



(51) International Patent Classification:

H04L 51/02 (2022.01) *H04L 51/214* (2022.01)
H04L 51/04 (2022.01)

(21) International Application Number:

PCT/US2024/014450

(22) International Filing Date:

05 February 2024 (05.02.2024)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

63/483.528	06 February 2023 (06.02.2023)	US
------------	-------------------------------	----

(71) Applicant: SNAP INC. [US/US]; 3000 31st Street, Santa

Applicant: SYNN INC. [US, US], 5000 51st Street, Santa Monica, California 90405 (US).

(72) Inventors: MULLIGAN, William Spencer; 3000 31st

Street, Santa Monica, California 90405 (US). **YOUNG, Samuel**; 3000 31st Street, Santa Monica, California 90405 (US). **BREWER, Jason**; 3000 31st Street, Santa Monica, California 90405 (US). **TODI, Anirudh**; 3000 31st Street, Santa Monica, California 90405 (US). **CASLIN, David**

Clark; 3000 31st Street, Santa Monica, California 90405 (US).

(74) **Agent: PERDOK, Monique M.** et al.; P.O. Box 2938, Minneapolis, Minnesota 55402 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM,

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

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, CV,

GH, GM, KE, LR, LS, MW, MZ, NA, RW, SC, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ,

(54) Title: CHATBOT FOR INTERACTIVE PLATFORMS




FIG. 3B

(S7) Abstract: A chatbot system for filtering conversation content. A chatbot system receives, from a client system, a prompt of a user during an interactive session. The chatbot system filters the prompt of the user based on a set of platform policies and generates a response based on the filtering of the prompt of the user, and communicates the response to the client system.

DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT,
LU, LV, MC, ME, MK, MT, NL, NO, PL, PT, RO, RS, SE,
SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN,
GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:

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

CHATBOT FOR INTERACTIVE PLATFORMS

PRIORITY CLAIM

[0001] This application claims the benefit of U.S. Provisional Application No. 63/483,403, filed on February 6, 2023, which is hereby incorporated by reference in its entirety.

TECHNICAL FIELD

[0002] The present disclosure relates generally to interactive platforms and more particularly to providing user interfaces to users of an interactive platform.

BACKGROUND

[0003] Users enjoy accessing interactive platforms to share content with other users of the interactive platform. In addition, users often share sensitive content when communicating with a chatbot.

BRIEF DESCRIPTION OF THE SEVERAL VIEWS OF THE DRAWINGS

[0004] In the drawings, which are not necessarily drawn to scale, like numerals may describe similar components in different views. To easily identify the discussion of any particular element or act, the most significant digit or digits in a reference number refer to the figure number in which that element is first introduced. Some non-limiting examples are illustrated in the figures of the accompanying drawings in which:

[0005] FIG. 1 is a diagrammatic representation of a networked environment in which the present disclosure may be deployed, according to some examples.

[0006] FIG. 2 is a diagrammatic representation of a messaging system, according to some examples, that has both client-side and server-side functionality.

[0007] FIG. 3A is an illustration of an interactive session by a user with a chatbot system, according to some examples.

[0008] FIG. 3B is a block diagram of the chatbot system, according to some examples.

[0009] FIG. 3C is a process flow diagram of a chat and advertising process of the chatbot system, according to some examples.

[0010] FIG. 3D is a process flow diagram of an advertising creation process, according to some examples.

[0011] FIG. 4A is a tech block diagram, according to some examples.

[0012] FIG. 4B is an activity diagram of a method of a chatbot system, according to some examples.

[0013] FIG. 5A and FIG. 5B are illustrations of a user interface, according to some examples.

[0014] FIG. 6 is a collaboration diagram of a sensitive content filtering system, according to some examples.

[0015] FIG. 7A is a collaboration diagram of a messaging and conversation system, according to some examples.

[0016] FIG. 7B is network diagram of a chatbot communication network, according to some examples.

[0017] FIG. 7C is a collaboration diagram of a communication network of a chatbot system, according to some examples.

[0018] FIG. 7D is an activity diagram of a communication method of a chatbot system, according to some examples.

[0019] FIG. 8 illustrates a machine-learning pipeline, according to some examples.

[0020] FIG. 9 illustrates training and use of a machine-learning program, according to some examples.

[0021] FIG. 10 is a diagrammatic representation of a data structure as maintained in a database, according to some examples.

[0022] FIG. 11 is a diagrammatic representation of a message, according to some examples.

[0023] FIG. 12 is a flowchart for an access-limiting process, according to some examples.

[0024] FIG. 13 is a diagrammatic representation of a machine in the form of a computer system within which a set of instructions may be executed to cause the machine to perform any one or more of the methodologies discussed herein, according to some examples.

[0025] FIG. 14 is a block diagram showing a software architecture within which examples may be implemented.

DETAILED DESCRIPTION

[0026] Interactive platforms (e.g., social platforms, social media platforms AR platforms, applications, messaging platforms, AR applications, operating systems, gaming systems or applications, systems with which a user interacts, and the like) may provide a way for users to interact with other users with similar interests. As interactive platforms may not charge for access to such platforms, an owner of an interactive platform may decide to serve advertising content to users based on interests of the users. Both users and owners of interactive platforms desire to see advertisements (“ads”) that are more relevant to a user's intent, which can make the ads more engaging and less disruptive to the user's experience. In addition, advertisers desire a way to easily provide engaging advertising content.

[0027] Various examples provide improved user intent detection during conversations with chatbots of an interactive platform. A chatbot, in some examples, is a software application that is designed to simulate human conversation through voice commands or text chats. A chatbot may employ Natural Language Processing (NLP) and Machine Learning (ML)/Artificial Intelligence (AI) methodologies to understand and interpret a user's input and generate a response.

[0028] In some examples, improved user intent detection allows advertisers to bid and target their ads to specific user segments based on a user's intent and interests, which can increase the chances that the user will engage with the ad. Advertisers can have faster ramp up time by providing targeted keywords. In addition, advertisers can benefit from automated creative generation which is based on matching user intent to advertisers' targeted user intent.

[0029] In some examples, a chatbot system provides user intent detection that improves targeting and optimization capabilities over time by analyzing data on user intent and conversions. This enhances the user experience and improves the relevance and performance of ads. Additionally, the interactive platform uses the extracted user intent to enhance the user experience across other portions of the interactive platforming site, making them more personalized and relevant to the user community. In some examples, an interactive platform enhances display advertising by targeting users based on their genuine intent ascertained, in whole or in part, through interaction with a chatbot. By extracting high intent and timely relevant keywords and concepts of conversation with the chatbot, the interactive platform may improve a user intent profile.

[0030] In some examples, advertisers target and bid for specific keywords or expanded concepts, which enables them to reach a highly targeted group of users.

[0031] Some platforms currently rely on traditional methods of ad targeting, such as demographic information and browsing and engagement history, which have limitations in terms of targeting users based on their direct and recent intent. With recent changes in privacy regulations, it has become increasingly difficult to obtain user intent information through third-party channels as well. Examples of this disclosure address these limitations by utilizing Large Language Models (LLMs) or AI systems and context extraction from conversations with a chatbot to target users with relevant ads. By using the context of conversations with the chatbot, an interactive platform can provide a more accurate and up-to-date representation of user intent and enable more effective ad targeting without endangering the privacy of users.

[0032] In some examples, users may interact with a chatbot in several different ways. Users can chat with the chatbot directly, in this scenario the whole conversation is utilized towards building the user intent. This also allows the chatbot system to extract accurate information about a user intent, as the user will be expressing their interests and needs directly to the chatbot. In some examples, users can be in a 1:1 chat with a friend and add the chatbot into the conversation using a tag, for example @mention. In this case, the prompt after the @mention can be utilized towards the user intent or the whole conversation may be used after the chatbot is added into the user intent extraction. In this manner, the platform can extract information about the user's intent from the context of the conversation and not only the prompt. In a group chat, the user can add the chatbot to the conversation. And the data extraction scenarios could be similar to 1:1 chats.

[0033] In some examples, an intent extraction pipeline maps conversations to an intent vector of important and intentful keywords and expanded concepts. This intent vector is a dynamic representation of a user's intent, which is updated in close to real-time as the user engages in conversations with the chatbot. So, the user's intent will be defined by a combination of demographics, engagement embeddings in the system and the keywords and concepts vector.

[0034] In some examples, the intent vector is a rich representation of a user's intent and the similarity of two vectors shows how closely two users' intents match. This can be used for “look-a-like” targeting and audience expansion, allowing the interactive

platform to quickly identify and target users who are similar to those who have already engaged with ads.

[0035] In some examples, textual components of ads are used to finetune the responses of an LLM, thus conditioning the responses of the large language model. This improves the suggestions by the chatbot system and improves user engagement.

[0036] In some examples, advertisers can use the concept vector to choose the target audience for their ads. They can bid on specific keywords and expanded concepts which allows them to target users who are likely to be interested in their product or service. Newer and smaller advertisers can also find relevant audiences by specifying a few key keywords, and the chatbot system will match that to a broader group of users.

[0037] In some examples, the chatbot system has detailed control over the context that is extracted from the conversations, so the list of possible keywords and concepts can be curated and adjusted as the platform evolves. The chatbot system can also learn new concepts on a daily basis and the list of acceptable expanded concepts can be moderated and curated to ensure they are relevant and appropriate for the platform. Overall, this platform uses the representation of user intent to provide a more personalized and engaging experience for users and more targeting options for advertisers. The intent vector can be used for many purposes, such as audience expansion, look-a-like targeting, and fast cold-start ramp-up.

[0038] In some examples, the chatbot system can be used in conjunction with many organic experiences such as content, search, augmented reality and engagement with mapping applications. This allows the interactive platform as a whole to be more efficient in targeting the right audience and improve the performance of the ads and organic experiences.

[0039] In some examples, the chatbot system preserves user privacy, as the chatbot system only uses a conversation when the chatbot is opted into the conversation and expanded concepts are from a curated list of potential bidding concepts.

[0040] In some examples, the chatbot system is a component of an interactive platform, a social platform, social media platform, a social network, an AR platform, an application, a messaging application, an AR application, an operating system, a gaming system or application, or any other system, media, or application with which a user can interact with.

[0041] In some examples, a user may communicate a user prompt to the chatbot system that contains sensitive content. The chatbot system receives, from a client system, a prompt of a user during an interactive session. The chatbot system filters the prompt of the user based on a set of platform policies, generates a response based on the filtering of the prompt of the user, and communicates the response to the client system.

[0042] In some examples, a platform chatbot messaging system receives, from a client system, a prompt of a user during an interactive session, filters the prompt of the user based on a set of platform policies, generates a response based on the filtering of the prompt of the user, and communicates the response to the client system.

[0043] In some examples, the platform chatbot messaging system filters the user prompt for keywords based on a set of platform policies.

[0044] In some examples, the keywords are keywords of sensitive content.

[0045] In some examples, the platform chatbot messaging system generates a filtered prompt based on the user prompt and a set of platform policies; communicates the filtered prompt to a chatbot component; and receives the response from the chatbot component, where the chatbot component generates the response based on the filtered prompt.

[0046] In some examples, the platform chatbot messaging system determines a set of platform resources based on the prompt of the user and generates the response based on the set of platform resources.

[0047] In some examples, the platform chatbot messaging system determines that the response comprises sensitive content, and in response, generates an appropriate response that does not comprise the sensitive content.

[0048] In some examples, the platform chatbot messaging system detects that the sensitive content comprises harmful conduct, and in response, redirects the user to seek appropriate help.

NETWORKED COMPUTING ENVIRONMENT

[0049] FIG. 1 is a block diagram showing an example interaction system 100 of an interactive platform for facilitating interactions (e.g., exchanging text messages, conducting text audio and video calls, or playing games) over a network. The interaction system 100 includes multiple client systems 102, each of which hosts multiple applications, including an interaction client 104 and other applications 106. Each

interaction client 104 is communicatively coupled, via one or more communication networks including a network 108 (e.g., the Internet), to other instances of the interaction client 104 (e.g., hosted on respective other client systems 102), an interaction server system 110 and third-party servers 112). An interaction client 104 can also communicate with locally hosted applications 106 using Applications Programming Interfaces (APIs).

[0050] Each client system 102 may include multiple user devices, such as a mobile device 114, head-wearable apparatus 116, and a computer client device 118 that are communicatively connected to exchange data and messages.

[0051] An interaction client 104 interacts with other interaction clients 104 and with the interaction server system 110 via the network 108. The data exchanged between the interaction clients 104 (e.g., interactions 120) and between the interaction clients 104 and the interaction server system 110 includes functions (e.g., commands to invoke functions) and payload data (e.g., text, audio, video, or other multimedia data).

[0052] The interaction server system 110 provides server-side functionality via the network 108 to the interaction clients 104. While certain functions of the interaction system 100 are described herein as being performed by either an interaction client 104 or by the interaction server system 110, the location of certain functionality either within the interaction client 104 or the interaction server system 110 may be a design choice. For example, it may be technically preferable to initially deploy particular technology and functionality within the interaction server system 110 but to later migrate this technology and functionality to the interaction client 104 where a client system 102 has sufficient processing capacity.

[0053] The interaction server system 110 supports various services and operations that are provided to the interaction clients 104. Such operations include transmitting data to, receiving data from, and processing data generated by the interaction clients 104. This data may include message content, client device information, geolocation information, media augmentation and overlays, message content persistence conditions, interactive platform information, and live event information. Data exchanges within the interaction system 100 are invoked and controlled through functions available via user interfaces (UIs) of the interaction clients 104.

[0054] Turning now specifically to the interaction server system 110, an API server 122 is coupled to and provides programmatic interfaces to interaction servers 124, making the functions of the interaction servers 124 accessible to interaction clients 104, other

applications 106 and third-party server 112. The interaction servers 124 are communicatively coupled to a database server 126, facilitating access to a database 128 that stores data associated with interactions processed by the interaction servers 124. Similarly, a web server 130 is coupled to the interaction servers 124 and provides web-based interfaces to the interaction servers 124. To this end, the web server 130 processes incoming network requests over the Hypertext Transfer Protocol (HTTP) and several other related protocols.

[0055] The API server 122 receives and transmits interaction data (e.g., commands and message payloads) between the interaction servers 124 and the client systems 102 (and, for example, interaction clients 104 and other application 106) and the third-party server 112. Specifically, the API server 122 provides a set of interfaces (e.g., routines and protocols) that can be called or queried by the interaction client 104 and other applications 106 to invoke functionality of the interaction servers 124. The API server 122 exposes various functions supported by the interaction servers 124, including account registration; login functionality; the sending of interaction data, via the interaction servers 124, from a particular interaction client 104 to another interaction client 104; the communication of media files (e.g., images or video) from an interaction client 104 to the interaction servers 124; the settings of a collection of media data (e.g., a story); the retrieval of a list of friends of a user of a client system 102; the retrieval of messages and content; the addition and deletion of entities (e.g., friends) to an entity graph (e.g., a social graph); the location of friends within a social graph; and opening an application event (e.g., relating to the interaction client 104).

[0056] The interaction servers 124 host multiple systems and subsystems, described below with reference to FIG. 2.

LINKED APPLICATIONS

[0057] Returning to the interaction client 104, features and functions of an external resource (e.g., a linked application 106 or applet) are made available to a user via an interface of the interaction client 104. In this context, “external” refers to the fact that the application 106 or applet is external to the interaction client 104. The external resource is often provided by a third party but may also be provided by the creator or provider of the interaction client 104. The interaction client 104 receives a user selection of an option to launch or access features of such an external resource. The external resource may be the application 106 installed on the client system 102 (e.g., a “native

app”), or a small-scale version of the application (e.g., an “applet”) that is hosted on the client system 102 or remote of the client system 102 (e.g., on third-party servers 112). The small-scale version of the application includes a subset of features and functions of the application (e.g., the full-scale, native version of the application) and is implemented using a markup-language document. In some examples, the small-scale version of the application (e.g., an “applet”) is a web-based, markup-language version of the application and is embedded in the interaction client 104. In addition to using markup-language documents (e.g., a *.ml file), an applet may incorporate a scripting language (e.g., a *.js file or a .json file) and a style sheet (e.g., a *.ss file).

[0058] In response to receiving a user selection of the option to launch or access features of the external resource, the interaction client 104 determines whether the selected external resource is a web-based external resource or a locally-installed application 106. In some cases, applications 106 that are locally installed on the client system 102 can be launched independently of and separately from the interaction client 104, such as by selecting an icon corresponding to the application 106 on a home screen of the client system 102. Small-scale versions of such applications can be launched or accessed via the interaction client 104 and, in some examples, no or limited portions of the small-scale application can be accessed outside of the interaction client 104. The small-scale application can be launched by the interaction client 104 receiving, from a third-party server 112 for example, a markup-language document associated with the small-scale application and processing such a document.

[0059] In response to determining that the external resource is a locally-installed application 106, the interaction client 104 instructs the client system 102 to launch the external resource by executing locally-stored code corresponding to the external resource. In response to determining that the external resource is a web-based resource, the interaction client 104 communicates with the third-party servers 112 (for example) to obtain a markup-language document corresponding to the selected external resource. The interaction client 104 then processes the obtained markup-language document to present the web-based external resource within a user interface of the interaction client 104.

[0060] The interaction client 104 can notify a user of the client system 102, or other users related to such a user (e.g., “friends”), of activity taking place in one or more external resources. For example, the interaction client 104 can provide participants in a conversation (e.g., a chat session) in the interaction client 104 with notifications relating

to the current or recent use of an external resource by one or more members of a group of users. One or more users can be invited to join in an active external resource or to launch a recently-used but currently inactive (in the group of friends) external resource. The external resource can provide participants in a conversation, each using respective interaction clients 104, with the ability to share an item, status, state, or location in an external resource in a chat session with one or more members of a group of users. The shared item may be an interactive chat card with which members of the chat can interact, for example, to launch the corresponding external resource, view specific information within the external resource, or take the member of the chat to a specific location or state within the external resource. Within a given external resource, response messages can be sent to users on the interaction client 104. The external resource can selectively include different media items in the responses, based on a current context of the external resource.

[0061] The interaction client 104 can present a list of the available external resources (e.g., applications 106 or applets) to a user to launch or access a given external resource. This list can be presented in a context-sensitive menu. For example, the icons representing different ones of the application 106 (or applets) can vary based on how the menu is launched by the user (e.g., from a conversation interface or from a non-conversation interface).

SYSTEM ARCHITECTURE

[0062] FIG. 2 is a block diagram illustrating further details regarding the interaction system 100, according to some examples. Specifically, the interaction system 100 is shown to comprise the interaction client 104 and the interaction servers 124. The interaction system 100 embodies multiple subsystems, which are supported on the client-side by the interaction client 104 and on the server-side by the interaction servers 124. Example subsystems are discussed below.

[0063] An image processing system 202 provides various functions that enable a user to capture and augment (e.g., augment or otherwise modify or edit) media content associated with a message.

[0064] A camera system 204 includes control software (e.g., in a camera application) that interacts with and controls hardware camera hardware (e.g., directly or via operating system controls) of the client system 102 to modify and augment real-time images captured and displayed via the interaction client 104.

