods and/or steps to determine the correctness and/or incorrectness of one or several received responses. In some embodiments, this can

include, for example, determining the correctness and/or incorrectness of a multiple choice response, a true/false response, a short answer response, an essay response, or the like. In some embodiments, the response processor can employ, for example, natural language processing, semantic analysis, or the like in determining the correctness or incorrectness of the received responses.

[0171] In some embodiments, the response process 676 can be performed by a response processor 678. The response processor 678 can be a hardware or software module of the content distribution network 100, and specifically of the server 102. In some embodiments, the response processor 678 can be embodied in the response system 406. In some embodiments, the response processor 678 can be communicatively connected to one or more of the modules of the presentation process 760 such as, for example, the presenter module 672 and/or the view module 674. In some embodiments, the response processor 678 can be communicatively connected with, for example, the message channel 412 and/or other components and/or modules of the content distribution network 100.

[0172] The summary model process 680 can comprise one or several methods and/or steps to generate and/or update one or several models. In some embodiments, this can include, for example, implementing information received either directly or indirectly from the response processor 678 to update one or several models. In some embodiments, the summary model process 680 can include the update of a model relating to one or several user attributes such as, for example, a user skill model, a user knowledge model, a learning style model, or the like. In some embodiments, the summary model process 680 can include the update of a model relating to one or several content attributes including attributes relating to a single content item and/or data packet and/or attributes relating to a plurality of content items and/or data packets. In some embodiments, these models can relate to an attribute of the one or several data packets such as, for example, difficulty, discrimination, required time, or the like.

[0173] In some embodiments, the summary model process 680 can be performed by the model engine 682. In some embodiments, the model engine 682 can be a hardware or software module of the content distribution network 100, and specifically of the server 102. In some embodiments, the model engine 682 can be embodied in the summary model system 404.

[0174] In some embodiments, the model engine 682 can be communicatively connected to one or more of the modules of the presentation process 760 such as, for example, the presenter module 672 and/or the view module 674, can be connected to the response processor 678 and/or the recommendation. In some embodiment, the model engine 682 can be communicatively connected to the message channel 412 and/or other components and/or modules of the content distribution network 100.

[0175] The packet selection process 684 can comprise one or several steps and/or methods to identify and/or select a data packet for presentation to a user. In some embodiments, this data packet can comprise a plurality of data packets. In some embodiments, this data packet can be selected according to one or several models updated as part of the summary model process 680. In some embodiments, this data packet can be selected according to one or several rules, probabilities, models, or the like. In some embodiments, the one or

several data packets can be selected by the combination of a plurality of models updated in the summary model process 680 by the model engine 682. In some embodiments, these one or several data packets can be selected by a recommendation engine 686. The recommendation engine 686 can be a hardware or software module of the content distribution network 100, and specifically of the server 102. In some embodiments, the recommendation engine 686 can be embodied in the packet selection system 402. In some embodiments, the recommendation engine 686 can be communicatively connected to one or more of the modules of the presentation process 670, the response process 676, and/or the summary model process 680 either directly and/or indirectly via, for example, the message channel.

[0176] In some embodiments, and as depicted in FIG. 9, a presenter module 672 can receive a data packet for presentation to a user device 106. This data packet can be received, either directly or indirectly from a recommendation engine 686. In some embodiments, for example, the presenter module 672 can receive a data packet for providing to a user device 106 from the recommendation engine 686, and in some embodiments, the presenter module 672 can receive an identifier of a data packet for providing to a user device 106 via a view module 674. This can be received from the recommendation engine 686 via a message channel 412. Specifically, in some embodiments, the recommendation engine 686 can provide data to the message channel 412 indicating the identification and/or selection of a data packet for providing to a user via a user device 106. In some embodiments, this data indicating the identification and/or selection of the data packet can identify the data packet and/or can identify the intended recipient of the data packet.

[0177] The message channel 412 can output this received data in the form of a data stream 690 which can be received by, for example, the presenter module 672, the model engine 682, and/or the recommendation engine 686. In some embodiments, some or all of: the presenter module 672, the model engine 682, and/or the recommendation engine 686 can be configured to parse and/or filter the data stream 690 to identify data and/or events relevant to their operation. Thus, for example, the presenter module 672 can be configured to parse the data stream for information and/or events relevant to the operation of the presenter module 672.

[0178] In some embodiments, the presenter module 672 can extract the data packet from the data stream 690 and/or extract data identifying the data packet and/or indicating the selecting of a data packet from the data stream. In the event that data identifying the data packet is extracted from the data stream 690, the presenter module 672 can request and receive the data packet from the database server 104, and specifically from the content library database 303. In embodiments in which data indicating the selection of a data packet is extracted from the data stream 690, the presenter module 672 can request and receive the data packet from the recommendation engine 686 and then request and receive the data packet from the database server 104, and specifically from the content library database 303, and in some embodiments in which data indicating the selection of a data packet is extracted from the data stream 690, the presenter module 672 can request and receive the data packet from the recommendation engine 686.

[0179] The presenter module can then provide the data packet and/or portions of the data packet to the view module 674. In some embodiments, for example, the presenter

module 672 can retrieve one or several rules and/or conditions that can be, for example, associated with the data packet and/or stored in the database server 104. In some embodiments, these rules and/or conditions can identify portions of a data packet for providing to the view module 674 and/or portions of a data packet to not provide to the view module 674. In some embodiments, for example, sensitive portions of a data packet, such as, for example, solution information to any questions associated with a data packet, is not provided to the view module 674 to prevent the possibility of undesired access to those sensitive portions of the data packet. Thus, in some embodiments, the one or several rules and/or conditions can identify portions of the data packet for providing to the view module 674 and/or portions of the data packet for not providing to the view module.

[0180] In some embodiments, the presenter module 672 can, according to the one or more rules and/or conditions, generate and transmit an electronic message containing all or portions of the data packet to the view module 674. The view module 674 can receive these all or portions of the data packet and can provide all or portions of this information to the user of the user device 106 associated with the view module 674 via, for example, the I/O subsystem 526. In some embodiments, as part of the providing of all or portions of the data packet to the user of the view module 674, one or several user responses can be received by the view module 674. In some embodiments, these one or several user responses can be received via the I/O subsystem 526 of the user device 106.

[0181] After one or several user responses have been received, the view module 674 can provide the one or several user responses to the response processor 678. In some embodiments, these one or several responses can be directly provided to the response processor 678, and in some embodiments, these one or several responses can be provided indirectly to the response processor 678 via the message channel 412.

[0182] After the response processor 678 receives the one or several responses, the response processor 678 can determine whether the responses are desired responses and/or the degree to which the received responses are desired responses. In some embodiments, the response processor can make this determination via, for example, use of one or several techniques, including, for example, natural language processing (NLP), semantic analysis, or the like.

[0183] In some embodiments, the response processor can determine whether a response is a desired response and/or the degree to which a response is a desired response with comparative data which can be associated with the data packet. In some embodiments, this comparative data can comprise, for example, an indication of a desired response and/or an indication of one or several undesired responses, a response key, a response rubric comprising one criterion or several criteria for determining the degree to which a response is a desired response, or the like. In some embodiments, the comparative data can be received as a portion of and/or associated with a data packet. In some embodiments, the comparative data can be received by the response processor 678 from the presenter module 672 and/or from the message channel 412. In some embodiments, the response data received from the view module 674 can comprise data identifying the user and/or the data packet or portion of the data packet with which the response is

associated. In some embodiments in which the response processor 678 merely receives data identifying the data packet and/or portion of the data packet associated with the one or several responses, the response processor 678 can request and/or receive comparative data from the database server 104, and specifically from the content library database 303 of the database server 104.