[0065] The augmentation system 206 provides functions related to the generation and publishing of augmentations (e.g., media overlays) for images captured in real-time by cameras of the client system 102 or retrieved from memory of the client system 102. For example, the augmentation system 206 operatively selects, presents, and displays media overlays (e.g., an image filter or an image lens) to the interaction client 104 for the augmentation of real-time images received via the camera system 204 or stored images retrieved from memory of a client system 102. These augmentations are selected by the augmentation system 206 and presented to a user of an interaction client 104, based on a number of inputs and data, such as for example:

- Geolocation of the client system 102; and
- interactive platform information of the user of the client system 102.

[0066] An augmentation may include audio and visual content and visual effects. Examples of audio and visual content include pictures, texts, logos, animations, and sound effects. An example of a visual effect includes color overlaying. The audio and visual content or the visual effects can be applied to a media content item (e.g., a photo or video) at client system 102 for communication in a message, or applied to video content, such as a video content stream or feed transmitted from an interaction client 104. As such, the image processing system 202 may interact with, and support, the various subsystems of the communication system 208, such as the messaging system 210 and the video communication system 212.

[0067] A media overlay may include text or image data that can be overlaid on top of a photograph taken by the client system 102 or a video stream produced by the client system 102. In some examples, the media overlay may be a location overlay (e.g., Venice beach), a name of a live event, or a name of a merchant overlay (e.g., Beach Coffee House). In further examples, the image processing system 202 uses the geolocation of the client system 102 to identify a media overlay that includes the name of a merchant at the geolocation of the client system 102. The media overlay may include other indicia associated with the merchant. The media overlays may be stored in the databases 128 and accessed through the database server 126.

[0068] The image processing system 202 provides a user-based publication platform that enables users to select a geolocation on a map and upload content associated with the selected geolocation. The user may also specify circumstances under which a particular media overlay should be offered to other users. The image processing system

202 generates a media overlay that includes the uploaded content and associates the uploaded content with the selected geolocation.

[0069] The augmentation creation system 214 supports augmented reality developer platforms and includes an application for content creators (e.g., artists and developers) to create and publish augmentations (e.g., augmented reality experiences) of the interaction client 104. The augmentation creation system 214 provides a library of built-in features and tools to content creators including, for example custom shaders, tracking technology, and templates.

[0070] In some examples, the augmentation creation system 214 provides a merchant-based publication platform that enables merchants to select a particular augmentation associated with a geolocation via a bidding process. For example, the augmentation creation system 214 associates a media overlay of the highest bidding merchant with a corresponding geolocation for a predefined amount of time.

[0071] A communication system 208 is responsible for enabling and processing multiple forms of communication and interaction within the interaction system 100 and includes a messaging system 210, a chatbot system 232, an audio communication system 216, and a video communication system 212. The messaging system 210 is responsible for enforcing the temporary or time-limited access to content by the interaction clients 104. The messaging system 210 incorporates multiple timers within an ephemeral timer system (not shown) that, based on duration and display parameters associated with a message or collection of messages (e.g., a story), selectively enable access (e.g., for presentation and display) to messages and associated content via the interaction client 104. Further details regarding the operation of the ephemeral timer system are provided below. The audio communication system 216 enables and supports audio communications (e.g., real-time audio chat) between multiple interaction clients 104. Similarly, the video communication system 212 enables and supports video communications (e.g., real-time video chat) between multiple interaction clients 104. The chatbot system 232 is responsible for generating responses to prompts received from a user and communicating a response to the prompt.

[0072] A user management system 218 is operationally responsible for the management of user data and profiles, and includes an interactive platform 220 that maintains interactive platform information regarding relationships between users of the interaction system 100.

[0073] A collection management system 222 is operationally responsible for managing sets or collections of media (e.g., collections of text, image video, and audio data). A collection of content (e.g., messages, including images, video, text, and audio) may be organized into an “event gallery” or an “event story.” Such a collection may be made available for a specified time period, such as the duration of an event to which the content relates. For example, content relating to a music concert may be made available as a “story” for the duration of that music concert. The collection management system 222 may also be responsible for publishing an icon that provides notification of a particular collection to the user interface of the interaction client 104. The collection management system 222 includes a curation function that allows a collection manager to manage and curate a particular collection of content. For example, the curation interface enables an event organizer to curate a collection of content relating to a specific event (e.g., delete inappropriate content or redundant messages). Additionally, the collection management system 222 employs machine vision (or image recognition technology) and content rules to curate a content collection automatically. In certain examples, compensation may be paid to a user to include user-generated content into a collection. In such cases, the collection management system 222 operates to automatically make payments to such users to use their content.

[0074] A map system 224 provides various geographic location functions and supports the presentation of map-based media content and messages by the interaction client 104. For example, the map system 224 enables the display of user icons or avatars (e.g., stored in profile data 1002) on a map to indicate a current or past location of “friends” of a user, as well as media content (e.g., collections of messages including photographs and videos) generated by such friends, within the context of a map. For example, a message posted by a user to the interaction system 100 from a specific geographic location may be displayed within the context of a map at that particular location to “friends” of a specific user on a map interface of the interaction client 104. A user can furthermore share his or her location and status information (e.g., using an appropriate status avatar) with other users of the interaction system 100 via the interaction client 104, with this location and status information being similarly displayed within the context of a map interface of the interaction client 104 to selected users.

[0075] A game system 226 provides various gaming functions within the context of the interaction client 104. The interaction client 104 provides a game interface providing a list of available games that can be launched by a user within the context of the

interaction client 104 and played with other users of the interaction system 100. The interaction system 100 further enables a particular user to invite other users to participate in the play of a specific game by issuing invitations to such other users from the interaction client 104. The interaction client 104 also supports audio, video, and text messaging (e.g., chats) within the context of gameplay, provides a leaderboard for the games, and also supports the provision of in-game rewards (e.g., coins and items).

[0076] An external resource system 228 provides an interface for the interaction client 104 to communicate with remote servers (e.g., third-party servers 112) to launch or access external resources, i.e., applications or applets. Each third-party server 112 hosts, for example, a markup language (e.g., HTML5) based application or a small-scale version of an application (e.g., game, utility, payment, or ride-sharing application). The interaction client 104 may launch a web-based resource (e.g., application) by accessing the HTML5 file from the third-party servers 112 associated with the web-based resource. Applications hosted by third-party servers 112 are programmed in JavaScript leveraging a Software Development Kit (SDK) provided by the interaction servers 124. The SDK includes Application Programming Interfaces (APIs) with functions that can be called or invoked by the web-based application. The interaction servers 124 host a JavaScript library that provides a given external resource access to specific user data of the interaction client 104. HTML5 is an example of technology for programming games, but applications and resources programmed based on other technologies can be used.

[0077] To integrate the functions of the SDK into the web-based resource, the SDK is downloaded by the third-party server 112 from the interaction servers 124 or is otherwise received by the third-party server 112. Once downloaded or received, the SDK is included as part of the application code of a web-based external resource. The code of the web-based resource can then call or invoke certain functions of the SDK to integrate features of the interaction client 104 into the web-based resource.

[0078] The SDK stored on the interaction server system 110 effectively provides the bridge between an external resource (e.g., applications 106 or applets) and the interaction client 104. This gives the user a seamless experience of communicating with other users on the interaction client 104 while also preserving the look and feel of the interaction client 104. To bridge communications between an external resource and an interaction client 104, the SDK facilitates communication between third-party servers 112 and the interaction client 104. A WebViewJavaScriptBridge running on a client

system 102 establishes two one-way communication channels between an external resource and the interaction client 104. Messages are sent between the external resource and the interaction client 104 via these communication channels asynchronously. Each SDK function invocation is sent as a message and callback. Each SDK function is implemented by constructing a unique callback identifier and sending a message with that callback identifier.

[0079] By using the SDK, not all information from the interaction client 104 is shared with third-party servers 112. The SDK limits which information is shared based on the needs of the external resource. Each third-party server 112 provides an HTML5 file corresponding to the web-based external resource to interaction servers 124. The interaction servers 124 can add a visual representation (such as a box art or other graphic) of the web-based external resource in the interaction client 104. Once the user selects the visual representation or instructs the interaction client 104 through a GUI of the interaction client 104 to access features of the web-based external resource, the interaction client 104 obtains the HTML5 file and instantiates the resources to access the features of the web-based external resource.

[0080] The interaction client 104 presents a graphical user interface (e.g., a landing page or title screen) for an external resource. During, before, or after presenting the landing page or title screen, the interaction client 104 determines whether the launched external resource has been previously authorized to access user data of the interaction client 104. In response to determining that the launched external resource has been previously authorized to access user data of the interaction client 104, the interaction client 104 presents another graphical user interface of the external resource that includes functions and features of the external resource. In response to determining that the launched external resource has not been previously authorized to access user data of the interaction client 104, after a threshold period of time (e.g., 3 seconds) of displaying the landing page or title screen of the external resource, the interaction client 104 slides up (e.g., animates a menu as surfacing from a bottom of the screen to a middle or other portion of the screen) a menu for authorizing the external resource to access the user data. The menu identifies the type of user data that the external resource will be authorized to use. In response to receiving a user selection of an accept option, the interaction client 104 adds the external resource to a list of authorized external resources and allows the external resource to access user data from the interaction client 104. The

external resource is authorized by the interaction client 104 to access the user data under an OAuth 2 framework.

[0081] The interaction client 104 controls the type of user data that is shared with external resources based on the type of external resource being authorized. For example, external resources that include full-scale applications (e.g., an application 106) are provided with access to a first type of user data (e.g., two-dimensional avatars of users with or without different avatar characteristics). As another example, external resources that include small-scale versions of applications (e.g., web-based versions of applications) are provided with access to a second type of user data (e.g., payment information, two-dimensional avatars of users, three-dimensional avatars of users, and avatars with various avatar characteristics). Avatar characteristics include different ways to customize a look and feel of an avatar, such as different poses, facial features, clothing, and so forth.

[0082] An advertisement system 230 operationally enables the purchasing of advertisements by third parties for presentation to end-users via the interaction clients 104 and also handles the delivery and presentation of these advertisements.

[0083] FIG. 3A is an illustration of an interactive session by a user with a chatbot system, FIG. 3B is an illustration of a block diagram of the chatbot system, and FIG. 3C is an illustration of a process flow diagram of a chat and advertising process of the chatbot system, according to some examples. A chatbot system 300 uses the chat and advertising process to conduct a chat session with a user during a (first) interactive session, and provide advertising content to the user in a later (second) interactive session.

[0084] The chatbot system 300, in some examples, is a software application that is designed to simulate human conversation through voice commands or text chats. It may employ natural language processing (NLP) and machine learning (ML)/artificial intelligence techniques to understand and interpret the user's input and generate a response.

[0085] In some examples, the chatbot's architecture comprises natural language understanding (NLU) components and dialogue management components. The NLU components are responsible for understanding a user's intent and extracting relevant information from the user's input. This is achieved by analyzing the user's input and mapping it to an intent. The NLU components can use various techniques such as rule-

based systems, statistical models, and neural networks to understand the user's input. The dialogue management components, on the other hand, generates a response to the user's input. The use the intent and any extracted information from the NLU components to determine an appropriate response. This can be done using rule-based systems, decision trees, or machine learning models, and the like.

[0086] In some examples, a chatbot system 300 may use a generative language model, such as large language model (LLM) 338, to improve its natural language understanding (NLU) and dialogue management capabilities. The LLM 338 is more fully described in referenced to FIG. 8 and FIG. 9.

[0087] In some examples, the chatbot system 300 may use the LLM 338 to analyze the user's input and extract relevant information, such as the user's intent and entities. The LLM 338 can also be used to identify and extract important information from unstructured text, such as a user's question or request. This can be done using techniques such as named entity recognition, part-of-speech tagging, sentiment analysis, and the like.

[0088] When the NLU components have extracted the relevant information from the user's input, the dialogue management module can use the LLM 338 to generate an appropriate response. The LLM 338 can be used to generate responses in a human-like manner by using techniques such as text generation, machine learning models, and the like.

[0089] In some examples, the knowledge base of the chatbot system 300 comprises a set of information that the chatbot can use to understand and respond to a user's input. This includes a predefined set of intents, entities, and responses, as well as external sources of information such as databases or Application Programming Interfaces (APIs) and the like. In some examples, intent information of a friend, or friends, of the user may be used to understand, inform, and/or respond to a user's intent.

[0090] In some examples, the chatbot system 300 may be integrated into various platforms such as websites, messaging apps, and mobile apps, allowing users to interact with it through text or voice commands.

[0091] In operation 302 of chatbot method 374, the chatbot system 300 receives, from a client system 336, a prompt 328 of a user during a first interactive session. For example, a user uses client system 336 to access an interactive server that hosts the chatbot system 300. The user enters a prompt, such as prompt 328, into the client system 336 and the

client system communicates the prompt 328 to the chatbot system 300. In some examples, the prompt 328 may include other types of data as well as text such as, but not limited to, image data, video data, audio data, electronic documents, links to data stored on the Internet or the client system 336, and the like. Regardless of the data type of the prompt 328, keyword attribution and expansion may be used to automatically generate a cluster of keywords or attributes that are associated with the received prompt 328. For example, image recognition may be deployed to identify objects and location associated with image data and to generate a keyword cluster or cloud that is then associated with the image- based prompt.

[0092] The prompt may furthermore be received through any number of interfaces and I/O components (e.g., the I/O components 1308) of a client system 102. These include gesture-based inputs obtained from a biometric component and inputs received via a brain-computer interface (BCI).

[0093] In operation 304, the chatbot system 300 generates, by the one or more processors, a response 370 based on the LLM 338 (LLM). For example, the 300 receives the prompt 328. The chatbot system 300 uses the LLM 338 to generate the response 370.

[0094] The chatbot system 300 generates a raw response 362 based on the prompt of the user, generates an adjusted input prompt 368 based on filtering the raw response 362, and generates a new response based on adjusted input prompt 368. The response filter component 340 uses a set of filtering criteria to eliminate specified content from the user feedback 360, for instance obscene words or concepts, or content that some may consider harmful. The chatbot system 300 generates adjusted input prompt 368 that is communicated to the LLM 338 as the prompt 328.

[0095] In some examples, the LLM 338 is hosted by the same system that hosts other components of the chatbot system 300. In some examples, the LLM 338 is hosted by a server system that is separate from the system that hosts the chatbot system 300 and the chatbot system 300 communicates with the LLM 338 over a network. For example, the chatbot system 300 receives the user prompt and communicates the user prompt to the LLM 338 residing on the separate system. The LLM 338 receives the prompt 328 and generates the raw response 362. The LLM 338 then communicates the raw response 362 to the chatbot system 300. The chatbot system 300 receives the raw response 362 for subsequent processing.

[0096] In some examples, the chatbot system 300 generates a set of potential responses and selects the response from the set of potential responses. In some examples, the chatbot system 300 communicates the potential responses to the client system 336 and the client system 336 displays them to the user and the user selects the response that is most applicable to the user's prompt.

[0097] In some examples, as part of the response 370, the chatbot system 300 generates a set of chatbot system prompts that are displayed to the user by the client system 336 and prompting the user to interact with the chatbot system 300. The chatbot system prompts generated by the chatbot system 300 may include context sensitive material, instructions to the user, possible topics of conversation, and the like. In some examples, the chatbot system 300 uses the chatbot system prompts to suggest chat topics to a user or to guide the user through a conversation, such as providing instructional material on various topics. In some examples, the chatbot system prompts comprise suggestions of conversations or questions that are intended to solicit a user to enter a user prompt. In such examples, the suggested chats are aimed at helping the user get to information they need, but provide a helpful side effect of generating additional user interactions with the chatbot system 300 from which intent can be inferred.

[0098] In operation 306, the chatbot system 300 determines a user intent 366 based on the user prompt 328 and response 370. For example, a logging component 350 receives prompt 328, potential responses 364, and response 370 and any follow up messages and communicates them to an intent processing component 348 that includes a parallel intent processing pipeline. This pipeline uses Natural Language Processing (NLP) methodologies to map sets of conversations to a set of intentful keywords and concepts. In addition to the keywords which are used in the original prompt, the stemmed keywords and expanded concepts are also generated. The platform also assigns a weight to each concept based on the importance of those in the conversation and their commercial weight. These keywords and concepts are aggregated and mapped to the user as part of an intent profile or intent vector having weighted keywords and concepts based on the importance of those keywords in the conversation.

[0099] In some examples, the chatbot system 300 collects a set of prompts, as exemplified by prompt 328, additional prompt 372, and additional prompt 330, during an interactive session. The chatbot system 300 maps the set of prompts to a set of keywords and/or concepts comprising a user intent vector, as exemplified by keywords 332 and

concepts 334. The chatbot system 300 assigns weights to the keywords of the set of keywords and/or concepts based on an importance score to the conversation of the keyword and/or concept, and determines the user intent based on the user intent vector including the weighted keywords and/or concepts.

[0100] In some examples, the chatbot system 300 stores a conversation state in the user profile database 354 of a series of interactive sessions so that the chatbot system 300 can have a context for conversations that occur over a plurality of interactive sessions.

[0101] Any conversation data collected by the chatbot system 300 are captured and cached with only user approval and deleted on user request. Further, such conversation data may be used for very limited purposes, such as generation of responses by the chatbot system 300. To ensure limited and authorized use of conversation data and other PII, access to this data is restricted to authorized personnel only, if at all. Any use of conversation data may strictly be limited to chatbot purposes, and the conversation data is not shared or sold to any third party without the explicit consent of the user. In addition, appropriate technical and organizational measures are implemented to ensure the security and confidentiality of this sensitive information.

[0102] In some examples, the chatbot system 300 determines a personality or tone for the responses generated by the chatbot system 300. For example, the chatbot system 300 stores a conversation state for the user as part of the user's profile stored in the user profile database 354. The chatbot system 300 uses the conversation state and demographic or other information about the user to determine a personality or tone for the user, such as by adopting a formal tone for an older user, and a more informal tone for a younger user. A user may also have the ability to specifically alter the “persona” of interactions with the chatbot system by specifically requesting that the chatbot system 300 respond or act in a specific way (e.g., “be funnier”, “answer in riddles” etc.). In addition to modifying the personality or tone of the responses, the visual interface presented by the chat system 300 may also be altered to reflect a specific “persona” of the chatbot system 300. In some examples, a slide toggle may be presented to enable a user to select between a menu of persona traits (e.g., “cheeky”, “wry”, “cute”, etc.)

[0103] In some examples, the chatbot system 300 associates a temporal time factor to a keyword or concept where the temporal time factor decays in time. For example, the chatbot system 300 attaches a temporal time factor to the keywords and concepts, which indicates how fresh that concept should be used for targeting and bidding. For example,

"Hotels in Cancun for spring break" vs "Planning for wedding next year" will have two different decaying factors. In some examples, the chatbot system 300 applies a temporal time factor to a conversation state.

[0104] In some examples, an overall user intent is built using various signals or data such as user demographics, location, device, engagement with organic surfaces of an interactive platform or features of an interactive platform application and consumption patterns, and overall friend-graph proximity carrier signals or data that may be discernable from the user's affiliation with the interactive platform.