[0184] After the comparative data has been received, the response processor 678 determines whether the one or several responses comprise desired responses and/or the degree to which the one or several responses comprise desired responses. The response processor can then provide the data characterizing whether the one or several responses comprise desired responses and/or the degree to which the one or several responses comprise desired responses to the message channel 412. The message channel can, as discussed above, include the output of the response processor 678 in the data stream 690 which can be constantly output by the message channel 412.

[0185] In some embodiments, the model engine 682 can subscribe to the data stream 690 of the message channel 412 and can thus receive the data stream 690 of the message channel 412 as indicated in FIG. 9. The model engine 682 can monitor the data stream 690 to identify data and/or events relevant to the operation of the model engine. In some embodiments, the model engine 682 can monitor the data stream 690 to identify data and/or events relevant to the determination of whether a response is a desired response and/or the degree to which a response is a desired response.

[0186] When a relevant event and/or relevant data are identified by the model engine, the model engine 682 can take the identified relevant event and/or relevant data and modify one or several models. In some embodiments, this can include updating and/or modifying one or several models relevant to the user who provided the responses, updating and/or modifying one or several models relevant to the data packet associated with the responses, and/or the like. In some embodiments, these models can be retrieved from the database server 104, and, in some embodiments, can be retrieved from the model data source 309 of the database server 104.

[0187] After the models have been updated, the updated models can be stored in the database server 104. In some embodiments, the model engine 682 can send data indicative of the event of the completion of the model update to the message channel 412. The message channel 412 can incorporate this information into the data stream 690 which can be received by the recommendation engine 686. The recommendation engine 686 can monitor the data stream 690 to identify data and/or events relevant to the operation of the recommendation engine 686. In some embodiments, the recommendation engine 686 can monitor the data stream 690 to identify data and/or events relevant to the updating of one or several models by the model engine 682.

[0188] When the recommendation engine 686 identifies information in the data stream 690 indicating the completion of the summary model process 680 for models relevant to the user providing the response and/or for models relevant to the data packet provided to the user, the recommendation engine 686 can identify and/or select a next data packet for providing to the user and/or to the presentation process 470. In some embodiments, this selection of the next data packet can be performed according to one or several rules and/or conditions. After the next data packet has been selected, the

recommendation engine 686 can provide information to the model engine 682 identifying the next selected data packet and/or to the message channel 412 indicating the event of the selection of the next content item. After the message channel 412 receives information identifying the selection of the next content item and/or receives the next content item, the message channel 412 can include this information in the data stream 690 and the process discussed with respect to FIG. 9 can be repeated.

[0189] With reference now to FIG. 10, a schematic illustration of a second embodiment of communication or processing that can be in the platform layer 654 and/or applications layer 656 via the message channel 412 is shown. In the embodiment depicted in FIG. 10, the data packet provided to the presenter module 672 and then to the view module 674 does not include a prompt for a user response and/or does not result in the receipt of a user response. As no response is received, when the data packet is completed, nothing is provided to the response processor 678, but rather data indicating the completion of the data packet is provided from one of the view module 674 and/or the presenter module 672 to the message channel 412. The data is then included in the data stream 690 and is received by the model engine 682 which uses the data to update one or several models. After the model engine 682 has updated the one or several models, the model engine 682 provides data indicating the completion of the model updates to the message channel 412. The message channel 412 then includes the data indicating the completion of the model updates in the data stream 690 and the recommendation engine 686, which can subscribe to the data stream 690, can extract the data indicating the completion of the model updates from the data stream 690. The recommendation engine 686 can then identify a next one or several data packets for providing to the presenter module 672, and the recommendation engine 686 can then, either directly or indirectly, provide the next one or several data packets to the presenter module 672.

[0190] In some embodiments, of the communication as shown in FIGS. 9 and 10, all communications between any of the presenter module 672, the response processor 678, the model engine 682, and the recommendation engine 686 can pass through the message channel 412. Alternatively, in some embodiments, some of the communications between any of the presenter module 672, the response processor 678, the model engine 682, and the recommendation engine 686 can pass through the message channel and others of the communications between any of the presenter module 672, the response processor 678, the model engine 682, and the recommendation engine 686 can be direct.

[0191] With reference now to FIG. 11, a schematic illustration of an embodiment of dual communication, or hybrid communication, in the platform layer 654 and/or applications layer 656 is shown. Specifically, in this embodiment, some communication is synchronous with the completion of one or several tasks and some communication is asynchronous. In the embodiment depicted in FIG. 11, the presenter module 972 communicates synchronously with the model engine 682 via a direct communication 692 and communicates asynchronously with the model engine 682 via the message channel 412.

[0192] In some embodiments, and as depicted in FIG. 11, the synchronous communication and/or the operation of the presenter module 672, the response processor 678, the model engine 682, and the recommendation engine 686 can

be directed and/or controlled by a controller. In some embodiments, this controller can be part of the server 102 and/or located in any one or more of the presenter module 672, the response processor 678, the model engine 682, and the recommendation engine 686. In some embodiments, this controller can be located in the presenter module 672, which presenter module can control communications with and between itself and the response processor 678, the model engine 682, and the recommendation engine 686, and the presenter module can thus control the functioning of the response processor 678, the model engine 682, and the recommendation engine 686.

[0193] Specifically, and with reference to FIG. 11, the presenter module 672 can receive and/or select a data packet for presentation to the user device 106 via the view module 674. In some embodiments, the presenter module 672 can identify all or portions of the data packet that can be provided to the view module 674 and portions of the data packet for retaining from the view module 674. In some embodiments, the presenter module can provide all or portions of the data packet to the view module 674. In some embodiments, and in response to the receipt of all or portions of the data packet, the view module 674 can provide a confirmation of receipt of the all or portions of the data packet and can provide those all or portions of the data packet to the user via the user device 106. In some embodiments, the view module 674 can provide those all or portions of the data packet to the user device 106 while controlling the inner loop of the presentation of the data packet to the user via the user device 106.

[0194] After those all or portions of the data packet have been provided to the user device 106, a response indicative of the completion of one or several tasks associated with the data packet can be received by the view module 674 from the user device 106, and specifically from the I/O subsystem 526 of the user device 106. In response to this receive, the view module 674 can provide an indication of this completion status to the presenter module 672 and/or can provide the response to the response processor 678.

[0195] After the response has been received by the response processor 678, the response processor 678 can determine whether the received response is a desired response. In some embodiments, this can include, for example, determining whether the response comprises a correct answer and/or the degree to which the response comprises a correct answer.

[0196] After the response processor has determined whether the received response is a desired response, the response processor 678 can provide an indicator of the result of the determination of whether the received response is a desired response to the presenter module 672. In response to the receipt of the indicator of whether the result of the determination of whether the received response is a desired response, the presenter module 672 can synchronously communicate with the model engine 682 via a direct communication 692 and can asynchronously communicate with model engine 682 via the message channel 412. In some embodiments, the synchronous communication can advantageously include two-way communication between the model engine 682 and the presenter module 672 such that the model engine 682 can provide an indication to the presenter module 672 when model updating is completed by the model engine.

[0197] After the model engine 682 has received one or both of the synchronous and asynchronous communications, the model engine 682 can update one or several models relating to, for example, the user, the data packet, or the like. After the model engine 682 has completed the updating of the one or several models, the model engine 682 can send a communication to the presenter module 672 indicating the completion of the updated one or several modules.

[0198] After the presenter module 672 receives the communication indicating the completion of the updating of the one or several models, the presenter module 672 can send a communication to the recommendation engine 686 requesting identification of a next data packet. As discussed above, the recommendation engine 686 can then retrieve the updated model and retrieve the user information. With the updated models and the user information, the recommendation engine can identify a next data packet for providing to the user, and can provide the data packet to the presenter module 672. In some embodiments, the recommendation engine 686 can further provide an indication of the next data packet to the model engine 682, which can use this information relating to the next data packet to update one or several models, either immediately, or after receiving a communication from the presenter module 672 subsequent to the determination of whether a received response for that data packet is a desired response.

[0199] With reference now to FIG. 12, a schematic illustration of one embodiment of the presentation process 670 is shown. Specifically, FIG. 12 depicts multiple portions of the presenter module 672, namely, the external portion 673 and the internal portion 675. In some embodiments, the external portion 673 of the presenter module 672 can be located in the server, and in some embodiments, the internal portion 675 of the presenter module 672 can be located in the user device 106. In some embodiments, the external portion 673 of the presenter module can be configured to communicate and/or exchange data with the internal portion 675 of the presenter module 672 as discussed herein. In some embodiments, for example, the external portion 673 of the presenter module 672 can receive a data packet and can parse the data packet into portions for providing to the internal portion 675 of the presenter module 672 and portions for not providing to the internal portion 675 of the presenter module 672. In some embodiments, the external portion 673 of the presenter module 672 can receive a request for additional data and/or an additional data packet from the internal portion 675 of the presenter module 672. In such an embodiment, the external portion 673 of the presenter module 672 can identify and retrieve the requested data and/or the additional data packet from, for example, the database server 104 and more specifically from the content library database 104.

[0200] With reference now to FIG. 13, a flowchart illustrating one embodiment of a process 440 for data management is shown. In some embodiments, the process 440 can be performed by the content management server 102, and more specifically by the presentation system 408 and/or by the presentation module or presentation engine. In some embodiments, the process 440 can be performed as part of the presentation process 670.

[0201] The process 440 begins at block 442, wherein a data packet is identified. In some embodiments, the data packet can be a data packet for providing to a student-user. In some embodiments, the data packet can be identified

based on a communication received either directly or indirectly from the recommendation engine 686.

[0202] After the data packet has been identified, the process 440 proceeds to block 444, wherein the data packet is requested. In some embodiments, this can include the requesting of information relating to the data packet such as the data forming the data packet. In some embodiments, this information can be requested from, for example, the content library database 303. After the data packet has been requested, the process 440 proceeds to block 446, wherein the data packet is received. In some embodiments, the data packet can be received by the presentation system 408 from, for example, the content library database 303.

[0203] After the data packet has been received, the process 440 proceeds to block 448, wherein one or several data components are identified. In some embodiments, for example, the data packet can include one or several data components which can, for example, contain different data. In some embodiments, one of these data components, referred to herein as a presentation component, can include content for providing to the student user, which content can include one or several requests and/or questions and/or the like. In some embodiments, one of these data components, referred to herein as a response component, can include data used in evaluating one or several responses received from the user device 106 in response to the data packet, and specifically in response to the presentation component and/or the one or several requests and/or questions of the presentation component. Thus, in some embodiments, the response component of the data packet can be used to ascertain whether the user has provided a desired response or an undesired response.

[0204] After the data components have been identified, the process 440 proceeds to block 450, wherein a delivery data packet is identified. In some embodiments, the delivery data packet can include the one or several data components of the data packets for delivery to a user such as the student-user via the user device 106. In some embodiments, the delivery packet can include the presentation component, and in some embodiments, the delivery packet can exclude the response packet. After the delivery data packet has been generated, the process 440 proceeds to block 452, wherein the delivery data packet is provided to the user device 106 and more specifically to the view module 674. In some embodiments, this can include providing the delivery data packet to the user device 106 via, for example, the communication network 120.

[0205] After the delivery data packet has been provided to the user device 106, the process 440 proceeds to block 454, wherein the data packet and/or one or several components thereof are sent to and/or provided to the response processor 678. In some embodiments, this sending of the data packet and/or one or several components thereof to the response processor can include receiving a response from the student-user, and sending the response to the student-user to the response processor simultaneous with the sending of the data packet and/or one or several components thereof to the response processor. In some embodiments, for example, this can include providing the response component to the response processor. In some embodiments, the response component can be provided to the response processor from the presentation system 408.

[0206] With reference now to FIG. 14, a flowchart illustrating one embodiment of a process 460 for evaluating a

response is shown. In some embodiments, the process can be performed as a part of the response process 676 and can be performed by, for example, the response system 406 and/or by the response processor 678. In some embodiments, the process 460 can be performed by the response system 406 in response to the receipt of a response, either directly or indirectly, from the user device 106 or from the view module 674.

[0207] The process 460 begins at block 462, wherein a response is received from, for example, the user device 106 via, for example, the communication network 120. After the response has been received, the process 460 proceeds to block 464, wherein the data packet associated with the response is received. In some embodiments, this can include receiving all or one or several components of the data packet such as, for example, the response component of the data packet. In some embodiments, the data packet can be received by the response processor from the presentation engine.

[0208] After the data packet has been received, the process 460 proceeds to block 466, wherein the response type is identified. In some embodiments, this identification can be performed based on data, such as metadata associated with the response. In other embodiments, this identification can be performed based on data packet information such as the response component.

[0209] In some embodiments, the response type can identify one or several attributes of the one or several requests and/or questions of the data packet such as, for example, the request and/or question type. In some embodiments, this can include identifying some or all of the one or several requests and/or questions as true/false, multiple choice, short answer, essay, or the like.

[0210] After the response type has been identified, the process 460 proceeds to block 468, wherein the data packet and the response are compared to determine whether the response comprises a desired response and/or an undesired response. In some embodiments, this can include comparing the received response and the data packet to determine if the received response matches all or portions of the response component of the data packet, to determine the degree to which the received response matches all or portions of the response component, to determine the degree to which the received response embodies one or several qualities identified in the response component of the data packet, or the like. In some embodiments, this can include classifying the response according to one or several rules. In some embodiments, these rules can be used to classify the response as either desired or undesired. In some embodiments, these rules can be used to identify one or several errors and/or misconceptions evidenced in the response. In some embodiments, this can include, for example: use of natural language processing software and/or algorithms; use of one or several digital thesauruses; use of lemmatization software, dictionaries, and/or algorithms; or the like.

[0211] After the data packet and the response have been compared, the process 460 proceeds to block 470 wherein response desirability is determined. In some embodiments this can include, based on the result of the comparison of the data packet and the response, whether the response is a desired response or is an undesired response. In some embodiments, this can further include quantifying the degree to which the response is a desired response. This determination can include, for example, determining if the

response is a correct response, an incorrect response, a partially correct response, or the like. In some embodiments, the determination of response desirability can include the generation of a value characterizing the response desirability and the storing of this value in one of the databases 104 such as, for example, the user profile database 301. After the response desirability has been determined, the process 460 proceeds to block 472, wherein an assessment value is generated. In some embodiments, the assessment value can be an aggregate value characterizing response desirability for one or more a plurality of responses. This assessment value can be stored in one of the databases 104 such as the user profile database 301.