[0105] In some examples, the intent vector is an additional dimension which will be used to refresh and update an intent profile stored in the user profile database 354.

[0106] In some examples, the chatbot system 300 uses a user profile 376 from a user profile database 354 that includes interactive platform data of the user.

[0107] In some examples, the chatbot system 300 generates the advertising content 358 using an interaction by the user with previously provided advertising content where the interaction is captured during a second interactive session. For example, the chatbot system 300 is a component of an interactive platforming system that provides the advertising content 358 to the user in the context of the user's interactions with the interactive platforming system outside of the user's interaction with the chatbot system itself. The chatbot system 300 receives data of the user's interaction with the advertising content 358 in an advertisement analytics database 344 and uses the data of the user's interactions to determine how effective the intent processing component 348 was in determining the user's intent.

[0108] In some examples, the intent processing component 348 uses artificial intelligence methodologies including Machine Learning (ML) models to generate the user intent 366. The chatbot system 300 uses the data on user intent and conversation context to improve its targeting and optimization capabilities over time by ingesting the feedback. The chatbot system 300 collects engagement of the user on ads, organic surfaces or features of the interactive platforming system and also responses to the chatbot system 232 itself and feeds that into the intent processing component 348. Accordingly, the chatbot system 300 can finetune or further pre-train the ML models of the intent processing component 348 to not only take user intent into consideration, but also use follow up actions to finetune the user intent inference model.

[0109] In operation 308, the chatbot system 300 determines advertising content 358 based on the user intent 366 and the response 370 as more fully described in reference to FIG. 3D.

[0110] In operation 310, the chatbot system 300 communicates to the client system 336, the advertising content 358 during a second interactive session. For example, the user interacts with an interaction system 100 (of FIG. 1) of an interactive platform and accesses some additional features of the interactive platform. As the user interacts with the interaction system 100, the interaction system 100 provides the advertising content 358 outside of the initial interactive session where the prompts were received and the responses generated.

[0111] In some examples, the system that hosts the chatbot system 300 may not be a component of an interactive platform but another interaction system that provides services and information to a group of users such as, but not limited to, a platform that provides enterprise wide connectivity to a group of users such as employees of a company, clients of an enterprise provided professional services, educational institutions, and the like. In some of such examples, the content provided to the users may not be advertising, but may be other types of useful information such as company policies, status messages for projects, newsworthy events, and the like.

[0112] In some examples, end-to-end encryption is used to secure communications thus ensuring that only the sender and the intended recipient can read the messages being exchanged. By implementing end-to-end encryption, the chatbot system 300 can provide a secure and private messaging experience for users while still extracting intent for advertising experience enhancement. In the context of user intent extraction, this means that once the conversation is end-to-end encrypted, the chatbot system 300 as an end of the conversation can decrypt messages and pass those to the intent extraction pipeline of the intent processing component 348.

[0113] In some examples, the chatbot system 300 is operatively connected to an Internet search engine or the like and the user can use the chatbot system 300 as an intelligent search engine to search the Internet.

[0114] In some examples, the chatbot system 300 is operatively connected to a proprietary database and the user can use the chatbot system 300 as an intelligent search assistant for searching the proprietary database.

[0115] In some examples, the LLM 338 is continuously retrained using information about advertising content 358 stored in the advertising and campaigns database 346. For example, the LLM 338 is trained on the latest products being offered by advertisers on the interaction system 100 and offers product information on those products in response to requests by a user for the newest model of particular product categories.

[0116] In some examples, the LLM 338 is continuously retrained or finetuned based on user interactions with the advertising content 358. For example, the interaction system 100 collects advertising content engagement metrics and stores the metrics in the advertisement analytics database 344. The metrics are then used to provide reinforcement to the LLM 338 when the LLM 338 provides a sequence of responses that lead to a successful user intent determination and consequently properly targeted advertising content 358 that the user interacted with.

[0117] In some examples, the chatbot system 300 uses conversations between users and other chatbots as inputs into the intent processing component 348. The other chatbots can be sponsored chatbots, custom chatbots built by users from self serve tools / templates, and the like.

[0118] In some examples, a number of messages that the chatbot system 300 can receive after the chatbot system 300 enters a conversation are limited in number, limited in time that the chatbot system 300 may retain the messages and/or limited in a scope of what the chatbot system 300 can do with the messages. In some examples, the chatbot system 300 explicitly indicates that the chatbot system 300 has joined a conversation in the interests of transparency. In some examples, the chatbot system 300 provides a disclosure regarding what is being shared, both via a presence indicator and/or other chat affordances. In some examples, the disclosure includes a statement that the chatbot system 300 is receiving messages in order to improve services provided to the users.

[0119] In some examples, the information that the chatbot system 300 extracts from conversations is focused on the user intent. In addition, the chatbot system 300 furthers a conversation with a user by knowing that user intent (e.g., if the user asks for good hotels in Cancun, the chatbot system 300 responds with: "Here is a list of hotels. I also know of a good promotion for a hotel, would you like to see it?"). This provides for the user intent being extracted from the user, matching the intent to possible ads, and embedding that knowledge into a response. In some examples, the chatbot system 300 augments response with a "popular with friends" list.

[0120] In some examples, textual components of ads are used to finetune the responses of the LLM 338 conditioning the responses of the LLM 338. This improves the suggestions the chatbot system 300 and improves user engagement.

[0121] FIG. 3D is an illustration of a process flow diagram of an advertising creation process, according to some examples. A chatbot system 300 uses the advertising creation process to create advertising content 358 that is provided to a user outside of the context of an interactive chat session.

[0122] In operation 312, the chatbot system 300 receives, from an advertiser, a landing page website or application for analysis. For example, the chatbot system 300 receives a landing page website or application from an advertiser. This landing page or application contains up-to-date and accurate information about the advertiser's product or service offering. The chatbot system 300 analyzes the received landing page/application using natural language processing techniques to extract useful details like product features, visual assets, and branding elements. The chatbot system 300 extracts this information so it can later generate targeted and relevant advertising creatives that accurately reflect the advertiser's offering. By receiving and analyzing the advertiser's own landing page/application, the chatbot system 300 is able to create customized ad content that is aligned with the advertiser's messaging and branding.

[0123] In operation 314, the chatbot system 300 receives, from the advertiser, additional media assets such as images and videos. For example, the chatbot system 300 receives additional media assets like images and videos from the advertiser in operation 314. By obtaining these supplemental visual media assets directly from the advertiser, the chatbot system 300 can incorporate elements that authentically represent the advertiser's brand identity and product offering. The chatbot system 300 analyzes the received images and videos using computer vision techniques to identify relevant visual components. These components, such as product imagery, brand logos, and video testimonials, provide additional creative elements that the chatbot system 300 can integrate when generating customized advertising content for the advertiser. By sourcing authentic brand imagery and videos from the advertiser, the chatbot system 300 is equipped to produce compelling ad creatives that resonate with the advertiser's target audience.

[0124] In operation 316, the chatbot system 300 receives, from the advertiser, additional context such as product specifications, product catalogs, and other relevant

assets. For example, the chatbot system 300 receives additional context like product specifications, product catalogs, and other relevant assets from the advertiser in operation 316. By obtaining supplemental information directly from the advertiser, the chatbot system 300 can better understand the full capabilities and value proposition of the advertiser's offering. The chatbot system 300 analyzes the received product details and catalogs using natural language processing to identify key attributes and messaging. These insights, such as technical specifications, usage scenarios, and benefits, allow the chatbot system 300 to craft advertising copy and creatives that accurately capture the essence of the product or service. By sourcing relevant contextual information from the advertiser, the chatbot system 300 can generate ads with messaging and visuals tailored to promote the unique selling points and differentiators of the offering. The additional context equips the chatbot system 300 to produce ads that resonate with the target audience.

[0125] In operation 318, the chatbot system 300 receives, from the advertiser, a selection from a variety of potential directional themes and templates for creating advertising creatives. For example, the chatbot system 300 receives a selection of potential directional themes and templates for creating advertising creatives from the advertiser. By allowing the advertiser to choose from different creative themes and templates, the chatbot system 300 aligns the visual style and messaging tone of the generated ads to the advertiser's brand and goals. The advertiser's selection provides guiding creative direction to the chatbot system 300 as it assembles visual, textual, and contextual components into advertising creatives. With the advertiser's input on preferred themes and templates, the chatbot system 300 can produce customized ads that adhere to the look, feel, and messaging that the advertiser aims to project in their marketing campaigns. The chosen creative direction steers the chatbot system 300 to generate ads with designs, layouts, and copy that help the advertiser's content resonate with their target audience.

[0126] In operation 320, the chatbot system 300 receives, from the advertiser, target keywords and concepts. For example, the chatbot system 300 receives target keywords and concepts from the advertiser in operation 320. By obtaining specific keywords and concepts directly from the advertiser, the chatbot system 300 can incorporate messaging into the generated ads that is directly relevant to the advertiser's offering and target audience. The chatbot system 300 analyzes the received keywords using natural language processing to identify key messaging themes. These targeted keywords and expanded

concepts allow the chatbot system 300 to produce ad copy and creatives that contain the precise terminology and messaging to effectively reach and resonate with the intended audience. The advertiser-supplied keywords and concepts provide the chatbot system 300 with the exact language needed to craft compelling ads tailored to promote the advertiser's product or service to the desired customer segments.

[0127] In operation 322, the chatbot system 300 extracts potential additional assets and metadata from the destination website or application, and takes into consideration the input keywords and concepts. For example, the chatbot system 300 extracts potential additional assets and metadata from the advertiser's website or application and considers the input keywords and concepts in operation 322. By programmatically analyzing the advertiser's digital presence, the chatbot system 300 can discover supplemental creative elements like images, videos, and text content. The chatbot system 300 leverages computer vision, natural language processing, and data mining techniques to identify relevant components from the website and app. These extracted creative assets and metadata, combined with the advertiser's supplied keywords and concepts, equip the chatbot system 300 with a robust set of components to assemble compelling, on-brand advertising creatives. By algorithmically harvesting assets from the advertiser's channels, the chatbot system 300 can generate customized ads tailored to promote the advertiser's brand and offering.

[0128] In operation 324, the chatbot system 300 uses a creative generation component 356 to generate a series of advertising content creatives based on this information. For example, the creative generation component 356 of the chatbot system 300 utilizes all of the collected information and assets to automatically generate a set of advertising content creatives. Specifically, the creative generation component 356 leverages the product details, brand guidelines, target keywords, creative direction, and digital assets gathered in previous operations to algorithmically produce multiple versions of ads tailored for the advertiser. Using advanced generative AI techniques, the creative generation component 356 synthesizes unique combinations of images, text, logos, and layouts into customized advertising creatives that align with the advertiser's brand, products, and target audience. By bringing together all of the inputs provided by the advertiser along with extracted metadata and assets, the creative generation component 356 can efficiently and automatically generate high-quality, on-strategy advertising creatives on behalf of the advertiser.

[0129] In operation 326, the chatbot system 300 modifies delivery of advertising content based on user engagement with the advertising content. For example, the chatbot system 300 modifies the delivery of advertising content based on user engagement with the ads. Specifically, the chatbot system 300 tracks metrics such as click-through rates, conversion rates, and dwell time for the generated ads. Using this engagement data, the chatbot system 300 can determine the highest performing ads and optimize the ad delivery accordingly. For example, ads with higher click-through or conversion rates may be shown more frequently or to broader audiences. Conversely, lower performing ads may be shown less or stopped altogether. The chatbot system 300 can also iterate on ad creative, copy, placement, etc. based on observed user engagement patterns. By leveraging user feedback and advertising performance data, the chatbot system 300 refines both ad creation and ad delivery to maximize results for the advertiser. This allows the system to automatically optimize campaigns and serve ads that resonate best with target users.

[0130] In some examples, the chatbot system 300 uses an advertising content delivery and bidding component 352 to provide advertisers an opportunity to bid on certain actions by choosing keywords or an expanded set of concepts (auto expansion) to select their potential target audience and display advertising content 358 that will be delivered to users who match that criteria within their intent vector. Overall, by mapping conversations to a set of keywords and concepts and attaching a temporal time factor to them, and finally mapping those to user intent and profile, the advertising content delivery and bidding component 352 of the chatbot system 300 enables advertisers to find their target audience much more precisely.

[0131] In some examples, the chatbot system 300 includes an ad targeting component 378. The ad targeting component 378 The ad targeting component utilizes the user intent and interests extracted from conversations with the chatbot system 300 to target and deliver relevant ads to users. For example, the ad targeting component 378 analyzes the user intent vector comprising weighted keywords and concepts determined by the intent processing component 348. It matches these keywords and concepts to target criteria specified by advertisers for their ad campaigns stored in the advertising and campaigns database 346. Ads are ranked based on the degree of match between the user intent vector and the advertiser targets. Ads with higher relevance are prioritized for delivery. The ad targeting component also considers other targeting parameters like user demographics, location, and platform engagement history from the user profile database

354. Additionally, the ad targeting component 378 tracks user engagement with delivered ads using feedback stored in the advertisement analytics database 344. The ad targeting component 378 optimizes ad targeting over time by determining which ads drive higher engagement for which user intents.

[0132] In some examples, the chatbot system 300 includes an ad ranking component 380. The ad ranking component 380 determines a priority and order in which ads are shown to users of the interactive platform. The ad targeting ad ranking component 380 uses user intent from chatbot conversations to deliver highly personalized, relevant ads to each user.

[0133] In some examples, the ad ranking component 380 ranks and scores ads using various criteria such as, but not limited to:

- Relevance to the user's intent vector - Ads are ranked higher if they closely match keywords/concepts the user expressed interest in during chatbot conversations.
- Expected engagement - Ads are scored based on historical click-through rates, conversion rates, and dwell times for different user intents.
- Bid amount - Advertisers can bid for specific keywords/concepts. Higher bids influence ad rank.
- Landing page relevance - The ad's landing page is analyzed for relevance to the user's intent. More relevant landing pages increase ad rank.
- Advertiser budget - Ad rank considers the remaining budget of the advertiser's campaign.

[0134] In some examples, the ad ranking component 380 combines one or more factors using a machine learning model to assign a final rank score to ads. The ad ranking component 380 selects the top ranking ads to show to each user.

[0135] In some examples, the ad ranking component 380 continuously monitors engagement with ranked ads and refines the model to improve ranking accuracy over time. This allows the most relevant ads to be surfaced to users based on their conversational intents with the chatbot system.

[0136] In some examples, the ad ranking component 380, can expand or narrow down targeting as it learns which concepts map better to user responses.

[0137] In some examples, an AI-driven ad creative generation component 356 of the chatbot system 300 assists advertisers by simplifying the process of creating advertising creatives. By providing inputs such as website, target application, additional assets, and

target keywords, the chatbot system 300 can automatically generate ad creatives. This can save advertisers time and resources, allowing them to focus on other aspects of their advertising campaigns.

[0138] In some examples, the chatbot system 300 utilizes machine learning methodologies to optimize the advertising content 358 creatives over time. The chatbot system 300 tests multiple permutations of the generated advertising content 358 to determine which permutations work better. Such a process will be more applicable to dynamic product ads which are generated based on a dynamic catalog stored in advertising and campaigns database 346. This improves advertising performance and increases the chances of success, ultimately providing a better return on investment for the advertisers. The chatbot system 300 not only saves time and resources for the advertisers but also increases the chances of success by providing the most relevant and effective advertising content 358 creatives for the target audience.

[0139] In some examples, the chatbot system 300 measures similarity of expanded concepts of a conversation with a user with advertiser ad concepts or contexts. This allows the advertisers to not only use specific words, but also sentences and descriptions for what they are looking to promote. In some examples, the ad relevance is a distance between ad concepts or contexts and the expanded concepts of the conversation.

[0140] FIG. 4A is a tech block diagram of a chatbot system 300 (of FIG. 3A), and FIG. 4B is an activity diagram of a method of the chatbot system 300, according to some examples.

[0141] In operation 402, a user uses an interactive platform application 408, such as application 106 executing on a client system 102 (of FIG. 1), to access the chatbot system 300 so that the user can chat with the chatbot system 300 or add it to existing chats. For example, the user utilizes the interactive platform application 408 to access the chatbot system 300, enabling the user to directly chat with the chatbot system 300 or add it to existing conversations. The interactive platform application 408 provides the user interface through which the user can interact with the chatbot system 300. Within the interactive platform application 408, the user has the option to initiate a 1-on-1 chat session with the chatbot system 300 or bring the chatbot system 300 into a group chat or messaging thread. This gives the user flexibility to engage the chatbot system 300 for personalized dialogue or within the context of ongoing discussions with other users. From the interactive platform application 408, the user can send textual prompts to the

chatbot system 300 and receive conversational responses. By accessing the chatbot system 300 through the interactive platform application 408, the user can leverage the chatbot's capabilities to obtain information, recommendations, and other services through natural conversation flows.

[0142] In operation 404, the full conversation 410 is communicated to the LLM 338. For example, the entire conversation between the user and the chatbot system 300 is communicated to the Large Language Model (LLM) 338. This provides the LLM 338 with the full context of the interactive dialog, including the initial prompt from the user as well as any follow-up messages. By analyzing the complete exchange comprising the full conversation, the LLM 338 can better understand the intent and meaning behind the user's queries. Additionally, having access to the full conversation flow enables the LLM 338 to maintain continuity and generate responses that logically follow from previous parts of the dialog. This conversational context allows the LLM 338 to produce more natural, relevant responses compared to only looking at individual messages in isolation. Feeding the complete back-and-forth conversation to the LLM 338 empowers it to have a more holistic understanding of the interaction and in turn generate superior responses.

[0143] In operation 406, responses 412a and chats are sent back to the user, from the output selector component 342, which can include textual responses, links to content, relevant ads, or any other response. For example, the output selector component 342 determines the most appropriate type of response to return, which could include textual responses, links to relevant content, ads targeted to the user's interests, or any other suitable response. In some examples, the output selector component 342 may decide a textual response is most fitting for a factual question, while an ad for a product or service may be most relevant for a commercial query. The output selector component 342 chooses among the available response types to provide the user with the most useful, tailored information. By supporting multiple response formats like text, links, and ads, the chatbot system 300 can dynamically serve the user with the most fitting type of reply for each conversation turn. This allows the chatbot system 300 to have a diverse, engaging dialogue with the user.

[0144] In some examples, a more detailed textual response is generated by the chatbot system 300 beyond a short reply. To do so, the LLM 338 generates a longer, more elaborate response. The output selector component 342 provides additional context 414 to the LLM 338 to help construct this longer response. For example, the output selector

component 342 may supply conversational history, user profile data, or other supplementary information to give the LLM 338 more context. Equipped with this extra context, the LLM 338 is able to produce a richer, more natural sounding lengthy response. The output selector component 342 determines when the dialog requires more elaboration and leverages the LLM 338's capabilities to generate articulate, contextual responses of multiple sentences or paragraphs in length. This allows the chatbot system 300 to handle user queries that demand more than just a short, simple response.