[0212] With reference now to FIG. 15, a schematic illustration of one embodiment of the content network 700 is shown. In some embodiments, the content network 700, also referred to herein as the object network can include a plurality of edges 701 linking a plurality of nodes 702. In some embodiments, the edges 701 can link the nodes 702 in prerequisite relationships such that a single edge 701 links a pair of nodes 702, one of which of the pair of nodes 702 is a prerequisite to the other of the pair of nodes 702. In FIG. 15, this prerequisite relationship is indicated by the direction of the arrow of the edge 701 linking a pair of nodes 702 such that the origin of the arrow is at the prerequisite node and the head of the arrow is at the non-prerequisite node.

[0213] In the embodiment depicted in FIG. 15, these nodes 702 include node 1, node 2, node 3, node 4, node 5, and node 6. As further depicted in FIG. 15, some of the nodes 702 such as, for example, node 1 may be unassociated with content whereas others of the nodes 702 such as, for example, node 2, node 3, node 4, node 5, and node 6 can be associated with content 704. Specifically, node 2 is associated with content b, node 3 is associated with content c, node 4 is associated with content letter d, node 5 is associated with placeholder content, and node 6 is associated with content f. In some embodiments, a placeholder node is a node whose content is not fixed but can be, for example, adaptively selected from, for example, a set of potential content or content objects that can be found in the database server 104. In some embodiments, this content can be adaptively selected based on one or several predictive or statistical models relating to content, the user, user skill level, or the like.

[0214] In some embodiments, the content of a placeholder node can be selected from one or several databases relevant to that placeholder node. These databases can be stored within, for example, the content library data store 303. In some embodiments, the content for presentation to a user at a placeholder node can be determined by the recommendation engine 686 based on statistical models which can be, for example, updated by the model engine based on data relating to one or several users past performance.

[0215] In some embodiments, some of the nodes 702 can be associated with a guard condition, also referred to herein as a gate condition, as indicated by conditions 706. In some embodiments, the guard condition can identify a prerequisite for entering and/or exiting a node 702. These guard conditions can, for example, predicate entry into a node 702 on a user's mastery and/or failure to master one or several previous nodes, objectives such as objective one 708, skill levels, or the like. In some embodiments, these guard conditions can predicate entry into and node 702 based on whether one or several responses to a previous node and/or

to an assessment were correct or incorrect, or on the ratio of correct incorrect responses to a previous node or assessment. In some embodiments, these guard conditions can predicate entry into a node 702 based on the completion of a previous node 702 such as, for example, entry into node 3 is predicated on the completion of node 2. In some embodiments, each of these guard conditions can be associated with a node 702 and/or on edge 701. In some embodiments, each guard condition is associated with a node 702. In some embodiments, this can result in an entry in a database of nodes 702 identifying a specific node 702 including data relating to and/or identifying one or several guard conditions and/or one or several pointers to those one or several guard conditions.

[0216] In some embodiments, some or all of the nodes 702 of the content network 700 can be arranged into one or several loops 703. In some embodiments, the duration of the loop can be based on one or several guard conditions determining when a user can exit the loop and/or on criteria identifying the minimum and/or maximum number of cycles that a user can remain in the loop. Specifically as shown in FIG. 700, node 1 comprises an indicator of the maximum number of cycles that a user can remain in the loops 703, and specifically indicates that the maximum number of cycles that a user can remain in the loops 703 is 2. In some embodiments, for example, the recommendation engine directing and/or controlling the movement of the user through the content network 700 can comprise a count that can be incremented each time that a user passes through a loop. This count can be compared to the indicator of the minimum and/or maximum number of cycles that user can remain in the loop to determine whether the user must be readmitted through the loop or whether the user should be prohibited from readmission into the loop.

[0217] As further seen in FIG. 15, node 1 comprises an edge 701 connecting node 1 to node 6. In some embodiments, when the count reaches the indicator of the maximum number of cycles that a user can remain in the loops, the user is then directed from node 1 to node 6.

[0218] With reference now to FIG. 16, a flowchart illustrating one embodiment of a process 800 for artificial intelligence based content recommendation and provisioning is shown. The process 800 can be performed by all or portions of the CDN 100 including, for example, the server 102, the AI agent 124, or the like.

[0219] The process 800 begins at block 802, wherein user identification information, also referred to herein as a user identifier, is received. In some embodiments, user identification information can include a username, password, login information, a unique user identifier, or the like. In some embodiments, the user identification information can be received by the user device 106 via the I/O subsystem 526 of the user device 106 or by the supervisor device 100 via the I/O subsystem 526 of the supervisor device 110 and can be provided to the server 102 via the communication subsystem(s) 532 of one or both of the user device 106 and the server 102, and communication network 120. After the user identification information has been received, the user associated with user identification information can be identified. In some embodiments, this can include comparing the received user identification information to information stored in the database server such as, for example, user profile database 301. In some embodiments, for example, user identification information can be uniquely associated

with user information stored in the user profile database 301 and the user can be identified by making the received user identification information with that user information.

[0220] At block 804 of the process 800, the user interface is launched. In some embodiments, the user interface can be configured to provide information to the user of the user device 106 and receive inputs from the user of the user device. The user interface can include one or several windows or fields in which content can be displayed and/or in which content can be entered. In some embodiments, the user interface can include one or several features such as, for example, toggles, buttons, switches, or the like through which the user of the user interface can input data. The launching of the user interface can include, for example, the generation and sending of a launch signal from the server 102 to the user device 106. In some embodiments, this launch signal can include one or several commands that can direct the user device 106, and particularly the I/O subsystem 526 of the user device 106 to launch a user interface when the launch signal is received.

[0221] At block 806 of the process 800, user metadata is received and/or retrieved. In some embodiments, the user metadata can be received and/or retrieved from the database server 104 and specifically from the user profile database 301. In some embodiments, the user metadata can identify one or several attributes of the user such as, for example, a user skill level, one or several learning styles, progress through content, location in a content network, location in a content progression, progress through a content progression, historic response data such as identification of one or several questions correctly or incorrectly answered, or the like. In some embodiments, all or portions of the user metadata can be based on a skill model, which can be dynamically adjusted by the model engine 682 in response to received responses. In some embodiments, the skill model can be based on a distribution such as, for example, a normal distribution and can be characterized by, for example, a center, a standard deviation, or the like.

[0222] After blocks 802-806, the process 800 proceeds to block 808 wherein content is selected and provided to the user device. In some embodiments, this can include the identification of the user location in the content network, the identification of one or several potential next nodes or potential next pieces of content, and the selection of one of those one or several potential next nodes or potential next pieces of content. In some embodiments, for example, the one or several next nodes or potential next pieces of content can be selected according to one or several guard conditions, according to one or several models such as an evidence model or skill model, or the like. After the content has been selected, the content can be provided to the user via the user device 106. In some embodiments, for example, the server 102 can provide the content to the user device 106 via the communication network 120. The user device 106 can then provide the content to the user via the user interface launched in block 804 and operating on the I/O subsystem 526 of the user device. In some embodiments, the content can be provided as part of the presentation process 670 by one or both of the presenter module 672 and the view module 674.

[0223] After the content has been selected and/or provided, the process 800 proceeds to block 810, wherein a response, also referred to herein as a response communication is received. In some embodiments, the response can be

received by the server 102 from the user device 106. Specifically, in some embodiments, the response can be communicated from the user device 106 to the server 102 via the communication network 120. In some embodiments, this communication can be part of the presentation process 670 and can include one or both of the presenter module 672 and the view module 674. The received response can be provided to the response processor 678 and, as part of the response process 670, the response can be evaluated as indicated in block 812 of FIG. 16. In some embodiments, the response can be evaluated to determine whether the user correctly or incorrectly responded to the content provided in block 808. In some embodiments, and as part of the evaluation process, an indicia of correctness can be associated with a response. In some embodiments, this indicia can indicate whether the received response is a correct response. In some embodiments, this indicia can include a first value indicative of correctness can be associated with a correct response and a second value indicative of incorrectness can be associated with an incorrect response.

[0224] After the response has been evaluated, the process proceeds to decision state 814, wherein it is determined if the received response was correct. In some embodiments this can include determining whether the indicia of correctness indicates that a correct response was received or whether the indicia of correctness indicates that an incorrect response was received. In some embodiments, this determination can be made by the server 102, and specifically by the response processor 678.

[0225] If it is determined that the response is incorrect, then the process 800 proceeds to block 816, wherein the user data, and specifically the user metadata is updated. In some embodiments, this can include updating the user data to reflect the incorrect response, a change in skill level caused by the incorrect response, an amount of lapsed time in responding, or the like. In some embodiments, the update of the user data can include updating one or several skill models associated with the user. This update can be performed by the server 102, and specifically by the model engine 682 as part of the summary model process 680.

[0226] After the user data has been updated, the process 800 proceeds to block 818, wherein all or portions of the user data are compared to an intervention threshold. In some embodiments, this comparison can include the retrieving of the intervention threshold from the database server 104, and specifically from the threshold database 310. In some embodiments, this intervention threshold can delineate between users designated for receipt of an intervention and users designated for nonreceipt of an intervention. In some embodiments, this threshold can specify a skill level such that when a user skill level drops below the threshold level, an intervention is indicated. In some embodiments, the intervention threshold can specify a rate of change, a slope, a derivative, or the like. In such an embodiment, when the student skill level is dropping faster than a threshold value, that student is indicated for intervention. The intervention threshold can be compared to the user data by the server 102.

[0227] After the intervention threshold has been compared to the user data, the process 800 proceeds to decision state 820, wherein it is determined whether to intervene. In some embodiments, this determination can be made based on the results of the comparison of the intervention threshold to the user data. This determination can be made by the server 102. In some embodiments, if a determination is made to inter-

vene, then an alert can be sent to the supervisor device **110**. In some embodiments, this alert can include data indicating the need for intervention, and in some embodiments, this alert can include data identifying a drop in the user skill level below a predetermined threshold value and/or a negative rate of change of the user skill level giving rise to the alert. In some embodiments, the alert can comprise code, to cause the supervisor device **110** to display all or portions of the alert content to the user upon receipt of the alert. In some embodiments, this can include code to cause the launch of a user interface on the supervisor device **110** to display all or portions of the alert content.

[0228] If it is determined to intervene, then the process **800** proceeds to block **822**, wherein the intervention is selected and delivered. In some embodiments, the intervention can be selected according to the user data and information associated with one or several potential interventions. In some embodiments, the intervention can be selected by the recommendation engine **686** and can be provided to the user via the user device **106**, and specifically via the I/O subsystem **526** of the user device **106**. In some embodiments, the intervention can include interaction with the AI agent.

[0229] After the intervention has been selected and delivered, the process **800** proceeds to block **824**, wherein the intervention is terminated. In some embodiments, the intervention can be terminated when the user has consumed all of the content associated with the intervention, has reached a desired skill level or mastery level, has provided one or several desired responses, or the like.

[0230] After the intervention has been terminated, or returning again to decision state **814**, if it is determined that the response is correct, then the process **800** proceeds to block **826**, wherein the user data, and specifically the user metadata is updated. In some embodiments, this can include updating the user data to reflect the correct response, a change in skill level caused by the correct response, an amount of lapsed time in responding, or the like, and in some embodiments, the user data can be updated to reflect the outcome of the intervention. In some embodiments, the update of the user data can include updating one or several skill models associated with the user. This update can be performed by the server **102**, and specifically by the model engine **682** as part of the summary model process **680**.

[0231] After the user data has been updated, or returning again to decision state **820**, if it is determined to not intervene, the process **800** proceeds to decision state **828**, wherein it is determined if there is additional content for delivery to the user. In some embodiments, this can include, for example, determining if all of a group of content has been provided to the user, if one or several time limits have been reached or surpassed, or the like. If it is determined that there is additional content, then the process **800** returns to block **808** and proceeds as outlined above.

[0232] If it is determined that there is no additional content, then the process **800** proceeds to block **830**, wherein a termination alert is generated and sent. In some embodiments, the termination alert can comprise data indicative of the user performance, the user skill level, a change in the user skill level, a lapsed time, or the like. In some embodiments, the termination alert can be sent to the user via the user device **106**, and in some embodiments, the termination alert can be sent to a supervisor such as a teacher via the supervisor device **110**. The termination alert can comprise

code, to cause the recipient device **106**, **110** to display all or portions of the alert content to the user upon receipt of the termination alert. In some embodiments, this can include code to cause the launch of a user interface on the user device **106** to display all or portions of the alert content.

[0233] With reference now to FIG. **17**, a flowchart illustrating one embodiment of a process **850** for selecting and delivering an intervention is shown. The process **850** can be performed as a part of, or in the place of block **822** of FIG. **16**. The process **850** can be performed by all or portions of the CDN **100**, and particularly by the server **102** and/or the AI agent **124**. The process **850** begins at block **852**, wherein the AI agent **124** is launched. In some embodiments, this can include the generating and sending of a launch signal from the server **102** to the AI agent **124**.

[0234] In some embodiments, the AI agent **124** can be launched in response to a request received from the user to launch the AI agent **124**, and in some embodiments, the AI agent **124** can be launched in response to the meeting and/or exceeding of one or several triggering conditions. In some embodiments the triggering condition can delineate between circumstances when an intervention is indicated, and particularly when an intervention with the AI agent **124** is indicated and situations in which an intervention is not indicated. The triggering conditions can detect and indicate for intervention based on a pattern of user interaction. The pattern of user interaction can arise from user interaction with content such as, for example, responses provided to content; interactions with the user device; interaction with an educational framework and/or a learning environment, or the like. These one or several triggering conditions can comprise, for example, one or several thresholds delineating between, for example: acceptable and unacceptable skill levels; acceptable and unacceptable responses; acceptable and unacceptable amounts of time on task; or the like. In some embodiments, the triggering conditions can comprise one or several threshold relating to user interactions with each other and/or with the supervisor, interaction with the CDN **100**, or the like.

[0235] After the AI agent is launched, the process **850** proceeds to block **854**, wherein a dialogue interface, also referred to herein as a dialogue user interface is launched. In some embodiments, the dialogue interface can comprise one or several window configured to display one or several text strings generated by the AI agent **124**. In some embodiments, these one or several windows can be further configured for input of one or several text strings from the user. In some embodiments, the dialogue interface can be launched within one or several windows in the user interface launched in block **804**.

[0236] After the dialogue interface has been launched, the process **850** proceeds to block **856**, wherein remediation dialogue is identified. The remediation dialogue can comprise content covering one or several topics. In some embodiments, this can include supplemental content which can be provided to the user and/or one or several prompts. In some embodiments, these one or several prompts can be questions the answers to which can be analyzed to determine: a skill level; to isolate difficulty, misconception, misunderstanding, knowledge gap, or the like; to determine engagement, or the like. In some embodiments, identification of the remediation dialogue can include identifying one or several topics for review via the intervention, and the retrieval and/or generation of text strings relevant to those

one or several topics for review. In some embodiments, the remediation dialogue can be identified, at least in part, based on the response received in block **810** that prompted the remediation.