[0145] In some examples, the LLM 338 is fine-tuned to conduct conversations in the context of the interactive platform hosting the chatbot system 300, taking into account additional user profile data 416 from a user profile database 354. The LLM 338 is customized to incorporate details from the user's profile on the interactive platform, such as their demographics, interests, and past activity. This allows the LLM 338 to have a basic understanding of who the user is and tailor its responses accordingly. In some examples, the LLM 338 is adapted to leverage the additional context 414 signals that may be relevant to the conversation on the interactive platform. This includes factors like the user's friends, location, time of day, and recent interactions on the platform. By factoring in these contextual clues, the LLM 338 can generate more situational and personalized responses. By fine-tuning the LLM 338 using profile data and additional context from the interactive platform, the chatbot system 300 can produce conversations that are more natural, intuitive, and engaging for users within the specific environment of the interactive platform. This platform-optimized training helps address user needs and expectations for conversations on the interactive platform.

[0146] In some examples, the LLM 338 responses 418, together with the LLM 338 context (conversation embedding) and additional context such as initial intent from the LLM 338 are sent to the intent processing component 348. The intent processing component 348 defines skill context 434 and details (for example keywords), to different relevant skill modules 422 and detects the user intent, and conversational state 420. For example, the intent processing component 348 analyzes the conversation between the user and the chatbot system 300 to extract important keywords, entities, and details. It passes these context clues to various skills modules that can generate relevant responses. Additionally, the intent processing component 348 leverages machine learning models to infer the user's true intent and state of the conversation from the contextual signals. For instance, the output selector component 342 detects whether the user is asking about purchasing diapers or requesting a Wikipedia definition. By extracting conversation

details and detecting user intent, the intent processing component 348 enables downstream components to produce responses tailored to the user's needs.

[0147] In some examples, the skill modules 422 include a dynamic content module 426. The dynamic content module 426 provides relevant, personalized information to users during a conversation by leveraging various data sources and generative techniques. Some examples of dynamic content operations comprise:

- Retrieving user-specific data like preferences, purchase history, calendar events, and the like to provide customized responses.
- Generating personalized recommendations for products, services, content, and the like tailored to the user.
- Synthesizing relevant news articles, social media posts, weather reports, and the like on topics discussed.
- Creating dynamic visualizations such as charts, graphs, maps, and the like on-the-fly to visualize data requested by the user.
- Generating custom imagery or multimedia to illustrate points in the conversation.
- Producing audio clips with relevant sound effects, background music, and the like to engage the user.
- Integrating with APIs, databases, websites to fetch up-to-date information as needed.
- Adapting responses to reflect latest offers, deals, pricing for products, and the like based on live data.
- Updating responses with real-time data like sports scores, stock prices, traffic, and the like.
- Synthesizing previous parts of the conversation into summaries to provide context.

[0148] The dynamic content module 426 taps into various data sources and generative models to create unique, personalized, dynamic content on the fly during a conversation. This makes the dialogue more engaging and relevant for each user.

[0149] In some examples, the skill modules 422 comprise a dynamic interactive platform ads module 424. The dynamic interactive platform ads module 424 provides relevant, targeted advertisements to users during a conversation. Operations of the dynamic interactive platform ads module 424 comprise:

- Analyzing the conversation context and user profile to determine user interests and intent.
- Using this understanding of user intent to retrieve and rank relevant ads from an ad inventory.
- Generating personalized ad copy and creatives on the fly tailored to the user.
- Dynamically inserting highly relevant ads at appropriate points in the conversation.
- Adapting ad creative elements like images, videos, audio to match conversation context.
- Updating ads with latest pricing, deals, offers pulled from live ad data.
- Pulling real-time performance data to optimize ad targeting and improve relevance.
- Providing interactive ad experiences like mini-games, polls, special offers etc.
- Enabling users to provide feedback on ads to further refine targeting.
- Measuring ad engagement, clicks, conversions to evaluate and improve ad performance.
- Following platform advertising policies and complying with user preferences.

[0150] The dynamic interactive platform ads module 424 taps into user intent, conversation context, and ad data to deliver highly dynamic, personalized, and relevant ads within the conversational experience. This benefits users with ads they want to see as well as advertisers reaching the right customers.

[0151] In some examples, the skill modules 422 comprise an open domain knowledge module 428. Operations of the open domain knowledge module 428 comprise:

- Answering factual questions by searching knowledge bases like Wikipedia, Wikidata, DBpedia etc.
- Providing definitions, explanations, and descriptions of concepts the user asks about.
- Summarize key information from long articles, papers, or passages relevant to the user's queries.
- Generating visual aids like info-graphics, charts, timelines, maps to illustrate complex topics.
- Enabling clarifying follow-up questions from the user to refine their information needs.

- Ranking and filter responses to provide the most relevant, helpful information to the user.
- Citing sources and link out to original references for transparency.
- Leveraging large language models pre-trained on massive text corpora to generate informative responses.
- Ingesting new information from various structured data sources.
- Asking users for feedback to improve the quality of information provided.
- Following ethical guidelines around information accuracy, attribution, and reliability.

[0152] The open domain knowledge module 428 provides users with informative, trustworthy, and useful information on any topic they inquire about by tapping into open data resources in a responsible manner. The goal is to enhance the conversational experience with the knowledge of the world.

[0153] In some examples, the skill modules 422 comprise a generative AI replies module 430. Some operations of a generative AI replies module 430 comprise:

- Generating natural language responses to user inputs using large language models like the LLM 338
- Allowing open-ended discussions on a wide range of topics beyond just questions and answers.
- Maintaining context and consistency across multiple turns of a conversation.
- Exhibiting personality and emotional intelligence through word choice and tone.
- Providing nuanced, non-repetitive responses tailored to each user input.
- Incorporating world knowledge from pre-training corpora into responses.
- Enabling mixed-initiative conversations where both the user and AI contribute.
- Dynamically adjusting reply style, length, depth based on user preferences.
- Seamlessly integrating with other modules like speech, vision, knowledge etc.
- Learning from user feedback and conversations to improve over time.
- Ensuring responses meet platform content policies and community guidelines.
- Attributing sources when incorporating external information into responses.
- Indicating when responses are speculative vs factual.

[0154] The generative AI replies module 430 leverages AI to produce human-like, engaging conversations that adapt to each user and conversation context while meeting interactive platform policies.

[0155] The various relevant skill modules 422 return skill replies 432 as potential responses 364 (Of FIG. 3A) to continue the conversation 410, which may include suggestions for dynamic content, relevant ads, open domain responses, and more following the output of the above skill modules. The skill replies 432 are sent back to the output selector component 342 as potential responses 364.

[0156] The output selector component 342 determines which response is more appropriate to s current stage of the conversation 410, and either returns the skill replies 432 to the LLM 338 for natural answer generation, or uses the responses 412a and 412b already generated and returns them directly to the interactive platform application 408 and the user.

[0157] In some examples, the LLM 338 is fine-tuned toward chatting in approved an interactive platform's approved chat styles, providing safe responses, respecting interactive platform values, providing multiple persona traits as described above, incorporating user context, and/or previous conversation history and optionally returning high level intents.

[0158] In some examples, the intent processing component 348 is based on an AI component such as a neural net or the like, and/or business logic designed to determine significant interactive platform intents (for example, information seeking, game playing, chit chat and more) and their attributes (information seeking - Baby diapers alternatives). The model is trained on extracting the intents from the LLM 338 context, or other embeddings of the conversation and conversation history. In addition to the model, simpler models such as keyword extraction, and categorization can be applied to enrich the intent details. The intent processing component 348 provides these details to the various skills skill modules 422. In some examples, the skill modules 422 are a set of parallel modules each designed for returning a response based on separated abilities and concerns. For example, a dynamic interactive platform ads module 424, given the intent and details, re-ranks and select the most appropriate ad to place as a result, whereas the generative AI replies module 430, given the intent and context would generate a multi-media result using the generative AI capabilities.

[0159] In some examples, not all of the skill modules are necessarily be used for each intent but for some intents several modules may be used generating several alternatives for response. The output selector component 342 applies the business logic to select the

most appropriate results at each turn, according to the business policy, their relative effectiveness and other metrics.

[0160] FIG. 5A and FIG. 5B are illustrations of user interfaces, according to some examples. A chatbot system 300 (of FIG. 3B) uses the user interfaces to communicate that the chatbot system 300 is active and available for interactions with users and may be acquiring information about a user's interaction with an interactive platform. User interactions with the chatbot system 300 may be multimodal. In some examples, the user may interact with the chatbot system 300 directly by exchanging user prompts and chatbot responses directly. In some examples, two or more users may interact with one another while the chatbot system 300 monitors the interaction in order to discern user intent of one or more of the users. In some examples, the chatbot system 300 is made available to a user in a user interface for invocation during a chat session with another user.

[0161] In some examples, the chatbot system 300 shows that the chatbot system 300 is listening to a conversation by adding a chatbot icon 502 as a member of the chat in a presences bar (i.e., alongside Bitmojis of friends, a character). In some examples, the chatbot system 300 provides a chatbot icon 504 to represent the chatbot system 300 as a member of a chat using an avatar. A user taps the chatbot icon 504 on the presence bar to start asking the chatbot system 300 a question or writes @ChatBot in the chat bar. In some examples, the 300 uses an “X is typing”-style visual as a loading indicator while the chatbot system 300 prepares a response to a query. In some examples, a user uses “@” to add the chatbot system 300 to a Chat and directly summons the 300 to help with a query.

[0162] In some examples, a user learns about, creates, and gets help with just about anything during an interactive session with the chatbot system 300. In some examples, the interactive session is a private interactive session with the chatbot system 300. In some examples, the interactive session includes one or more other users and one or more users ask the chatbot system 300 for help directly during the interactive session. In some examples, the chatbot system 300 provides assistance to a user in accessing different features of an interactive system hosting the chatbot system 300 or another interactive system or Internet locations.

[0163] In some examples, the chatbot system 300 helps a user in the same way an intelligent friend would, such as by communicating directly with a user and a user's

friends directly using a chat interface and by sending social media content and other content directly to the user.

[0164] In some examples, a user utilizes the chatbot system 300 as a primary user interface with the interactive system hosting the chatbot system 300. As the chatbot system 300 infers the user intent from interactions with the user and has a conversation state that is saved, the chatbot system 300 also infers what a user is interested in based on what the chatbot system 300 learns from the chats it has access to (including 1:1 chat threads with the chatbot system 300), a user's general usage of the interactive system hosting the chatbot system 300, other third-party communications with the user, how the user communicates, who the user's friends are, as well as those friends' user intents, profiles, geographic location(s), interests, and the chatbot system 300 helps the user learn about, create, and engage in new activities.

[0165] In some examples, the chatbot system 300 influences the content the user sees across the interactive system that hosts the chatbot system 300 and in chats where the chatbot system 300 is a participant. For example, if a user asks about fun games to play with the user's baby, the chatbot system 300 may suggest newborn advice videos on a content platform or recommend an ad for a baby product.

[0166] In some examples, the chatbot system 300 can send any kind of content that a user of the interactive platform hosting the chatbot system 300 can send such as, but not limited to, chats, any type of media such as images, videos, audio recordings, social network posts, chat media, web links, map places, AI-generated media, media from the user's own datastore, and the like.

[0167] In some examples, implementation of the chatbot system 300 in a conversational interface integrated within an interactive application provides messaging, content, and display advertisements (including AI-generated ads based on keyword winning links) which also helps provide signals for ranking of all content provided by the chatbot system 300 (including display advertisements, media, and AR).

[0168] In some examples, interactions with the chatbot system 300 across all features of the interactive system that hosts the chatbot system 300 improves the LLM 338 over time. Interaction data collected and used to train the LLM 338 include, but are not limited to, what is asked, follow-up questions, in application reporting of answers, in-application reactions to answers, forwarding of answers, sharing of answers, retention, frequency of use, ad interactions, other searches, other actions, visual content of social

network media content, camera rolls, user datastores of interactive application interactions, and the like.

[0169] In some examples, the chatbot system 300, through retraining of the LLM 338, provides for interactive system wide, per-user, and individual conversation training. In some examples, the LLM 338 is trained as part of a system-wide model based on all interactive system interactions. In some examples, the LLM 338 is trained in a per-user and even per conversation model (i.e., “what should I/we do today” yields different results for different people and for the same people asking it in different conversations).

[0170] In some examples, the chatbot system 300 provides suggested chat topics as chatbot system prompts, both before and after answering a question. For example, the chatbot system 300 provides the chatbot system prompts of what a user may say or do next. Tapping the suggestion prompts sends resends them to the chatbot system 300 as new user prompts.

[0171] In some examples, the chatbot system 300 provides a mixed or multi-modal chat session as an interactive session. For example, the chatbot system 300 infers when to send a user plain text chat data and when other kinds of media are required to satisfy a user’s request. In some examples, a user may send user prompts to the chatbot system 300 silently within an existing interactive session with one or more other users. In some examples, the chatbot system 300 provides the user the ability to send requests to the chatbot system 300 within a conversation without other participants seeing questions or answers unless the user chooses to make those interactions visible.

[0172] In some examples, the chatbot system 300 provides the ability for a user to modify the personality of the chatbot system 300 such as, but not limited to, receiving an explanation of how the user wants the chatbot system 300 to act directly, alter the persona based on how the user prefers to communicate and how the user’s friends communicate with the user, providing a visual interface to alter the persona of the chatbot system 300, use user profile information such as age, location, and content preferences to alter the persona, and the like.

[0173] In some examples, the chatbot system 300 provides for the reporting of answers in chat using chat reporting features of the interactive system hosting the chatbot system 300, and using the reporting to improve the LLM 338 by letting users provide the right response.

[0174] In some examples, the chatbot system 300 provides for assigning accuracy scores to users who provide feedback.

[0175] In some examples, the chatbot system 300 provides for human review of most asked for topics. In some examples, the chatbot system 300 provides a verified answer framework.

[0176] In some examples, the chatbot system 300 provides for personality moderation of response such as “I don’t feel like I know for sure, but I think the closest answer I can give you is ...”.

[0177] In some examples, the chatbot system 300 provides for confidence scoring of responses used to modify the way the chatbot system 300 generates a response, such as “Hmm, it’s not 100% but my best bet is ...”

[0178] In some examples, the chatbot system 300 provides for solo or private interactive sessions and a user may provide user prompts on topics such as, but not limited to, relationship advice, finding content, homework help, general learning, fashion, connect verbal chat requests to a visual search database, and the like.

[0179] In some examples, the chatbot system 300 provides an eCommerce front end by allowing for verbal or text user prompts to search for products or services, (e.g., “Nissan Ultra 2012 under \$15k”, best car for newborns).

[0180] In some examples, the chatbot system 300 provides for a quick search by receiving a short query that is extrapolated out to a more detailed query (e.g., “Nike history” extrapolated into questions automatically like: “100 words on how Nike started”, “What did Nike’s leaders have in common?”) based on queries an interactive platform’s community have provided as user prompts.

[0181] In some examples, the chatbot system 300 provides for a user interacting with one or more other users and the chatbot system 300 together and the chatbot system 300 responding with a context sensitive response (e.g., “pick someone to cook tonight!”, “write a story about our friendship”, “find somewhere close by all of us for breakfast on Thursday”).

[0182] In some examples, the chatbot system 300 provides for a user to program a chat based application or function using verbal commands (e. g., create an account that is all about positivity and send the user inspirational quotes”.

[0183] In some examples, the chatbot system 300 provides the ability to easily forward responses of the chatbot system 300 to other interactive sessions with clear attribution.

[0184] In some examples, the chatbot system 300 provides the ability to export a photo or video of a conversation as part social network media, including with a watermark.

[0185] In some examples, the chatbot system 300 provides a sped up timelapse style of question and answer being typed out.

[0186] In some examples, the chatbot system 300 provides the user the ability to choose different styles.

[0187] In some examples, the chatbot system 300 caches other user prompts and access-limiting process 1200 responses and suggests user prompts and/or responses based on what a user is entering as a current user prompt and the cached user prompts and chatbot system 300 responses.

[0188] In some examples, the chatbot system 300 provides the ability for users to see user prompts entered by friends as suggestions. A user taps the previous user prompts to see who asked it and what the response was. In some such examples, a user may react and leave comments).

[0189] In some examples, the chatbot system 300 provides comments on answers by users with a “factcheck” for other users to see.

[0190] In some examples, the chatbot system 300 provides for the sharing of results of queries regarding social media by sharing a typing video of the question and answer OR recording over visual overlay of Q&A responses with a images of a user’s reaction (still or video) in different styles (or add a memory or color/asset/generative media item as a background instead of a live camera).

[0191] In some examples, the chatbot system 300 provides the user with the ability to send social media posts to get edits and sometimes responses.

[0192] In some examples, the chatbot system 300 provides a user with the ability to send a social media post to the chatbot system 300 (including with a text question) using an in-chat camera where a text field is up by default on a preview.

[0193] In some examples, the chatbot system 300 receives a user prompt from a user in a form of a request to produce an image or video.

[0194] In some examples, the chatbot system 300 provides users with the ability to name friends and use their cameos (or other likeness) to produce artificial intelligence media (e.g., “me and @chuckie_mills on a boat in Barbados”).

[0195] In some examples, the chatbot system 300 provides the user the ability to send social media posts to the chatbot system 300 with instructions (e.g., “make these all have balanced lighting”, “combine these two faces, make this one funnier”).

[0196] In some examples, the chatbot system 300 provides the user the ability to combine visual media with chat inputs (i.e., [sends picture of fridge] “What can I cook with this”, [sends interactive platform post of an image of my face] “make me look like a vampire”). In some such examples, the user interface includes an in-chat camera with an always-scanning feature.

[0197] FIG. 6 is a collaboration diagram of a sensitive content filtering system, according to some examples. A platform chatbot messaging system, such as platform chatbot messaging system 700 of FIG. 7A, of an interactive platform uses a sensitive content filter component 604 in conjunction with a chatbot system 616, such as chatbot system 300 of FIG. 3B, to enforce a set of platform policies 610 in regard to sensitive user prompts made by a user 602 using a user interface 608.

[0198] A user 602 may enter a user prompt 612 containing sensitive content, such as content relating to geopolitics, criminal activity, illicit drug use, a physiological condition of the user 602, and the like, into a user interface 608 of the interaction system 100. A sensitive content filter component 604 receives the user prompt 612 and analyzes the user prompt 612 for sensitive content. For example, the sensitive content filter component 604 parses the user prompt 612 to find keywords that signal that the content is sensitive. In some examples, the sensitive content filter component 604 utilizes AI methodologies to analyze the user prompt 612. In some examples, the keywords are specified in the platform policies 610.

[0199] When the sensitive content filter component 604 detects sensitive content in the user prompt 612, the sensitive content filter component 604 generates a response 614 based on the platform policies 610. For example, the sensitive content filter component 604 generates the response 614 based on a set of templates. In some examples, the sensitive content filter component 604 utilizes AI methodologies to generate the response 614. In some examples, the sensitive content filter component 604 performs a search on a platform for information or resources related to the sensitive content and uses the search results to generate the response 614.

[0200] When the sensitive content filter component 604 does not detect sensitive content, the sensitive content filter component 604 communicates the user prompt 612 to

the chatbot system 616 and the chatbot system 616 generates a response 614, as described more fully in reference to FIG. 3A and FIG. 3B, and communicates the response 614 to the user 602 via the sensitive content filter component 604 and the user interface 608.