[0237] After the remediation dialogue has been identified, the process **850** proceeds to decision state **858**, wherein it is determined if supplemental content is available and/or is part of the remediation dialogue. In some embodiments, this can include reviewing metadata of content that is included in the remediation dialogue to determine if any of the remediation dialogue is supplemental content. In some embodiments, the determination of decision state **858** can further include determining whether to provide supplemental content. This decision can, in some embodiments, be made by the AI agent **124** based on one or several inputs received from the user and/or the user skill level. In some embodiments in which the user skill level is very low, supplemental content can be provided to increase the user skill level. The AI agent **124** and/or the server **102** can make the determination of decision state **858**.

[0238] If it is determined that supplemental content is desired, then the process **850** proceeds to block **860**, wherein the supplemental content is identified. In some embodiments, the supplemental content can be identified by the AI agent **124** based on information received from the user, information retrieved about the user, and/or meta data associated with the intervention or components thereof. Specifically, in some embodiments, the supplemental content, also referred to herein as the delivered content, can be selected based on the user skill level as identified in the user metadata, and specifically in the skill model at the time of the selection of the delivered content. In some embodiments, the supplemental content can be further selected based on one or several requests received from the user such as, for example, a request for content explaining a topic, subject, or the like. In some embodiments, these requests can be part of the remediation dialogue. In some such embodiments, the supplemental content can be selected based on the received request and the user skill level.

[0239] After the supplemental content has been identified, the process **850** proceeds to block **862**, wherein the supplemental content is delivered. In some embodiments, this can include generating and/or sending a message from the server **102** and/or the AI agent **124** to the user device **106**. In some embodiments, this message can comprise data comprising the supplemental content. This message can be sent to the server **102** via the communication network **120**.

[0240] After the supplemental content has been delivered, or, returning again to decision state **858**, if it is determined that no supplemental content is desired, then the process **850** proceeds to block **864**, wherein one or several prompts are generated and/or delivered. In some embodiments, these prompts can comprise one or several questions. The one or several prompts can be delivered via an electronic communication between the user device **106** and the AI agent **124** and/or the server **102**. In some embodiments these electronic communication can comprise the one or several prompts.

[0241] After the one or several prompts have been generated and/or delivered, the process **850** proceeds to block **866**, wherein a response communication is received by the AI agent **124** and/or the server **102** from the user device **106**. In some embodiments, the response communication can comprise one or several text strings generated from user inputs to the user device **106**. In some embodiments these

one or several text strings can be generated via user inputs to the user device **106** via, for example, a keyboard, a microphone, or the like. In some embodiments, these text strings can be generated with one or several natural language processing algorithms including, for example, one or several speech recognition algorithms based on electronic signals generated, for example, by a microphone of the user device **106** from user speech.

[0242] After the response communication has been received, the process **850** proceeds to block **868**, wherein the received response is evaluated. In some embodiments, the received response can be evaluated by a component of the CDN **100** such as, for example, the server **102**, the AI agent **124**, and/or the response processor **678**. In some embodiments, the evaluation of the response can comprise the application of one or several natural language processing algorithms to one or several text strings comprising the response to determine the meaning of the response. These one or several algorithms can include, for example, one or several natural language understanding algorithms, one or several topic segmentation algorithms, or the like. In some embodiments, the evaluation of the response can include the natural language processing of the received response and the determination of understanding, misconceptions, skill level, or the like based on the output of the natural language processing.

[0243] After the response has been evaluated, the process **850** proceeds to block **870**, wherein a user performance metric is generated. In some embodiments, the user performance metric can comprise a user skill level, a change in the user skill level resulting from the intervention and specifically from the remediation dialogue, a duration of time spent in intervention, or the like. The user performance metric can be generated by a component of the CDN **100** such as, for example, the server **102**, the AI agent **124**, and/or the response processor **678**. In some embodiments the generated user performance metric can be stored in the user profile database **301**. Specifically, in some embodiments, the user performance metric can be used to update metadata associated with the user and/or one or several evidence models associated with the user such as a skill model.

[0244] After the user performance metric has been generated, the process **850** proceeds to block **872**, wherein the user performance metric is compared to a termination threshold. In some embodiments, the termination threshold can delineate between when to terminate the remediation dialogue and when to continue the remediation dialogue. In some embodiments, for example, the termination threshold can identify a time duration that when met or exceeded is associated with termination of the remediation dialogue. In some embodiments, the termination can identify a skill level, or a positive change in a skill level that when met or exceeded is associated with termination of the remediation dialogue. The termination threshold can be stored in, for example, the database server **104**, and specifically in the threshold database **310**.

[0245] After the termination threshold has been compared to the user performance metric, the process **850** proceeds to decision state **874**, wherein it is determined whether to terminate the remediation dialogue. In some embodiments, this determination is made based on the comparison of the termination threshold to the user performance metric. This determination can be made by a component of the CDN **100** such as, for example, the server **102**, the AI agent **124**,

and/or the response processor **678**. If it is determined not to terminate the remediation dialogue, then the process **850** returns to block **856** and continues as outlined above. If it is determined to terminate the remediation dialogue, then the process **850** continues to block **876** and proceeds to block **826** of FIG. **16**. In some embodiments, and as part of block **876**, an indication of the user performance metric such as, for example, an indication of the user skill level can be provided from, for example, the AI agent **124** to the server **102**.

[0246] In some embodiments, an process **850** can be terminated according to the meeting and/or exceeding of one or several halting criteria. The user termination threshold can be, in one embodiment, a halting criterion. In some embodiments, for example, the halting criterion can specify one or several thresholds for delineating between when to terminate the process **850** and when to continue the process **850**. In some embodiments, these thresholds can relate to, for example, the amount of time spent on the process **850**, a student rate of progress based on the process **850**, a student skill level, or the like. In some embodiments, one or several components of the CDN **100** can compare one or several user attributes to one or several halting criteria, and specifically to the one or several thresholds associated with those one or several halting criteria. If termination of the process **850** is indicated based on the comparison, then the process **850** can move to block **876**. If termination of the intervention is not indicated, then the process **850** returns to block **856** until termination of the intervention is indicated.

[0247] In some embodiments, the process **850** can include multiple rounds of interrogation, including providing a prompt and/or supplemental content, and response. In some embodiments, each of these round can include looping through some or all of steps **856-874** of FIG. **17**. As depicted in FIG. **17**, the user performance metric can be generated as a part of that loops, and specifically can be generated after the providing of a prompt and the receipt of a response to that prompt. Thus, in some embodiments, user skill level can be calculated subsequent to some or all rounds of interrogation and response.

[0248] With reference now to FIGS. **18** and **19**, flowcharts illustrating one embodiment of a process **900** for content selection and delivery are shown. The process **900** can be performed by all or portions of the CDN **100** including, for example, the server **102**, the database server **104**, the user device **106**, the supervisor device **100**, and/or the AI agent **124**. The process **900** begins at block **902**, wherein user identification information, also referred to herein as a user identifier, is received. In some embodiments, user identification information can include a username, password, login information, a unique user identifier, or the like. In some embodiments, the user identification information can be received by the user device **106** via the I/O subsystem **526** of the user device **106** or by the supervisor device **100** via the I/O subsystem **526** of the supervisor device **110** and can be provided to the server **102** via the communication subsystem(s) **532** of one or both of the user device **106** and the server **102**, and communication network **120**.