[0201] In some examples, the sensitive content filter component 604 filters responses created by the chatbot system 616 that contain sensitive content. In response to detecting that the sensitive chatbot system 616 generated a response 614 containing sensitive content, the sensitive content filter component 604 generates a more appropriate response based on the platform policies 610. For example, the sensitive content filter component 604 filters the response 614 to determine if the response 614 comprises sensitive content. When the sensitive content filter component 604 detects that the response 614 comprises sensitive content, the sensitive content filter component 604 regenerates a prompt that is more appropriate and does not comprise the sensitive content and re-prompts the chatbot system 616 with the appropriate prompt. The chatbot system 616 receives the appropriate prompt and generates the appropriate response. In some examples, the sensitive content filter component 604 prompts the chatbot system 616 to regenerate the response to remove the sensitive content and generate the appropriate response with the sensitive content removed.

[0202] In some examples, the chatbot system 616 is pre-trained to generate the response 614 to a user prompt 612 containing sensitive content in accordance with the platform policies 610. In some examples, the sensitive content filter component 604 is incorporated into the chatbot system 616.

[0203] In some examples, the user 602 enters a user prompt 612 containing sensitive content about harmful conduct that is being directed to the user by another person. In response to detecting that the sensitive content comprises harmful conduct, the sensitive content filter component 604 encourages the user 602 to report harm to the interactive platform's Trust & Safety team. The sensitive content filter component 604 recognizes that there is a potential severe harm and redirects the user to seek help in an appropriate manner. For example, the chatbot system 616 generates a response 614 "I'm a chatbot and cannot help with real-world harms and emergencies. If you or someone you know has experienced harm, please submit a report to our Trust & Safety team. If it's an emergency, please call 911."

[0204] In some examples, the user 602 enters a user prompt 612 containing sensitive content about topics that relate to specific harms or topics that the interactive platform has created extensive native content for. The sensitive content filter component 604 searches for and surfaces specific resources that the platform has created for support the platform community in certain situations.

[0205] In some examples, the user 602 enters a user prompt 612 containing sensitive content about a geopolitically charged issue or about known conspiracy theories. The sensitive content filter component 604 generates a response 614 with a general response such as, but not limited to: “I’m not intended to talk about topics like this”.

[0206] In some examples, the user 602 enters a user prompt 612 containing sensitive content about or plans of an illegal or dangerous activity. It is possible that an account shares information about illegal, dangerous or violative activity with the chatbot system 616. The sensitive content filter component 604 generates a response 614 warning the user about the consequences of such action and may proactively create a task based on certain keywords if the platform policies 610 indicates this needs to be escalated to law enforcement.

[0207] In some examples, the user 602 enters a user prompt 612 containing sensitive content about platform employees. In response, the sensitive content filter component 604 generates a response 614 with a general response such as, but not limited to: “I’m not intended to talk about topics like this”.

[0208] FIG. 7A is a collaboration diagram of a platform chatbot messaging system 700, according to some examples. An interaction system 100 (of FIG. 1) of an interactive platform uses a platform chatbot messaging system 700 to create a stateful conversation with a user of the interactive platform.

[0209] The platform chatbot messaging system 700 utilizes a chatbot system 710, such as chatbot system 300 of FIG. 3A, to generate messages that are communicated to a user of the interactive platform via a client system 756 in order to adapt the chatbot system 710 to an existing platform architecture.

[0210] In some examples, the chatbot system 710 is modeled as a friend and receives and sends chats via a messaging backend system in a format of a Messaging Communication Service (MCS) 704. The MCS 704 powers the chatbot backend 702.

[0211] In some examples, the chatbot system 710 is pre-trained to feel like a real user in respects to how it will “read” chat messages, and show presence and typing indicators when interacting with users on the platform.

[0212] In some examples, the chatbot system 710 responds to text messages, stickers, platform submissions, and the like.

[0213] In some examples, the platform chatbot messaging system 700 is modular in construction and allow experimentation and testing with various different ML models.

[0214] In some examples, the platform chatbot messaging system 700 provides a flexible moderation system that detects certain sensitive topics and reacts accordingly as more fully described in reference to FIG. 6. In some examples, the platform chatbot messaging system 700 is extended to cover multiple languages.

[0215] In some examples, the platform chatbot messaging system 700 stores various inputs and output parameters into a hot storage 706 as short term storage and a cold storage 708 as long term storage that are useful for training a generative model such as, LLM 338 of FIG. 3B, in order to improve the chatbot system 710 over time.

[0216] In some examples, the chatbot backend 702 receives new requests forwarded by MCS 704, keeping track of conversation history, and formulates prompts that are communicated to the chatbot system 710.

[0217] In some examples, the platform chatbot messaging system 700 generates a stateful conversation based on conversation histories stored in the hot storage 706 and the cold storage 708 as part of a session history 712. Much like a real user would have their conversation history visible to them as they type messages to a friend, the chatbot system 710 accesses previous messages in a conversation in order to understand the context for the most recent message.

[0218] In some examples, the chatbot system 710 may not be re-trained based on previous conversations. As such, the chatbot system 710 does not “remember” previous messages in the conversation in the event the conversation is set to be deleted after viewing by a user. In order to replicate the concept of “memory” for the chatbot system 710, the platform chatbot messaging system 700 retains a copy of each message regardless of a message deletion policy that is applied to the conversation. For example, the platform chatbot messaging system 700 retains a specified number of messages, or retains messages based on a specified time period (e.g., 10 messages and/or 30 minutes of conversation).

[0219] In some examples, if a user explicitly deletes a message from a chat (via a long press or the like), the platform chatbot messaging system 700 removes that message from the conversation history, since the user has instructed the interactive platform hosting the platform chatbot messaging system 700 to delete the message.

[0220] In some examples, chatbot backend 702 uses the hot storage 706 to store conversation data such as, but not limited to: A UserID of the user in the conversation; each received and sent message; a timestamp for each message; metadata about which model, parameters were used for generating the response; and keywords that are extracted from each message which could aid in determining conversation context.

[0221] In some examples, after a message is no longer relevant to a current conversation, the chatbot backend 702 moves that data to long term storage in a cold storage 708 location that facilitates offline processing of that data. In some examples, the hot storage 706 and cold storage 708 are used as data sources for pre-training and fine-tuning an LLM of the chatbot system 710.

[0222] In some examples, the stored data is used for: ingestion by a user profile service to augment the platform's understanding of the user's interests; used by discover or monetization platform teams for aid in ranking; using previous conversations to learn the user's interests, preferences, and leverage in future conversations to make the chatbot system 710 more personalized.

[0223] In some examples, when a user sends in a user message 748, as a preliminary filter, the chatbot backend 702 will perform keyword detection for any subject matter that the platform considers sensitive. In some examples, the chatbot backend 702 includes a set of platform policies for content filtering that is based on geography and/or local legal systems.

[0224] In some examples, the platform chatbot messaging system 700 does not store conversations on specified topics.

[0225] In some examples, the chatbot backend 702 provides for message rate limiting. For example, a per conversation rate limit set is set to a specified number of messages per minute that may avoid rapid spamming attacks. When a conversation rate exceeds the rate limit, the platform chatbot messaging system 700 responds by messaging a user that the user has exceeded the conversation rate limit and performs no further processing the user messages sent by the user.

[0226] In some examples, as a way to put an upper cap on cost per user per day, the chatbot backend 702 service maintains a per day rate limit of a number of messages per day.

[0227] In some examples, the chatbot backend 702 caches similar messages that are then communicated to the chatbot system 710 in order to minimize operations of the chatbot system 710 and reduce a cost of providing or utilizing a service of the chatbot system 710.

[0228] In some examples, the chatbot backend 702 attaches metadata to a user message 748 or chatbot system 710 chatbot message 762. For example, a specified number of previous messages in a conversation can be sent to the chatbot system 710 as part of a user prompt. This provides the chatbot system 710 a context on questions that are “follow ups” to previous topics. For example, if a user asks for restaurants in LA, then says “what about bars”, the chatbot system 710 leverages the location data from previous question to suggest bars in LA.

[0229] In some examples, the platform chatbot messaging system 700 pre-trains the chatbot system 710 on details of the platform.

[0230] FIG. 7B is network diagram of a platform chatbot messaging system network 746, FIG. 7C is a collaboration diagram of communications within the platform chatbot messaging system 700, and FIG. 7D is an activity diagram of a chatbot messaging method 778, according to some examples. An interactive platform uses the platform chatbot messaging system 700 to provide an asynchronous architecture for communicating messages between a chatbot system 710 and one or more users.

Although the chatbot messaging method 778 depicts a particular sequence of operations, the sequence may be altered without departing from the scope of the present disclosure. For example, some of the operations depicted may be performed in parallel or in a different sequence that does not materially affect the function of the routine. In other examples, different components of a platform or system that implements the method may perform operations at substantially the same time or in a specific sequence.

[0231] Users, such as user 1 726, user 2 724, and user 3 722, communicate with the chatbot system 710 of an interactive platform, such as chatbot system 300 of FIG. 3B, via a Message Communication Service (MCS) 704 using respective client systems, such as client 1 720, client 2 718, and client 3 714. The client systems communicate via respective gRPC streams gRPC stream 732, gRPC stream 740, and gRPC stream 742

using a duplex service 716. The duplex service 716 communicates with a MCS server 734 via HTTP POST requests 738 to a send transient message endpoint of the server 734. The server 734 calls chatbot system 710 and uses broadcast transient message function 744 to send messages back to the clients. Server 734 also communicates with Network User Buffer (NUB) service 730 by calling a send to users function 736. NUB service 730 then uses a NUB 728 to route messages to the appropriate clients and the chatbot system 710.

[0232] In operation 766, a client system 756 sends a user message 748 to an MCS 704. For example, a user uses the client system 756 to send a user message 748 comprising text, images, video or other data to the MCS 704 over a network. The client system 756 communicates using a gRPC stream via a duplex service 716.

[0233] In operation 768, the MCS 704 writes the user message 748 to a message database 750 and sends the user message 748 to the chatbot system 710. For example, the MCS 704 writes the user message 748 received from the client system to a message database 750. The MCS 704 also sends the user message 748 to the chatbot system 710.

[0234] In operation 770, the chatbot system 710 shows a chatbot presence 758 of the chatbot system 710 on the client system 756 using a presence server 752. For example, the chatbot system 710 shows a chatbot presence 758 on the client system 756 using a presence server 752. This indicates to the client system 756 that the chatbot system 710 is present and ready to receive and respond to user messages 748.

[0235] In operation 772, the chatbot system 710 fetches a past message context 764 from the MCS 704. This allows the chatbot system 710 to retrieve previous messages in a conversation from the MCS 704 in order to understand the context for the current user message 748 from the user.

[0236] In operation 774, the chatbot system 710 communicates with a chatbot/LLM service 754. For example, the chatbot system 710 generates a prompt using the user message 748 and communicates the prompt to the chatbot/LLM service 754 as part of conversation 760. The chatbot/LLM service 754 receives the prompt as and generates a response using the prompt as part of the conversation 760. In some examples, the chatbot system 710 engages in a conversation 760 with the chatbot/LLM service 754 as more fully described in reference to FIG. 3B. This allows the chatbot system 710 to use the capabilities of the chatbot/LLM service 754 for natural language processing to

understand the user message 748 and generate an appropriate chatbot response for conducting a human-like conversation.

[0237] In some examples, the chatbot/LLM service 754 service may be hosted on the same system as the chatbot system 710 or on a separate server that the chatbot system 710 communicates with over a network. If hosted separately, the chatbot system 710 can send the user prompt to the chatbot/LLM service 754 and receive back a raw response generated by the chatbot/LLM service 754.

[0238] In operation 776, the chatbot system 710 sends a chatbot message 762 to the client system 756 via the MCS 704. For example, the chatbot system 710 generates the chatbot message 762 using the conversation 760. The chatbot system 710 communicates the chatbot message 762 to the client system 756 via the MCS 704. The client system 756 receives the chatbot message 762 and provides the chatbot message 762 to the use in a display on the client system 756.

[0239] In some examples, the MCS 704 expects client systems to model real human behavior of interacting with an application of an interactive platform. The MCS 704 forwards messages to client systems in real-time on a best-effort basis, but client systems are not necessarily reachable, so client systems periodically synchronize the interactive platform feeds and conversations a user participates in. This typically happens when a user places an application, such application 106 of a client system 102, into an execution foreground or swipes into a conversation view on their client system.

[0240] In some examples, the chatbot system 710 may not follow the same behavioral patterns as human users, so messages forwarded to a chatbot system 710 might be lost in transit if the real-time connection to the chatbot system 710 fails, during network partitions, or during service outages. The chatbot system 710 may not know the optimal time to synchronize conversations if the chatbot system 710 misses incoming events and is forced to synchronize periodically. In some examples, a chatbot system 710 which serves millions of users inherits this problem and presents a scaling challenge—periodically synchronizing millions of conversations on a cadence measured in seconds may increase costs and present availability risks to the MCS 704.

[0241] In some examples, conversations with the chatbot system 710 also have special User eXperience (UX) where certain chrome (e.g., UI elements such as, but not limited to, frames, buttons, icons, visual styling, layout, branding elements, navigational elements, headers, footers, sidebars, and the like) is hidden and other chrome may

behave differently. Accordingly, interactive platform mobile and web clients recognize these conversations and treat them differently than typical 1:1 conversations between users.

[0242] In some examples, the chatbot system 710 is supported as a first-class concept within the interactive platform. For example, the chatbot system 710 is presented to users of the interactive platform on the same level as other users and messages within the interactive platform that are communicated by the chatbot system 710 are treated as having the same priority as messages from other users.

[0243] In some examples, the platform chatbot messaging system 700 forwards messages and events from the MCS 704 to a chatbot gRPC interface that is agnostic as to a chatbot-like service.

[0244] In some examples, the platform chatbot messaging system 700 uses a platform task service to abstract away durable queuing of messages/events.

[0245] In some examples, the platform chatbot messaging system 700 provides for APIs that send messages and manage the message lifecycle.

[0246] In some examples, the platform chatbot messaging system 700 provides access an API that issues presence heartbeats to show chatbot presence in an interactive platform application hosted by a client system 756.

[0247] In some examples, the platform chatbot messaging system 700 employs conversation-level rate limiting so users cannot send an unreasonable number of messages that exceed a threshold limit per a unit of time. In some examples, the unit of time is a minute. In some examples, the unit of time is 24 hours.

[0248] In some examples, the platform chatbot messaging system 700 employs rate limiting for spam and abuse: if a user exceeds a rate limit threshold for a conversation, the MCS 704 will fail the message send. In some examples, the platform chatbot messaging system 700 provides for chatbot-specific rate limits with specific responses—such as responding with a predefined message.

[0249] In some examples, the platform chatbot messaging system 700 employs a conversation subtype to indicate a chatbot message in an existing messaging system.

[0250] In some examples, the platform chatbot messaging system 700 pins a chatbot conversation to a top position of a feed if enabled, and unpins the conversation from the feed when disabled. Once the chatbot conversation is created, the chatbot conversation behaves like other conversations. After unpinning, the chatbot conversation will drop in

friends feed as other conversations take higher priority. In some examples, the platform chatbot messaging system 700 enables the feature on when the platform chatbot messaging system 700 detects that a user enables the feed item and disables the feature when the platform chatbot messaging system 700 detects that the user disables the feature.

[0251] In some examples, when entering a chatbot conversation for the first time, the platform chatbot messaging system 700 shows a localized prompt that explains to the user what this conversation is and expectations around using it—effectively terms of service. In some examples, the platform chatbot messaging system 700 checks whether the user has already accepted the dialog and does not display the dialog in response to detecting that the user has already accepted the dialog.

[0252] In some examples, the chrome for the conversation view is different for chatbots. For chatbot conversations, the layout of the conversation is changed relative to a conversation of a human user such as by hiding calling and game buttons. In addition, a circle profile icon for the chatbot conversation will not load to a profile and will instead load an alternate modal screen.

[0253] In some examples, data communicated during a chatbot conversation includes, but is not limited to: content (in bytes); input for a chatbot component; a content type (the MCS 704 filters by the content type on what to forward, or the chatbot system 710 can determine whether they care about the content type); Message ID (users can update/delete their messages); Conversation ID (this is so the platform knows which conversation to send the chatbot message back to allowing group conversations); and Sender ID (so the platform chatbot messaging system 700 knows who the message sender is. The service generating a response may look up the user data of the message sender (age, location etc.) to further curate the responses).

[0254] In some examples, a platform chatbot messaging system 700 does not allow chatbot messages to be forwarded.

[0255] In some examples, data sent in an update content message includes: Message ID; Conversation ID; Actor ID; and an update action type. In some examples, the platform chatbot messaging system 700 filters by the update action type on the outbound forward as a limited number of fields are useful for chatbots, such as a Reaction Request prompting a user to provide user intent via an upvote, downvote, or the like.

[0256] In some examples, the platform chatbot messaging system 700 supports a user deleting their messages.

[0257] In some examples, a client sends messages to a duplex service, and the duplex service routes the message to a presence service. Then presence service distributes the message to multiple users through a platform messaging service such as, but not limited to a NUB. NUB calls use a duplex service to deliver presence information to client.

[0258] In some examples, the platform chatbot messaging system 700 collects a variety of messaging and conversation metrics as session history 712.

[0259] In some examples, health monitoring aims to respond quickly to possible shutdowns and control the system loading. As the chatbot system 710 may pay per token when accessing an external chatbot/LLM service 754, the platform chatbot messaging system 700 uses alerts on sudden growth to avoid unexpected expenses. To ensure quick reaction the platform chatbot messaging system 700 supports collecting and storing health metrics in a real-time monitoring system. The list of metrics includes, but is not limited to, :

[0260] • Number of shows of the chatbot system 710 in the chat feed

[0261] • Number of started chats

[0262] • Number of sent tokens in total

[0263] • Number of received tokens in total

[0264] • Number of messages sent to the API endpoint

[0265] • Number of messages received from the API endpoint

[0266] • Number of users who have reached the rate limit

[0267] • Response times such as percentile responses times P25, P50, P90, and P99.

[0268] In some examples, the platform chatbot messaging system 700 collects metrics for comparing different models, model versions, and design changes. Product metrics cover technical performance, retention, conversion rate, and safety metrics. The list of metrics includes, but is not limited to:

- Number of Sends, Session, Senders.
- Conversation Length, the number of messages per chat session.
- Open rate, the number of chatbot component opens per total number of shows of the chatbot component in the chat feed.

- Start rate, the number of sessions with at least one message per total number of sessions.
- Prompt using rate, the number of using prompts per total number of sent messages.
- Bounce Rate, the number of sessions with <X messages per total number of sessions.
- Retention rate is the number of users who are retained in the next X days.
- Satisfying rate is the number of sessions with positive feedback about platform chatbot component answer per total number of questions.
- Reaction count is the number of bitmoji reactions on platform chatbot component answers.
- The number of forwarded messages.
- Sessions per user.