[0249] After the user identification information has been received, the process **900** proceeds to block **904**, wherein the user history, or in other words, user metadata is received and/or retrieved. In some embodiments, the user metadata can be received and/or retrieved from the database server **104** and specifically from the user profile database **301**. In

some embodiments, the user metadata can identify one or several attributes of the user such as, for example, a user skill level, one or several learning styles, progress through content, location in a content network, location in a content progression, progress through a content progression, historic response data such as identification of one or several questions correctly or incorrectly answered, or the like. In some embodiments, all or portions of the user metadata can be based on a skill model, which can be dynamically adjusted by the model engine **682** in response to received responses. In some embodiments, the skill model can be based on a distribution such as, for example, a normal distribution and can be characterized by, for example, a center, a standard deviation, or the like.

[0250] After the user metadata has been received, the process **900** proceeds to block **906**, wherein the user's location in the content network and/or the content progression is determined. In some embodiments, this location can be determined based on the received user metadata. Specifically, in some embodiments, the user metadata can identify one or several content networks and/or content progressions in which the user is located and can identify one or several last nodes completed by the user, current nodes at which the user is located, and/or next uncompleted nodes. Based on this information in the metadata, the server **102** or other component of the CDN **100** can determine the location of the user.

[0251] After the user's location has been determined, the process **900** proceeds to block **908**, wherein a node is identified for delivery to the user. In some embodiments, this node can be the node at which the user is currently located and/or the next node in the content network identified by the user's location in the content network. After the node for delivery has been identified, the process **900** proceeds to block **916** wherein one or several future potential next nodes are identified. In some embodiments, this can include identifying potential next nodes for some or all of the outcomes of the node for delivery identified in block **908**. Specifically, in some embodiments, one or several future potential next nodes can be identified for the event that a correct response is provided to the node identified in block **908** and one or several future potential next nodes can be identified for the event that an incorrect response is provided to the node identified in block **908**. In some embodiments, the delivery node can comprise a placeholder node, and in some embodiments, the delivery node is not a placeholder node. In such embodiments, a placeholder node is a node having content adaptively selected from a plurality of selectable content options, and a non-placeholder node is directly and/or statically linked to a single piece of content. In some embodiments in which the delivery node comprises a placeholder node, identification of the delivery node can further include identification of the content for delivery with the placeholder node. In some embodiments, the identification of the content for delivery with the placeholder node can include the selection of content according to metadata associated with the user and/or according to metadata associated with the content of the placeholder node. In some embodiments, the content of the placeholder node can comprise an intervention such as, for example, a remediation dialogue.

[0252] After the future potential next nodes have been identified, the process **900** proceeds to block **918**, wherein the gate conditions associated with those future potential next nodes are identified and/or retrieved. In some embodi-

ments, these gate conditions can be retrieved from the database server **104**, and specifically from the content library database **303**.

[0253] After the gate conditions have been retrieved, the process **900** proceeds to block **922** and continues with block **924** of FIG. **19**, wherein node data is delivered. In some embodiments, the node data can include the content of the delivery node, the identified future potential next nodes, and the gate conditions associated with those future potential next nodes. In some embodiments, the node data can be delivered from the server **102** to the user device **106**. The node data can be delivered to the user device **106** in one communication or in several communications.

[0254] In some embodiments, for example, the node data can be delivered to the user device in a first communication containing the delivery node and a second communication containing the potential next nodes and the gate conditions for the potential next nodes. In some embodiments, these communications can be timed to maximize efficiency of the content distribution network **100**. In one embodiment, for example, the first communication can be delivered to the user device and the delivery node contained in the first communication can be provided to the user simultaneous with the delivery of the second communication containing the potential next nodes and the gate conditions for the potential next nodes.

[0255] In some embodiments, and upon receipt of the node data, the user device **106** can provide the delivery node identified in block **908** to the user via the I/O subsystem **526** as indicated in block **926**. In embodiments in which the content of the node comprises an intervention or a remediation dialogue, the user device **106** can deliver the intervention according to steps **822**, **824** of FIG. **16**, and/or according to process **850** of FIG. **17**.

[0256] The user device can then receive a response from the user and can, in some embodiments, evaluate the received response as indicated in block **928**. The user device can then identify a next node from the potential next nodes received with the node data based on the gate conditions also received with the node data. The user device **106** can provide the next node identified from the received potential next nodes to the user and can generate and send an update communication. In some embodiments, the update communication is sent to the server **102** from the user device **106**, and the update communication identifies the response, the response evaluation, and the provided next node.

[0257] At block **930**, the server **102** receives the update communication from the user device **106**. The process then proceeds to block **932**, wherein the user metadata and/or the user model is updated and stored. In some embodiments, this update can be performed by the server **102**, and specifically by the model engine **682**. The updated user metadata and/or user model can be stored in the database server **104**, and specifically in the user profile database **301**.

[0258] After the user metadata and/or model have been updated, the process **900** proceeds to block **936**, wherein it is determined whether to intervene. In some embodiments, this can include comparing the user metadata and/or user model to an intervention threshold. More specifically, in some embodiments, the intervention determination can be made according to blocks **818**, **820** of FIG. **16**. In some embodiments,

[0259] After the intervention determination has been made, the process **900** proceeds to blocks **938**, **940** wherein

an intervention is selected, provided, and terminated. In some embodiments, this can include the steps **852-874** of process **850** shown in FIG. **17**. With specific reference to block **938**, an intervention is selected and provided at this step. In some embodiments, the intervention can be selected according to the user data and information associated with one or several potential interventions. In some embodiments, the intervention can be selected by the recommendation engine **686** and can be provided to the user via the user device **106**, and specifically via the I/O subsystem **526** of the user device **106**. In some embodiments, the intervention can include one or several of: interaction with the AI agent **124**; targeted remediation content; scaffolding hints; question generation; answer specific feedback; text and/or text summaries; follow-up questions, or the like.

[0260] After the intervention has been selected and delivered, the process **900** proceeds to block **940**, wherein the intervention is terminated. In some embodiments, the intervention can be terminated when the user has consumed all of the content associated with the intervention, has reached a desired skill level or mastery level, has provided one or several desired responses, or the like. In some embodiments, an intervention can be terminated according to the meeting and/or exceeding of one or several halting criteria. In some embodiments, for example, each intervention can have a halting criterion. In some embodiments, for example, a halting criterion can specify one or several thresholds for delineating between when to terminate an intervention and when to continue an intervention. In some embodiments, these thresholds can relate to, for example, the amount of time spent on the intervention, a student rate of progress based on the intervention, a student skill level, or the like. In some embodiments, one or several components of the CDN **100** can compare one or several user attributes to one or several halting criteria, and specifically to the one or several thresholds associated with those one or several halting criteria. If termination of the intervention is indicated based on the comparison, then the intervention can be terminated. If termination of the intervention is not indicated, then the process **900** returns to block **938** until termination of the intervention is indicated.

[0261] After the intervention has been terminated, or returning again to decision state **936**, if it is determined to not intervene, the process **900** proceeds to block **942** and returns to block **916** of FIG. **18**. In such an embodiment future potential next nodes identified in block **916** are the future potential next nodes of the next node provided as part of block **928**.

[0262] A number of variations and modifications of the disclosed embodiments can also be used. Specific details are given in the above description to provide a thorough understanding of the embodiments. However, it is understood that the embodiments may be practiced without these specific details. For example, well-known circuits, processes, algorithms, structures, and techniques may be shown without unnecessary detail in order to avoid obscuring the embodiments.