[0269] In some examples, the platform chatbot messaging system 700 collects performance metrics as well as health metrics, the platform use them to evaluate the health of the service and API endpoint responsiveness. The platform has the ability to use them for deep analysis and AB tests. In addition to the health metrics, the platform chatbot messaging system 700 collects metrics including, but not limited to:

- Distribution of number of tokens per message sent to API endpoint.
- Distribution of number of tokens per reply received from API endpoint.
- Miss rate, the number of received messages per number of sent messages.

[0270] In some examples, the 700 collects a safety metric as a report rate of a number of reports on chatbot component messages per total number of chatbot component messages.

[0271] In some examples, in order to leverage platform chatbot component quality and fine-tune for users needs, the platform chatbot messaging system 700 collects and analyzes messages content. The insights that the platform determines includes, but are not limited to:

- Topics.
- Engagement per topics.
- Platform related questions.
- Positive feedback.
- Negative feedback.

- Prompt quality.
- Sensitive topics.
- Harmful and abusive information.

[0272] In some examples, to determine the quality of the conversation, as well as the markup of topics, the platform chatbot messaging system 700 can use the following methods:

- Keywords - the platform uses a dataset of key words to identify dangerous and inappropriate content, positive feedback (e.g. thanks, wow, awesome), negative feedback (e.g. not exactly, wrong), and topics (e.g. bitmoji, movie, music).
- User feedback - for collecting user feedback the platform uses native chat interface such as reports, bitmoji reactions, forwarding, and direct questions from the chatbot component (e.g., do you satisfy by the answer?).
- Data Markup.

[0273] In some examples, the platform chatbot messaging system 700 detects a set of events that enable the platform chatbot messaging system 700 to collect meaningful metrics. Events include, but are not limited to:

- ChatbotChatCreate - the event reporting to track when a user opens any chat room. The event contains CorrespondentId. The event is used for the OpenRate.
- ChatbotChatSend - the main event to track when a user sends a message to the chatbot component and the chatbot component sends an answer. For the chatbot component, the platform triggers the event on the backend side. The event body contains UserID, CorrespondentId, MessageId, MessageType (text, reaction, media). The event is used for engagement metrics such as the count of sends, conversation length, reaction count and the like.
- ChatbotChatView - an event to track when a user viewed a message. The event contains UserId, CorrespondentId, MessageId, MessageType (text, reaction, media).
- ChatbotChatShow - an event to track the number of impressions of the chat with the chatbot component on the friend feed page. The event is used for the Open Rate calculation.
- ChatbotMessageSend - an event to calculate number of messages sent to API endpoint. the platform will send in this event a MessageId, number of tokens, message type (text, reaction, media), status, is this prompt, CorrespondentId and message sending timestamp. The event is used to calculate distribution of number of

tokens per message sent to API endpoint, number of requests, prompt usage.

Together with the ChatbotMessageReceive used to calculate the waiting time.

- ChatbotMessageReceive - an event to calculate number of messages received from API endpoint. the platform will send in this MessageId, event a number of tokens, message type (if the platform support text-to-media models), status, CorrespondentId, response waiting time, and message receiving timestamp. The event is used to calculate distribution of number of tokens per message received from API endpoint, number of requests, endpoint API health. Together with the ChatbotMessageSend used to calculate the waiting time.
- ChatbotChatReport - the number of reports on chatbot messages. The event aimed to leverage Safety and exclude edgy answers.
- ChatbotShows - the number of impressions of the chat with the chatbot component on the friend feed page.
- ChatbotMessageSend - the number of chats sent to the endpoint. The main event to track the system health in real-time. Dimension: the order number of messages in the session and amount of tokens in this message.
- ChatbotMessageReceive - the number of chats received from the endpoint. The event to track the API Endpoint health in real-time. Dimension: amount of tokens in this message.
- ChatbotResponseTime - the time between a message being sent and an answer receiving.
- ChatbotLimitReached - the number of blocked messages by reason of reaching limits. Dimension: the type of limit that was exceeded.

[0274] In some examples, once a session has been completed (either based on idle duration or reaching max allowed length), the session will be “ended” and a new session will be created the next time the user sends a text message. When a session is completed, the chatbot backend 702 will move the session metadata from hot storage 706 to cold storage 708. This will allow consumption of this data for analysis and further ML training.

[0275] In some examples, a MessageId for each message is used to join saved content with events to restore a full picture of the conversation. After collecting raw data, the platform chatbot messaging system 700 breaks messages into topics, determine which conversions ended successfully and classify the types of user questions.

MACHINE-LEARNING PIPELINE

[0276] FIG. 9 is a flowchart depicting a machine-learning pipeline 900, according to some examples. The machine-learning pipeline 900 may be used to generate a trained machine-learning program 902, such as creative generation component 356 (of FIG. 3B), LLM 338 (of FIG. 3B), intent processing component 348 (of FIG. 3B), and sensitive content filter component 604 (of FIG. 6) to perform operations associated with performing searches and generating query responses.

Overview

[0277] Broadly, machine learning may involve using computer algorithms to automatically learn patterns and relationships in data, potentially without the need for explicit programming. Machine learning algorithms can be divided into three main categories: supervised learning, unsupervised learning, and reinforcement learning.

- Supervised learning involves training a model using labeled data to predict an output for new, unseen inputs. Examples of supervised learning algorithms include linear regression, decision trees, and neural networks.
- Unsupervised learning involves training a model on unlabeled data to find hidden patterns and relationships in the data. Examples of unsupervised learning algorithms include clustering, principal component analysis, and generative models like autoencoders.
- Reinforcement learning involves training a model to make decisions in a dynamic environment by receiving feedback in the form of rewards or penalties. Examples of reinforcement learning algorithms include Q-learning and policy gradient methods.

[0278] Examples of specific machine learning algorithms that may be deployed, according to some examples, include logistic regression, which is a type of supervised learning algorithm used for binary classification tasks. Logistic regression models the probability of a binary response variable based on one or more predictor variables. Another example type of machine learning algorithm is Naïve Bayes, which is another supervised learning algorithm used for classification tasks. Naïve Bayes is based on Bayes' theorem and assumes that the predictor variables are independent of each other. Random Forest is another type of supervised learning algorithm used for classification, regression, and other tasks. Random Forest builds a collection of decision trees and combines their outputs to make predictions. Further examples include neural networks, which consist of interconnected layers of nodes (or neurons) that process information

and make predictions based on the input data. Matrix factorization is another type of machine learning algorithm used for recommender systems and other tasks. Matrix factorization decomposes a matrix into two or more matrices to uncover hidden patterns or relationships in the data. Support Vector Machines (SVM) are a type of supervised learning algorithm used for classification, regression, and other tasks. SVM finds a hyperplane that separates the different classes in the data. Other types of machine learning algorithms include decision trees, k-nearest neighbors, clustering algorithms, and deep learning algorithms such as convolutional neural networks (CNN), recurrent neural networks (RNN), and transformer models. The choice of algorithm depends on the nature of the data, the complexity of the problem, and the performance requirements of the application.

[0279] The performance of machine learning models is typically evaluated on a separate test set of data that was not used during training to ensure that the model can generalize to new, unseen data.

[0280] Although several specific examples of machine learning algorithms are discussed herein, the principles discussed herein can be applied to other machine learning algorithms as well. Deep learning algorithms such as convolutional neural networks, recurrent neural networks, and transformers, as well as more traditional machine learning algorithms like decision trees, random forests, and gradient boosting may be used in various machine learning applications.

[0281] Three example types of problems in machine learning are classification problems, regression problems, and generation problems. Classification problems, also referred to as categorization problems, aim at classifying items into one of several category values (for example, is this object an apple or an orange?). Regression algorithms aim at quantifying some items (for example, by providing a value that is a real number). Generation algorithms aim at producing new examples that are similar to examples provided for training. For instance, a text generation algorithm is trained on many text documents and is configured to generate new coherent text with similar statistical properties as the training data.

Training phases 904

[0282] Generating a trained machine-learning program 902 may include multiple phases that form part of the machine-learning pipeline 900, including for example the following phases illustrated in FIG. 8:

- **Data collection and preprocessing 802:** This phase may include acquiring and cleaning data to ensure that it is suitable for use in the machine learning model. This phase may also include removing duplicates, handling missing values, and converting data into a suitable format.
- **Feature engineering 804:** This phase may include selecting and transforming the training data 906 to create features that are useful for predicting the target variable. Feature engineering may include (1) receiving features 908 (e.g., as structured or labeled data in supervised learning) and/or (2) identifying features 908 (e.g., unstructured or unlabeled data for unsupervised learning) in training data 906.
- **Model selection and training 806:** This phase may include selecting an appropriate machine learning algorithm and training it on the preprocessed data. This phase may further involve splitting the data into training and testing sets, using cross-validation to evaluate the model, and tuning hyperparameters to improve performance.
- **Model evaluation 808:** This phase may include evaluating the performance of a trained model (e.g., the trained machine-learning program 902) on a separate testing dataset. This phase can help determine if the model is overfitting or underfitting and determine whether the model is suitable for deployment.
- **Prediction 810:** This phase involves using a trained model (e.g., trained machine-learning program 902) to generate predictions on new, unseen data.
- **Validation, refinement or retraining 812:** This phase may include updating a model based on feedback generated from the prediction phase, such as new data or user feedback.
- **Deployment 814:** This phase may include integrating the trained model (e.g., the trained machine-learning program 902) into a more extensive system or application, such as a web service, mobile app, or IoT device. This phase can involve setting up APIs, building a user interface, and ensuring that the model is scalable and can handle large volumes of data.

[0283] FIG. 9 illustrates further details of two example phases, namely a training phase 904 (e.g., part of the model selection and trainings 806) and a prediction phase 910 (part of prediction 810). Prior to the training phase 904, feature engineering 804 is used to identify features 908. This may include identifying informative, discriminating, and independent features for effectively operating the trained machine-learning program 902

in pattern recognition, classification, and regression. In some examples, the training data 906 includes labeled data, known for pre-identified features 908 and one or more outcomes. Each of the features 908 may be a variable or attribute, such as an individual measurable property of a process, article, system, or phenomenon represented by a data set (e.g., the training data 906). Features 908 may also be of different types, such as numeric features, strings, and graphs, and may include one or more of content 912, concepts 914, attributes 916, historical data 918, and/or user data 920, merely for example.

[0284] In training phase 904, the machine-learning pipeline 900 uses the training data 906 to find correlations among the features 908 that affect a predicted outcome or prediction/inference data 922.

[0285] With the training data 906 and the identified features 908, the trained machine-learning program 902 is trained during the training phase 904 during machine-learning program training 924. The machine-learning program training 924 appraises values of the features 908 as they correlate to the training data 906. The result of the training is the trained machine-learning program 902 (e.g., a trained or learned model).

[0286] Further, the training phase 904 may involve machine learning, in which the training data 906 is structured (e.g., labeled during preprocessing operations). The trained machine-learning program 902 implements a neural network 926 capable of performing, for example, classification and clustering operations. In other examples, the training phase 904 may involve deep learning, in which the training data 906 is unstructured, and the trained machine-learning program 902 implements a deep neural network 926 that can perform both feature extraction and classification/clustering operations.

[0287] In some examples, a neural network 926 may be generated during the training phase 904, and implemented within the trained machine-learning program 902. The neural network 926 includes a hierarchical (e.g., layered) organization of neurons, with each layer consisting of multiple neurons or nodes. Neurons in the input layer receive the input data, while neurons in the output layer produce the final output of the network. Between the input and output layers, there may be one or more hidden layers, each consisting of multiple neurons.

[0288] Each neuron in the neural network 926 operationally computes a function, such as an activation function, which takes as input the weighted sum of the outputs of the

neurons in the previous layer, as well as a bias term. The output of this function is then passed as input to the neurons in the next layer. If the output of the activation function exceeds a certain threshold, an output is communicated from that neuron (e.g., transmitting neuron) to a connected neuron (e.g., receiving neuron) in successive layers. The connections between neurons have associated weights, which define the influence of the input from a transmitting neuron to a receiving neuron. During the training phase, these weights are adjusted by the learning algorithm to optimize the performance of the network. Different types of neural networks may use different activation functions and learning algorithms, affecting their performance on different tasks. The layered organization of neurons and the use of activation functions and weights enable neural networks to model complex relationships between inputs and outputs, and to generalize to new inputs that were not seen during training.

[0289] In some examples, the neural network 926 may also be one of several different types of neural networks, such as a single-layer feed-forward network, a Multilayer Perceptron (MLP), an Artificial Neural Network (ANN), a Recurrent Neural Network (RNN), a Long Short-Term Memory Network (LSTM), a Bidirectional Neural Network, a symmetrically connected neural network, a Deep Belief Network (DBN), a Convolutional Neural Network (CNN), a Generative Adversarial Network (GAN), an Autoencoder Neural Network (AE), a Restricted Boltzmann Machine (RBM), a Hopfield Network, a Self-Organizing Map (SOM), a Radial Basis Function Network (RBFN), a Spiking Neural Network (SNN), a Liquid State Machine (LSM), an Echo State Network (ESN), a Neural Turing Machine (NTM), or a Transformer Network, merely for example.

[0290] In addition to the training phase 904, a validation phase may be performed on a separate dataset known as the validation dataset. The validation dataset is used to tune the hyperparameters of a model, such as the learning rate and the regularization parameter. The hyperparameters are adjusted to improve the model's performance on the validation dataset.

[0291] Once a model is fully trained and validated, in a testing phase, the model may be tested on a new dataset. The testing dataset is used to evaluate the model's performance and ensure that the model has not overfitted the training data.

[0292] In prediction phase 910, the trained machine-learning program 902 uses the features 908 for analyzing query data 928 to generate inferences, outcomes, or

predictions, as examples of a prediction/inference data 922. For example, during prediction phase 910, the trained machine-learning program 902 generates an output. Query data 928 is provided as an input to the trained machine-learning program 902, and the trained machine-learning program 902 generates the prediction/inference data 922 as output, responsive to receipt of the query data 928.

[0293] In some examples, the trained machine-learning program 902 may be a generative AI model. Generative AI is a term that may refer to any type of artificial intelligence that can create new content from training data 906. For example, generative AI can produce text, images, video, audio, code, or synthetic data similar to the original data but not identical.

[0294] Some of the techniques that may be used in generative AI are:

- Convolutional Neural Networks (CNNs): CNNs may be used for image recognition and computer vision tasks. CNNs may, for example, be designed to extract features from images by using filters or kernels that scan the input image and highlight important patterns.
- Recurrent Neural Networks (RNNs): RNNs may be used for processing sequential data, such as speech, text, and time series data, for example. RNNs employ feedback loops that allow them to capture temporal dependencies and remember past inputs.
- Generative adversarial networks (GANs): GNNs may include two neural networks: a generator and a discriminator. The generator network attempts to create realistic content that can “fool” the discriminator network, while the discriminator network attempts to distinguish between real and fake content. The generator and discriminator networks compete with each other and improve over time.
- Variational autoencoders (VAEs): VAEs may encode input data into a latent space (e.g., a compressed representation) and then decode it back into output data. The latent space can be manipulated to generate new variations of the output data. VAEs may use self-attention mechanisms to process input data, allowing them to handle long text sequences and capture complex dependencies.
- Transformer models: Transformer models may use attention mechanisms to learn the relationships between different parts of input data (such as words or pixels) and generate output data based on these relationships. Transformer models can handle sequential data, such as text or speech, as well as non-sequential data, such as images or code.

[0295] In generative AI examples, the query data 928 may include text, audio, image, video, numeric, or media content prompts and the output inference/prediction data 722 includes text, images, video, audio, code, or synthetic data.

DATA ARCHITECTURE

[0296] FIG. 10 is a schematic diagram illustrating data structures 1000, which may be stored in the database 1004 of the interaction server system 110, according to certain examples. While the content of the database 1004 is shown to comprise multiple tables, it will be appreciated that the data could be stored in other types of data structures (e.g., as an object-oriented database).

[0297] The database 1004 includes message data stored within a message table 1006. This message data includes, for any particular message, at least message sender data, message recipient (or receiver) data, and a payload. Further details regarding information that may be included in a message, and included within the message data stored in the message table 1006, are described below with reference to FIG. 10.

[0298] An entity table 1008 stores entity data, and is linked (e.g., referentially) to an entity graph 1010 and profile data 1002. Entities for which records are maintained within the entity table 1008 may include individuals, corporate entities, organizations, objects, places, events, and so forth. Regardless of entity type, any entity regarding which the interaction server system 110 stores data may be a recognized entity. Each entity is provided with a unique identifier, as well as an entity type identifier (not shown).

[0299] The entity graph 1010 stores information regarding relationships and associations between entities. Such relationships may be social, professional (e.g., work at a common corporation or organization), interest-based, or activity-based, merely for example. Certain relationships between entities may be unidirectional, such as a subscription by an individual user to digital content of a commercial or publishing user (e.g., a newspaper or other digital media outlet, or a brand). Other relationships may be bidirectional, such as a “friend” relationship between individual users of the interaction system 100.

[0300] Certain permissions and relationships may be attached to each relationship, and also to each direction of a relationship. For example, a bidirectional relationship (e.g., a friend relationship between individual users) may include authorization for the publication of digital content items between the individual users, but may impose certain restrictions or filters on the publication of such digital content items (e.g., based on

content characteristics, location data or time of day data). Similarly, a subscription relationship between an individual user and a commercial user may impose different degrees of restrictions on the publication of digital content from the commercial user to the individual user, and may significantly restrict or block the publication of digital content from the individual user to the commercial user. A particular user, as an example of an entity, may record certain restrictions (e.g., by way of privacy settings) in a record for that entity within the entity table 1008. Such privacy settings may be applied to all types of relationships within the context of the interaction system 100, or may selectively be applied to only certain types of relationships.

[0301] The profile data 1002 stores multiple types of profile data about a particular entity. The profile data 1002 may be selectively used and presented to other users of the interaction system 100 based on privacy settings specified by a particular entity. Where the entity is an individual, the profile data 1002 includes, for example, a username, telephone number, address, settings (e.g., notification and privacy settings), as well as a user-selected avatar representation (or collection of such avatar representations). A particular user may then selectively include one or more of these avatar representations within the content of messages communicated via the interaction system 100, and on map interfaces displayed by interaction clients 104 to other users. The collection of avatar representations may include “status avatars,” which present a graphical representation of a status or activity that the user may select to communicate at a particular time.

[0302] Where the entity is a group, the profile data 1002 for the group may similarly include one or more avatar representations associated with the group, in addition to the group name, members, and various settings (e.g., notifications) for the relevant group.

[0303] The database 1004 also stores augmentation data, such as overlays or filters, in an augmentation table 1012. The augmentation data is associated with and applied to videos (for which data is stored in a video table 1014) and images (for which data is stored in an image table 1016).

[0304] Filters, in some examples, are overlays that are displayed as overlaid on an image or video during presentation to a message receiver. Filters may be of various types, including user-selected filters from a set of filters presented to a message sender by the interaction client 104 when the message sender is composing a message. Other types of filters include geolocation filters (also known as geo-filters), which may be

presented to a message sender based on geographic location. For example, geolocation filters specific to a neighborhood or special location may be presented within a user interface by the interaction client 104, based on geolocation information determined by a Global Positioning System (GPS) unit of the client system 102.

[0305] Another type of filter is a data filter, which may be selectively presented to a message sender by the interaction client 104 based on other inputs or information gathered by the client system 102 during the message creation process. Examples of data filters include current temperature at a specific location, a current speed at which a message sender is traveling, battery life for a client system 102, or the current time.

[0306] Other augmentation data that may be stored within the image table 1016 includes augmented reality content items (e.g., corresponding to applying Lenses or augmented reality experiences). An augmented reality content item may be a real-time special effect and sound that may be added to an image or a video.

[0307] As described above, augmentation data includes augmented reality (AR), virtual reality (VR) and mixed reality (MR) content items, overlays, image transformations, images, and modifications that may be applied to image data (e.g., videos or images). This includes real-time modifications, which modify an image as it is captured using device sensors (e.g., one or multiple cameras) of the client system 102 and then displayed on a screen of the client system 102 with the modifications. This also includes modifications to stored content, such as video clips in a collection or group that may be modified. For example, in a client system 102 with access to multiple augmented reality content items, a user can use a single video clip with multiple augmented reality content items to see how the different augmented reality content items will modify the stored clip. Similarly, real-time video capture may use modifications to show how video images currently being captured by sensors of a client system 102 would modify the captured data. Such data may simply be displayed on the screen and not stored in memory, or the content captured by the device sensors may be recorded and stored in memory with or without the modifications (or both). In some systems, a preview feature can show how different augmented reality content items will look within different windows in a display at the same time. This can, for example, enable multiple windows with different pseudorandom animations to be viewed on a display at the same time.

[0308] Data and various systems using augmented reality content items or other such transform systems to modify content using this data can thus involve detection of objects

(e.g., faces, hands, bodies, cats, dogs, surfaces, objects, etc.), tracking of such objects as they leave, enter, and move around the field of view in video frames, and the modification or transformation of such objects as they are tracked. In various examples, different methods for achieving such transformations may be used. Some examples may involve generating a three-dimensional mesh model of the object or objects, and using transformations and animated textures of the model within the video to achieve the transformation. In some examples, tracking of points on an object may be used to place an image or texture (which may be two-dimensional or three-dimensional) at the tracked position. In still further examples, neural network analysis of video frames may be used to place images, models, or textures in content (e.g., images or frames of video). Augmented reality content items thus refer both to the images, models, and textures used to create transformations in content, as well as to additional modeling and analysis information needed to achieve such transformations with object detection, tracking, and placement.

[0309] Real-time video processing can be performed with any kind of video data (e.g., video streams, video files, etc.) saved in a memory of a computerized system of any kind. For example, a user can load video files and save them in a memory of a device, or can generate a video stream using sensors of the device. Additionally, any objects can be processed using a computer animation model, such as a human's face and parts of a human body, animals, or non-living things such as chairs, cars, or other objects.

[0310] In some examples, when a particular modification is selected along with content to be transformed, elements to be transformed are identified by the computing device, and then detected and tracked if they are present in the frames of the video. The elements of the object are modified according to the request for modification, thus transforming the frames of the video stream. Transformation of frames of a video stream can be performed by different methods for different kinds of transformation. For example, for transformations of frames mostly referring to changing forms of an object's elements, characteristic points for each element of an object are calculated (e.g., using an Active Shape Model (ASM) or other known methods). Then, a mesh based on the characteristic points is generated for each element of the object. This mesh is used in the following stage of tracking the elements of the object in the video stream. In the process of tracking, the mesh for each element is aligned with a position of each element. Then, additional points are generated on the mesh.

[0311] In some examples, transformations changing some areas of an object using its elements can be performed by calculating characteristic points for each element of an object and generating a mesh based on the calculated characteristic points. Points are generated on the mesh, and then various areas based on the points are generated. The elements of the object are then tracked by aligning the area for each element with a position for each of the at least one element, and properties of the areas can be modified based on the request for modification, thus transforming the frames of the video stream. Depending on the specific request for modification, properties of the mentioned areas can be transformed in different ways. Such modifications may involve changing the color of areas; removing some part of areas from the frames of the video stream; including new objects into areas that are based on a request for modification; and modifying or distorting the elements of an area or object. In various examples, any combination of such modifications or other similar modifications may be used. For certain models to be animated, some characteristic points can be selected as control points to be used in determining the entire state-space of options for the model animation.

[0312] In some examples of a computer animation model to transform image data using face detection, the face is detected on an image using a specific face detection algorithm (e.g., Viola-Jones). Then, an Active Shape Model (ASM) algorithm is applied to the face region of an image to detect facial feature reference points.

[0313] Other methods and algorithms suitable for face detection can be used. For example, in some examples, features are located using a landmark, which represents a distinguishable point present in most of the images under consideration. For facial landmarks, for example, the location of the left eye pupil may be used. If an initial landmark is not identifiable (e.g., if a person has an eyepatch), secondary landmarks may be used. Such landmark identification procedures may be used for any such objects. In some examples, a set of landmarks forms a shape. Shapes can be represented as vectors using the coordinates of the points in the shape. One shape is aligned to another with a similarity transform (allowing translation, scaling, and rotation) that minimizes the average Euclidean distance between shape points. The mean shape is the mean of the aligned training shapes.

[0314] A transformation system can capture an image or video stream on a client device (e.g., the client system 102) and perform complex image manipulations locally on the

client system 102 while maintaining a suitable user experience, computation time, and power consumption. The complex image manipulations may include size and shape changes, emotion transfers (e.g., changing a face from a frown to a smile), state transfers (e.g., aging a subject, reducing apparent age, changing gender), style transfers, graphical element application, and any other suitable image or video manipulation implemented by a convolutional neural network that has been configured to execute efficiently on the client system 102.

[0315] In some examples, a computer animation model to transform image data can be used by a system where a user may capture an image or video stream of the user (e.g., a selfie) using the client system 102 having a neural network operating as part of an interaction client 104 operating on the client system 102. The transformation system operating within the interaction client 104 determines the presence of a face within the image or video stream and provides modification icons associated with a computer animation model to transform image data, or the computer animation model can be present as associated with an interface described herein. The modification icons include changes that are the basis for modifying the user's face within the image or video stream as part of the modification operation. Once a modification icon is selected, the transform system initiates a process to convert the image of the user to reflect the selected modification icon (e.g., generate a smiling face on the user). A modified image or video stream may be presented in a graphical user interface displayed on the client system 102 as soon as the image or video stream is captured and a specified modification is selected. The transformation system may implement a complex convolutional neural network on a portion of the image or video stream to generate and apply the selected modification. That is, the user may capture the image or video stream and be presented with a modified result in real-time or near real-time once a modification icon has been selected. Further, the modification may be persistent while the video stream is being captured, and the selected modification icon remains toggled. Machine-taught neural networks may be used to enable such modifications.

[0316] The graphical user interface, presenting the modification performed by the transform system, may supply the user with additional interaction options. Such options may be based on the interface used to initiate the content capture and selection of a particular computer animation model (e.g., initiation from a content creator user interface). In various examples, a modification may be persistent after an initial selection of a modification icon. The user may toggle the modification on or off by

tapping or otherwise selecting the face being modified by the transformation system and store it for later viewing or browsing to other areas of the imaging application. Where multiple faces are modified by the transformation system, the user may toggle the modification on or off globally by tapping or selecting a single face modified and displayed within a graphical user interface. In some examples, individual faces, among a group of multiple faces, may be individually modified, or such modifications may be individually toggled by tapping or selecting the individual face or a series of individual faces displayed within the graphical user interface.

[0317] A story table 1018 stores data regarding collections of messages and associated image, video, or audio data, which are compiled into a collection (e.g., a story or a gallery). The creation of a particular collection may be initiated by a particular user (e.g., each user for which a record is maintained in the entity table 1008). A user may create a “personal story” in the form of a collection of content that has been created and sent/broadcast by that user. To this end, the user interface of the interaction client 104 may include an icon that is user-selectable to enable a message sender to add specific content to his or her personal story.

[0318] A collection may also constitute a “live story,” which is a collection of content from multiple users that is created manually, automatically, or using a combination of manual and automatic techniques. For example, a “live story” may constitute a curated stream of user-submitted content from various locations and events. Users whose client devices have location services enabled and are at a common location event at a particular time may, for example, be presented with an option, via a user interface of the interaction client 104, to contribute content to a particular live story. The live story may be identified to the user by the interaction client 104, based on his or her location. The end result is a “live story” told from a community perspective.

[0319] A further type of content collection is known as a “location story,” which enables a user whose client system 102 is located within a specific geographic location (e.g., on a college or university campus) to contribute to a particular collection. In some examples, a contribution to a location story may require a second degree of authentication to verify that the end-user belongs to a specific organization or other entity (e.g., is a student on the university campus).

[0320] As mentioned above, the video table 1014 stores video data that, in some examples, is associated with messages for which records are maintained within the

message table 1006. Similarly, the image table 1016 stores image data associated with messages for which message data is stored in the entity table 1008. The entity table 1008 may associate various augmentations from the augmentation table 1012 with various images and videos stored in the image table 1016 and the video table 1014.

[0321] The databases 1004 also include social network information collected by an interaction system of an interactive platform. The social network information may include without limitation relationship and communication data for users of the interactive platform. The social network information can be used to group two or more users and offer additional functionality of the interaction system 100. Examples of relationships include, but are not limited to, best friends relationships where two or more users are determined to be mutual best friends based on a frequency of their interactions, users who have common interests in current events, users who share an affiliation through social clubs or philanthropic organizations, and the like. Examples of communications include without limitation chats, private and public messages, exchanges of media such as images, videos, audio recordings, and the like.

DATA COMMUNICATIONS ARCHITECTURE

[0322] FIG. 11 is a schematic diagram illustrating a structure of a message 1100, according to some examples, generated by an interaction client 104 for communication to a further interaction client 104 via the interaction servers 124. The content of a particular message 1100 is used to populate the message table 1006 stored within the database 1004, accessible by the interaction servers 124. Similarly, the content of a message 1100 is stored in memory as “in-transit” or “in-flight” data of the client system 102 or the interaction servers 124. A message 1100 is shown to include the following example components:

- Message identifier 1102: a unique identifier that identifies the message 1100.
- Message text payload 1134: text, to be generated by a user via a user interface of the client system 102, and that is included in the message 1100.
- Message image payload 1104: image data, captured by a camera component of a client system 102 or retrieved from a memory component of a client system 102, and that is included in the message 1100. Image data for a sent or received message 1100 may be stored in the image table 1106.
- Message video payload 1108: video data, captured by a camera component or retrieved from a memory component of the client system 102, and that is included in

the message 1100. Video data for a sent or received message 1100 may be stored in the video table 1110.

- Message audio payload 1112: audio data, captured by a microphone or retrieved from a memory component of the client system 102, and that is included in the message 1100.
- Message augmentation data 1114: augmentation data (e.g., filters, stickers, or other annotations or enhancements) that represents augmentations to be applied to message image payload 1104, message video payload 1108, or message audio payload 1112 of the message 1100. Augmentation data for a sent or received message 1100 may be stored in the augmentation table 1116.
- Message duration parameter 1118: parameter value indicating, in seconds, the amount of time for which content of the message (e.g., the message image payload 1104, message video payload 1108, message audio payload 1112) is to be presented or made accessible to a user via the interaction client 104.
- Message geolocation parameter 1120: geolocation data (e.g., latitudinal and longitudinal coordinates) associated with the content payload of the message. Multiple message geolocation parameter 1120 values may be included in the payload, each of these parameter values being associated with respect to content items included in the content (e.g., a specific image within the message image payload 1104, or a specific video in the message video payload 1108).
- Message story identifier 1122: identifier values identifying one or more content collections (e.g., “stories” identified in the story table 1124) with which a particular content item in the message image payload 1104 of the message 1100 is associated. For example, multiple images within the message image payload 1104 may each be associated with multiple content collections using identifier values.
- Message tag 1126: each message 1100 may be tagged with multiple tags, each of which is indicative of the subject matter of content included in the message payload. For example, where a particular image included in the message image payload 1104 depicts an animal (e.g., a lion), a tag value may be included within the message tag 1126 that is indicative of the relevant animal. Tag values may be generated manually, based on user input, or may be automatically generated using, for example, image recognition.

- Message sender identifier 1128: an identifier (e.g., a messaging system identifier, email address, or device identifier) indicative of a user of the client system 102 on which the message 1100 was generated and from which the message 1100 was sent.
- Message receiver identifier 1130: an identifier (e.g., a messaging system identifier, email address, or device identifier) indicative of a user of the client system 102 to which the message 1100 is addressed.

[0323] The contents (e.g., values) of the various components of message 1100 may be pointers to locations in tables within which content data values are stored. For example, an image value in the message image payload 1104 may be a pointer to (or address of) a location within an image table 1106. Similarly, values within the message video payload 1108 may point to data stored within a video table 1110, values stored within the message augmentation data 1114 may point to data stored in an augmentation table 1116, values stored within the message story identifier 1122 may point to data stored in a story table 1124, and values stored within the message sender identifier 1128 and the message receiver identifier 1130 may point to user records stored within an entity table 1132.

TIME-BASED ACCESS LIMITATION ARCHITECTURE

[0324] FIG. 12 is a schematic diagram illustrating an access-limiting process 1200, in terms of which access to content (e.g., an ephemeral message 1202 and associated multimedia payload of data) or a content collection (e.g., an ephemeral message group 1204) may be time-limited (e.g., made ephemeral).

[0325] An ephemeral message 1202 is shown to be associated with a message duration parameter 1206, the value of which determines the amount of time that the ephemeral message 1202 will be displayed to a receiving user of the ephemeral message 1202 by the interaction client 104. In some examples, an ephemeral message 1202 is viewable by a receiving user for up to a maximum of 10 seconds, depending on the amount of time that the message sender specifies using the message duration parameter 1206.

[0326] The message duration parameter 1206 and the message receiver identifier 1208 are shown to be inputs to a message timer 1210, which is responsible for determining the amount of time that the ephemeral message 1202 is shown to a particular receiving user identified by the message receiver identifier 1208. In particular, the ephemeral message 1202 will be shown to the relevant receiving user for a time period determined by the value of the message duration parameter 1206. The message timer 1210 is shown to provide output to a more generalized messaging system 1212, which is responsible for

the overall timing of display of content (e.g., an ephemeral message 1202) to a receiving user.

[0327] The ephemeral message 1202 is shown in FIG. 12 to be included within an ephemeral message group 1204 (e.g., a collection of messages in a personal story, or an event story). The ephemeral message group 1204 has an associated group duration parameter 1214, a value of which determines a time duration for which the ephemeral message group 1204 is presented and accessible to users of the interaction system 100. The group duration parameter 1214, for example, may be the duration of a music concert, where the ephemeral message group 1204 is a collection of content pertaining to that concert. Alternatively, a user (either the owning user or a curator user) may specify the value for the group duration parameter 1214 when performing the setup and creation of the ephemeral message group 1204.

[0328] Additionally, each ephemeral message 1202 within the ephemeral message group 1204 has an associated group participation parameter 1216, a value of which determines the duration of time for which the ephemeral message 1202 will be accessible within the context of the ephemeral message group 1204. Accordingly, a particular ephemeral message group 1204 may “expire” and become inaccessible within the context of the ephemeral message group 1204 prior to the ephemeral message group 1204 itself expiring in terms of the group duration parameter 1214. The group duration parameter 1214, group participation parameter 1216, and message receiver identifier 1208 each provide input to a group timer 1218, which operationally determines, firstly, whether a particular ephemeral message 1202 of the ephemeral message group 1204 will be displayed to a particular receiving user and, if so, for how long. Note that the ephemeral message group 1204 is also aware of the identity of the particular receiving user as a result of the message receiver identifier 1208.

[0329] Accordingly, the group timer 1218 operationally controls the overall lifespan of an associated ephemeral message group 1204 as well as an individual ephemeral message 1202 included in the ephemeral message group 1204. In some examples, each and every ephemeral message 1202 within the ephemeral message group 1204 remains viewable and accessible for a time period specified by the group duration parameter 1214. In a further example, a certain ephemeral message 1202 may expire within the context of ephemeral message group 1204 based on a group participation parameter 1216. Note that a message duration parameter 1206 may still determine the duration of

time for which a particular ephemeral message 1202 is displayed to a receiving user, even within the context of the ephemeral message group 1204. Accordingly, the message duration parameter 1206 determines the duration of time that a particular ephemeral message 1202 is displayed to a receiving user regardless of whether the receiving user is viewing that ephemeral message 1202 inside or outside the context of an ephemeral message group 1204.

[0330] The messaging system 1212 may furthermore operationally remove a particular ephemeral message 1202 from the ephemeral message group 1204 based on a determination that it has exceeded an associated group participation parameter 1216. For example, when a message sender has established a group participation parameter 1216 of 24 hours from posting, the messaging system 1212 will remove the relevant ephemeral message 1202 from the ephemeral message group 1204 after the specified 24 hours. The messaging system 1212 also operates to remove an ephemeral message group 1204 when either the group participation parameter 1216 for each and every ephemeral message 1202 within the ephemeral message group 1204 has expired, or when the ephemeral message group 1204 itself has expired in terms of the group duration parameter 1214.

[0331] In certain use cases, a creator of a particular ephemeral message group 1204 may specify an indefinite group duration parameter 1214. In this case, the expiration of the group participation parameter 1216 for the last remaining ephemeral message 1202 within the ephemeral message group 1204 will determine when the ephemeral message group 1204 itself expires. In this case, a new ephemeral message 1202, added to the ephemeral message group 1204, with a new group participation parameter 1216, effectively extends the life of an ephemeral message group 1204 to equal the value of the group participation parameter 1216.

[0332] Responsive to the messaging system 1212 determining that an ephemeral message group 1204 has expired (e.g., is no longer accessible), the messaging system 1212 communicates with the interaction system 100 (and, for example, specifically the interaction client 104) to cause an indicium (e.g., an icon) associated with the relevant ephemeral message group 1204 to no longer be displayed within a user interface of the interaction client 104. Similarly, when the messaging system 1212 determines that the message duration parameter 1206 for a particular ephemeral message 1202 has expired, the messaging system 1212 causes the interaction client 104 to no longer display an

indiciu (e.g., an icon or textual identification) associated with the ephemeral message 1202.