[0263] Implementation of the techniques, blocks, steps and means described above may be done in various ways. For example, these techniques, blocks, steps and means may be implemented in hardware, software, or a combination thereof. For a hardware implementation, the processing units may 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.

[0264] Also, it is noted that the embodiments may be described as a process which is depicted as a flowchart, a flow diagram, a swim diagram, a data flow diagram, a structure diagram, or a block diagram. Although a depiction may describe the operations as a sequential process, many of the operations can be performed in parallel or concurrently. In addition, the order of the operations may be re-arranged. A process is terminated when its operations are completed, but could have additional steps not included in the figure. A process may correspond to a method, a function, a procedure, a subroutine, a subprogram, etc. When a process corresponds to a function, its termination corresponds to a return of the function to the calling function or the main function.

[0265] Furthermore, embodiments may be implemented by hardware, software, scripting languages, firmware, middleware, microcode, hardware description languages, and/or any combination thereof. When implemented in software, firmware, middleware, scripting language, and/or microcode, the program code or code segments to perform the necessary tasks may be stored in a machine readable medium such as a storage medium. A code segment or machine-executable instruction may 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, data structures, and/or program statements. A code segment may be coupled to another code segment or a hardware circuit by passing and/or receiving information, data, arguments, parameters, and/or memory contents. Information, arguments, parameters, data, etc. may be passed, forwarded, or transmitted via any suitable means including memory sharing, message passing, token passing, network transmission, etc.

[0266] For a firmware and/or software implementation, the methodologies may be implemented with modules (e.g., procedures, functions, and so on) that perform the functions described herein. Any machine-readable medium tangibly embodying instructions may be used in implementing the methodologies described herein. For example, software codes may be stored in a memory. Memory may be implemented within the processor or external to the processor. As used herein, 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.

[0267] Moreover, as disclosed herein, the term “storage medium” may 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, and/or various other storage mediums capable of storing that contain or carry instruction(s) and/or data.

[0268] 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 recommendation system comprising:
memory comprising:

a user profile database comprising information relating to a plurality of users, wherein the information relating to the plurality of users comprises unique user history data associated with each of the users in the plurality of users; and

a content library database comprising a plurality of nodes linked in a plurality of sequential relationships in a content network, wherein some of the plurality of nodes are associated with guard conditions, wherein each guard condition identifies at least one of: an entrance condition controlling entrance into the node associated with the guard condition; and an exit condition controlling exit from the node associated with the guard condition;

a user device comprising:

a network interface configured to exchange data via the communication network; and

an I/O subsystem configured to convert electrical signals to user interpretable outputs via a user interface; and

one or more servers configured to:

receive user identification information, wherein the user identification information identifies a user;

determine a location of the user identified by the received user identification information in the content network;

select a next node based on the location of the user in the content network and on at least some of the guard conditions;

deliver next node content to the user device via a communication network, wherein the next node content is adaptively selected from a plurality of selectable content options when the next node is a placeholder node, and wherein the next content is directly linked to the next node when the next node is not a placeholder node.

2. The system of claim 1, wherein at least some of the guard conditions regulate entry into a learning sequence comprising some of the plurality of nodes.

3. The system of claim 1, wherein at least some of the guard conditions regulate exit from a learning sequence comprising some of the plurality of nodes.

4. The system of claim 1, further comprising an artificial intelligence engine configured to receive inputs from at least the user device, and wherein at least one of the selectable content options of the placeholder node comprises a dialogue with the artificial intelligence engine.

5. The system of claim 4, wherein the artificial intelligence engine is configured to:

launch the dialogue;

deliver interrogation and content based on the received user responses; and

determine termination of the dialogue based on a user performance metric indicative of a user skill level, wherein the dialogue is determined for termination when the user performance metric identifies a sufficient user skill level.

6. The system of claim 5, wherein the remediation dialogue comprises multiple levels of interrogation and response.

7. The system of claim 6, wherein the artificial intelligence engine is configured to reevaluate the user skill level subsequent to each level of interrogation and response.

8. The system of claim 7, wherein the reevaluation of the user skill level comprises application of a natural language processing algorithm to the received response, wherein the natural language processing algorithm comprises a speech recognition algorithm and natural language understanding algorithm.

9. The system of claim 1, wherein the one or several servers are further configured to determine a user skill level based on user responses received relating to the delivered next node content.

10. The system of claim 1, further comprising a supervisor device in communicating connection with the one or several servers via the communication network, wherein the one or several servers are further configured to send an alert to the supervisor device when the determined user skill level drops below a predetermined threshold value, wherein the alert comprises computer code configured to direct the supervisor device to display the alert upon receipt.

11. A method of content selection and delivery, the method comprising:

receiving user identification information at one or several servers from a user device via a communication network, wherein the user identification information identifies a user;

determining with the one or several servers a location of the user identified by the received user identification information in the content network;

selecting with the one or several servers a next node based on the location of the user in the content network and on at least some of the guard conditions; and

delivering with the one or several servers next node content to the user device via the communication network, wherein the next node content is adaptively selected from a plurality of selectable content options when the next node is a placeholder node, and wherein the next content is directly linked to the next node when the next node is not a placeholder node.

12. The method of claim 11, wherein at least some of the guard conditions regulate entry into a learning sequence comprising some of the plurality of nodes.

13. The method of claim 11, wherein at least some of the guard conditions regulate exit from a learning sequence comprising some of the plurality of nodes.

14. The method of claim 11, wherein the next content comprises an intervention.

15. The method of claim 14, wherein the intervention comprises at least one of: targeted remediation content; a

scaffolding hint; question generation; answer specific feedback; a text; a text summary; and a follow-up question.

16. The method of claim 11, wherein the next content directly linked to the next node comprises a dialogue with an artificial intelligence engine, wherein the artificial intelligence engine is configured to receive inputs from at least the user device.

17. The method of claim 11, wherein at least one of the selectable content options of the placeholder node comprises a dialogue with an artificial intelligence engine, wherein the artificial intelligence engine is configured to receive inputs from at least the user device.

18. The method of claim 17, further comprising: launching the dialogue with the artificial intelligence engine;

delivering interrogation and content based on the received user responses from the artificial intelligence engine to the user device; and

terminating the dialogue with the artificial intelligence engine based on at least one halting criterion.

19. The method of claim 18, wherein the halting criteria comprises a termination threshold, and wherein terminating the dialogue comprises: generating a user performance metric indicative of at least one of: a skill level; a correct response; and an incorrect response; and comparing the user performance metric to the termination threshold.

20. The method of claim 19, wherein the dialogue is terminated when the user performance metric identifies a sufficient user skill level.

21. The method of claim 18, wherein the remediation dialogue comprises multiple levels of interrogation and response.

22. The method of claim 21, further comprising reevaluating the user skill level with the artificial intelligence engine subsequent to each level of interrogation and response.

23. The method of claim 22, wherein the reevaluation of the user skill level comprises application by the one or several servers of a natural language processing algorithm to the received response, wherein the natural language processing algorithm comprises a speech recognition algorithm and natural language understanding algorithm.

24. The method of claim 11, further comprising determining a user skill level with the one or several servers based on user responses received relating to the delivered next node content.

25. The method of claim 11, further comprising sending an alert from the one or several servers to a supervisor device when the determined user skill level drops below a predetermined threshold value, wherein the alert comprises computer code directing the supervisor device to display the alert upon receipt.

* * * * *