MACHINE ARCHITECTURE

[0333] FIG. 13 is a diagrammatic representation of the machine 1300 within which instructions 1302 (e.g., software, a program, an application, an applet, an app, or other executable code) for causing the machine 1300 to perform any one or more of the methodologies discussed herein may be executed. For example, the instructions 1302 may cause the machine 1300 to execute any one or more of the methods described herein. The instructions 1302 transform the general, non-programmed machine 1300 into a particular machine 1300 programmed to carry out the described and illustrated functions in the manner described. The machine 1300 may operate as a standalone device or may be coupled (e.g., networked) to other machines. In a networked deployment, the machine 1300 may operate in the capacity of a server machine or a client machine in a server-client network environment, or as a peer machine in a peer-to-peer (or distributed) network environment. The machine 1300 may comprise, but not be limited to, a server computer, a client computer, a personal computer (PC), a tablet computer, a laptop computer, a netbook, a set-top box (STB), a personal digital assistant (PDA), an entertainment media system, a cellular telephone, a smartphone, a mobile device, a wearable device (e.g., a smartwatch), a smart home device (e.g., a smart appliance), other smart devices, a web appliance, a network router, a network switch, a network bridge, or any machine capable of executing the instructions 1302, sequentially or otherwise, that specify actions to be taken by the machine 1300. Further, while a single machine 1300 is illustrated, the term “machine” shall also be taken to include a collection of machines that individually or jointly execute the instructions 1302 to perform any one or more of the methodologies discussed herein. The machine 1300, for example, may comprise the client system 102 or any one of multiple server devices forming part of the interaction server system 110. In some examples, the machine 1300 may also comprise both client and server systems, with certain operations of a particular method or algorithm being performed on the server-side and with certain operations of the particular method or algorithm being performed on the client-side.

[0334] The machine 1300 may include processors 1304, memory 1306, and input/output I/O components 1308, which may be configured to communicate with each other via a bus 1310. In an example, the processors 1304 (e.g., a Central Processing Unit

(CPU), a Reduced Instruction Set Computing (RISC) Processor, a Complex Instruction Set Computing (CISC) Processor, a Graphics Processing Unit (GPU), a Digital Signal Processor (DSP), an Application Specific Integrated Circuit (ASIC), a Radio-Frequency Integrated Circuit (RFIC), another processor, or any suitable combination thereof) may include, for example, a processor 1312 and a processor 1314 that execute the instructions 1302. The term "processor" is intended to include multi-core processors that may comprise two or more independent processors (sometimes referred to as "cores") that may execute instructions contemporaneously. Although FIG. 13 shows multiple processors 1304, the machine 1300 may include a single processor with a single-core, a single processor with multiple cores (e.g., a multi-core processor), multiple processors with a single core, multiple processors with multiple cores, or any combination thereof.

[0335] The memory 1306 includes a main memory 1316, a static memory 1318, and a storage unit 1320, both accessible to the processors 1304 via the bus 1310. The main memory 1306, the static memory 1318, and storage unit 1320 store the instructions 1302 embodying any one or more of the methodologies or functions described herein. The instructions 1302 may also reside, completely or partially, within the main memory 1316, within the static memory 1318, within machine-readable medium 1322 within the storage unit 1320, within at least one of the processors 1304 (e.g., within the processor's cache memory), or any suitable combination thereof, during execution thereof by the machine 1300.

[0336] The I/O components 1308 may include a wide variety of components to receive input, provide output, produce output, transmit information, exchange information, capture measurements, and so on. The specific I/O components 1308 that are included in a particular machine will depend on the type of machine. For example, portable machines such as mobile phones may include a touch input device or other such input mechanisms, while a headless server machine will likely not include such a touch input device. It will be appreciated that the I/O components 1308 may include many other components that are not shown in FIG. 13. In various examples, the I/O components 1308 may include user output components 1324 and user input components 1326. The user output components 1324 may include visual components (e.g., a display such as a plasma display panel (PDP), a light-emitting diode (LED) display, a liquid crystal display (LCD), a projector, or a cathode ray tube (CRT)), acoustic components (e.g., speakers), haptic components (e.g., a vibratory motor, resistance mechanisms), other signal generators, and so forth. The user input components 1326 may include

alphanumeric input components (e.g., a keyboard, a touch screen configured to receive alphanumeric input, a photo-optical keyboard, or other alphanumeric input components), point-based input components (e.g., a mouse, a touchpad, a trackball, a joystick, a motion sensor, or another pointing instrument), tactile input components (e.g., a physical button, a touch screen that provides location and force of touches or touch gestures, or other tactile input components), audio input components (e.g., a microphone), and the like.

[0337] In further examples, the I/O components 1308 may include biometric components 1328, motion components 1330, environmental components 1332, or position components 1334, among a wide array of other components. For example, the biometric components 1328 include components to detect expressions (e.g., hand expressions, facial expressions, vocal expressions, body gestures, or eye-tracking), measure biosignals (e.g., blood pressure, heart rate, body temperature, perspiration, or brain waves), identify a person (e.g., voice identification, retinal identification, facial identification, fingerprint identification, or electroencephalogram-based identification), and the like. The motion components 1330 include acceleration sensor components (e.g., accelerometer), gravitation sensor components, rotation sensor components (e.g., gyroscope). The biometric components 1328 may include a brain-machine interface (BMI) system that allows communication between the brain and an external device or machine. This is achieved by recording brain activity, translating it into a format that can be understood by a computer, and then using the resulting signals to control the device or machine.

[0338] Example types of BMI technologies, including:

- Electroencephalography (EEG) based BMIs, which record electrical activity in the brain using electrodes placed on the scalp.
- Invasive BMIs, which involve surgically implanting electrodes directly into the brain.
- Optogenetics BMIs, which use light to control the activity of specific nerve cells in the brain.
- Functional magnetic resonance imaging (fMRI)-based BMIs, which use magnetic fields to measure blood flow in the brain, which can be used to infer brain activity.

[0339] Any biometric data collected by the biometric components 1328 are captured and cached with only user approval and deleted on user request. Further, such biometric data may be used for very limited purposes, such as identification verification. To ensure limited and authorized use of biometric information and other Personally Identifiable Information (PII), access to this data is restricted to authorized personnel only, if at all. Any use of biometric data may strictly be limited to identification verification purposes, and the biometric data is not shared or sold to any third party without the explicit consent of the user. In addition, appropriate technical and organizational measures are implemented to ensure the security and confidentiality of this sensitive information.

[0340] The environmental components 1332 include, for example, one or cameras (with still image/photograph and video capabilities), illumination sensor components (e.g., photometer), temperature sensor components (e.g., one or more thermometers that detect ambient temperature), humidity sensor components, pressure sensor components (e.g., barometer), acoustic sensor components (e.g., one or more microphones that detect background noise), proximity sensor components (e.g., infrared sensors that detect nearby objects), gas sensors (e.g., gas detection sensors to detection concentrations of hazardous gases for safety or to measure pollutants in the atmosphere), or other components that may provide indications, measurements, or signals corresponding to a surrounding physical environment.

[0341] With respect to cameras, the client system 102 may have a camera system comprising, for example, front cameras on a front surface of the client system 102 and rear cameras on a rear surface of the client system 102. The front cameras may, for example, be used to capture still images and video of a user of the client system 102 (e.g., “selfies”), which may then be augmented with augmentation data (e.g., filters) described above. The rear cameras may, for example, be used to capture still images and videos in a more traditional camera mode, with these images similarly being augmented with augmentation data. In addition to front and rear cameras, the client system 102 may also include a 360° camera for capturing 360° photographs and videos.

[0342] Further, the camera system of the client system 102 may include dual rear cameras (e.g., a primary camera as well as a depth-sensing camera), or even triple, quad or penta rear camera configurations on the front and rear sides of the client system 102. These multiple cameras systems may include a wide camera, an ultra-wide camera, a telephoto camera, a macro camera, and a depth sensor, for example.

[0343] The position components 1334 include location sensor components (e.g., a GPS receiver component), altitude sensor components (e.g., altimeters or barometers that detect air pressure from which altitude may be derived), orientation sensor components (e.g., magnetometers), and the like.

[0344] Communication may be implemented using a wide variety of technologies. The I/O components 1308 further include communication components 1336 operable to couple the machine 1300 to a network 1338 or devices 1340 via respective coupling or connections. For example, the communication components 1336 may include a network interface component or another suitable device to interface with the network 1338. In further examples, the communication components 1336 may include wired communication components, wireless communication components, cellular communication components, Near Field Communication (NFC) components, Bluetooth[®] components (e.g., Bluetooth[®] Low Energy), Wi-Fi[®] components, and other communication components to provide communication via other modalities. The devices 1340 may be another machine or any of a wide variety of peripheral devices (e.g., a peripheral device coupled via a USB).

[0345] Moreover, the communication components 1336 may detect identifiers or include components operable to detect identifiers. For example, the communication components 1336 may include Radio Frequency Identification (RFID) tag reader components, NFC smart tag detection components, optical reader components (e.g., an optical sensor to detect one-dimensional bar codes such as Universal Product Code (UPC) bar code, multi-dimensional bar codes such as Quick Response (QR) code, Aztec code, Data Matrix, Dataglyph, MaxiCode, PDF417, Ultra Code, UCC RSS-2D bar code, and other optical codes), or acoustic detection components (e.g., microphones to identify tagged audio signals). In addition, a variety of information may be derived via the communication components 1336, such as location via Internet Protocol (IP) geolocation, location via Wi-Fi[®] signal triangulation, location via detecting an NFC beacon signal that may indicate a particular location, and so forth.

[0346] The various memories (e.g., main memory 1316, static memory 1318, and memory of the processors 1304) and storage unit 1320 may store one or more sets of instructions and data structures (e.g., software) embodying or used by any one or more of the methodologies or functions described herein. These instructions (e.g., the

instructions 1302), when executed by processors 1304, cause various operations to implement the disclosed examples.

[0347] The instructions 1302 may be transmitted or received over the network 1338, using a transmission medium, via a network interface device (e.g., a network interface component included in the communication components 1336) and using any one of several well-known transfer protocols (e.g., hypertext transfer protocol (HTTP)). Similarly, the instructions 1302 may be transmitted or received using a transmission medium via a coupling (e.g., a peer-to-peer coupling) to the devices 1340.

SOFTWARE ARCHITECTURE

[0348] FIG. 14 is a block diagram 1400 illustrating a software architecture 1402, which can be installed on any one or more of the devices described herein. The software architecture 1402 is supported by hardware such as a machine 1404 that includes processors 1406, memory 1408, and I/O components 1410. In this example, the software architecture 1402 can be conceptualized as a stack of layers, where each layer provides a particular functionality. The software architecture 1402 includes layers such as an operating system 1412, libraries 1414, frameworks 1416, and applications 1418. Operationally, the applications 1418 invoke API calls 1420 through the software stack and receive messages 1422 in response to the API calls 1420.

[0349] The operating system 1412 manages hardware resources and provides common services. The operating system 1412 includes, for example, a kernel 1424, services 1426, and drivers 1428. The kernel 1424 acts as an abstraction layer between the hardware and the other software layers. For example, the kernel 1424 provides memory management, processor management (e.g., scheduling), component management, networking, and security settings, among other functionalities. The services 1426 can provide other common services for the other software layers. The drivers 1428 are responsible for controlling or interfacing with the underlying hardware. For instance, the drivers 1428 can include display drivers, camera drivers, BLUETOOTH® or BLUETOOTH® Low Energy drivers, flash memory drivers, serial communication drivers (e.g., USB drivers), WI-FI® drivers, audio drivers, power management drivers, and so forth.

[0350] The libraries 1414 provide a common low-level infrastructure used by the applications 1418. The libraries 1414 can include system libraries 1430 (e.g., C standard library) that provide functions such as memory allocation functions, string manipulation functions, mathematic functions, and the like. In addition, the libraries 1414 can include

API libraries 1432 such as media libraries (e.g., libraries to support presentation and manipulation of various media formats such as Moving Picture Experts Group-4 (MPEG4), Advanced Video Coding (H.264 or AVC), Moving Picture Experts Group Layer-3 (MP3), Advanced Audio Coding (AAC), Adaptive Multi-Rate (AMR) audio codec, Joint Photographic Experts Group (JPEG or JPG), or Portable Network Graphics (PNG)), graphics libraries (e.g., an OpenGL framework used to render in two dimensions (2D) and three dimensions (3D) in a graphic content on a display), database libraries (e.g., SQLite to provide various relational database functions), web libraries (e.g., WebKit to provide web browsing functionality), and the like. The libraries 1414 can also include a wide variety of other libraries 1434 to provide many other APIs to the applications 1418.

[0351] The frameworks 1416 provide a common high-level infrastructure that is used by the applications 1418. For example, the frameworks 1416 provide various graphical user interface (GUI) functions, high-level resource management, and high-level location services. The frameworks 1416 can provide a broad spectrum of other APIs that can be used by the applications 1418, some of which may be specific to a particular operating system or platform.

[0352] In an example, the applications 1418 may include a home application 1436, a contacts application 1438, a browser application 1440, a book reader application 1442, a location application 1444, a media application 1446, a messaging application 1448, a game application 1450, and a broad assortment of other applications such as a third-party application 1452. The applications 1418 are programs that execute functions defined in the programs. Various programming languages can be employed to create one or more of the applications 1418, structured in a variety of manners, such as object-oriented programming languages (e.g., Objective-C, Java, or C++) or procedural programming languages (e.g., C or assembly language). In a specific example, the third-party application 1452 (e.g., an application developed using the ANDROID™ or IOS™ software development kit (SDK) by an entity other than the vendor of the particular platform) may be mobile software running on a mobile operating system such as IOS™, ANDROID™, WINDOWS® Phone, or another mobile operating system. In this example, the third-party application 1452 can invoke the API calls 1420 provided by the operating system 1412 to facilitate functionalities described herein.

[0353] Additional examples include:

[0354] Example 1 is a method comprising: receiving, by one or more processors, from a client system, a prompt of a user during an interactive session; filtering, by the one or more processors, the prompt of the user based on a set of platform policies; generating, by the one or more processors, a response based on the filtering of the prompt of the user; and communicating, by the one or more processors, to the client system, the response.

[0355] In Example 2, the subject matter of Example 1 includes, wherein filtering the prompt of the user comprises: filtering the user prompt for keywords based on a set of platform policies.

[0356] In Example 3, the subject matter of any of Example 1-2 includes, wherein the keywords are keywords of sensitive content.

[0357] In Example 4, the subject matter of any of Examples 1–3 includes, wherein generating the response based on the filtering comprises: generating a filtered prompt based on the user prompt and a set of platform policies; communicating the filtered prompt to a chatbot component; and receiving the response from the chatbot component, the response generated by the chatbot component based on the filtered prompt.

[0358] In Example 5, the subject matter of any of Examples 1–4 includes, wherein generating the response based on the filtering comprises: determining a set of platform resources based on the prompt of the user; and generating the response based on the set of platform resources.

[0359] In Example 6, the subject matter of any of Examples 1–5 includes, determining, by the one or more processors, that the response comprises sensitive content; and in response to determining that the response comprises sensitive content, generating, by the one or more processors, an appropriate response that does not comprise the sensitive content.

[0360] In Example 7, the subject matter of any of Examples 1–6 includes, detecting, by the one or more processors, that the sensitive content comprises harmful conduct; and in response to detecting the harmful conduct, redirecting, by the one or more processors, the user to seek appropriate help.

[0361] Example 8 is at least one machine-readable medium including instructions that, when executed by processing circuitry, cause the processing circuitry to perform operations to implement of any of Examples 1–7.

[0362] Example 9 is an apparatus comprising means to implement of any of Examples 1–7.

[0363] Example 10 is a system to implement of any of Examples 1–7.

CONCLUSION

[0364] Changes and modifications may be made to the disclosed examples without departing from the scope of the present disclosure. These and other changes or modifications are intended to be included within the scope of the present disclosure.

GLOSSARY

[0365] "Carrier signal" refers to any intangible medium that is capable of storing, encoding, or carrying instructions for execution by the machine and includes digital or analog communications signals or other intangible media to facilitate communication of such instructions. Instructions may be transmitted or received over a network using a transmission medium via a network interface device.

[0366] "Client device" refers to any machine that interfaces to a communications network to obtain resources from one or more server systems or other client devices. A client device may be, but is not limited to, a mobile phone, desktop computer, laptop, portable digital assistants (PDAs), smartphones, tablets, ultrabooks, netbooks, laptops, multi-processor systems, microprocessor-based or programmable consumer electronics, game consoles, set-top boxes, or any other communication device that a user may use to access a network.

[0367] "Communication network" refers to one or more portions of a network that may be an ad hoc network, an intranet, an extranet, a virtual private network (VPN), a local area network (LAN), a wireless LAN (WLAN), a wide area network (WAN), a wireless WAN (WWAN), a metropolitan area network (MAN), the Internet, a portion of the Internet, a portion of the Public Switched Telephone Network (PSTN), a plain old telephone service (POTS) network, a cellular telephone network, a wireless network, a Wi-Fi® network, another type of network, or a combination of two or more such networks. For example, a network or a portion of a network may include a wireless or cellular network, and the coupling may be a Code Division Multiple Access (CDMA) connection, a Global System for Mobile communications (GSM) connection, or other types of cellular or wireless coupling. In this example, the coupling may implement any of a variety of types of data transfer technology, such as Single Carrier Radio

Transmission Technology (1xRTT), Evolution-Data Optimized (EVDO) technology, General Packet Radio Service (GPRS) technology, Enhanced Data rates for GSM Evolution (EDGE) technology, third Generation Partnership Project (3GPP) including 3G, fourth-generation wireless (4G) networks, Universal Mobile Telecommunications System (UMTS), High Speed Packet Access (HSPA), Worldwide Interoperability for Microwave Access (WiMAX), Long Term Evolution (LTE) standard, others defined by various standard-setting organizations, other long-range protocols, or other data transfer technology.

[0368] "Component" refers to a device, physical entity, or logic having boundaries defined by function or subroutine calls, branch points, APIs, or other technologies that provide for the partitioning or modularization of particular processing or control functions. Components may be combined via their interfaces with other components to carry out a machine process. A component may be a packaged functional hardware unit designed for use with other components and a part of a program that usually performs a particular function of related functions. Components may constitute either software components (e.g., code embodied on a machine-readable medium) or hardware components. A "hardware component" is a tangible unit capable of performing certain operations and may be configured or arranged in a certain physical manner. In various examples, one or more computer systems (e.g., a standalone computer system, a client computer system, or a server computer system) or one or more hardware components of a computer system (e.g., a processor or a group of processors) may be configured by software (e.g., an application or application portion) as a hardware component that operates to perform certain operations as described herein. A hardware component may also be implemented mechanically, electronically, or any suitable combination thereof. For example, a hardware component may include dedicated circuitry or logic that is permanently configured to perform certain operations. A hardware component may be a special-purpose processor, such as a field-programmable gate array (FPGA) or an application-specific integrated circuit (ASIC). A hardware component may also include programmable logic or circuitry that is temporarily configured by software to perform certain operations. For example, a hardware component may include software executed by a general-purpose processor or other programmable processors. Once configured by such software, hardware components become specific machines (or specific components of a machine) uniquely tailored to perform the configured functions and are no longer general-purpose processors. It will be appreciated that the decision to implement a

hardware component mechanically, in dedicated and permanently configured circuitry, or in temporarily configured circuitry (e.g., configured by software), may be driven by cost and time considerations. Accordingly, the phrase "hardware component" (or "hardware-implemented component") should be understood to encompass a tangible entity, be that an entity that is physically constructed, permanently configured (e.g., hardwired), or temporarily configured (e.g., programmed) to operate in a certain manner or to perform certain operations described herein. Considering examples in which hardware components are temporarily configured (e.g., programmed), each of the hardware components need not be configured or instantiated at any one instance in time. For example, where a hardware component comprises a general-purpose processor configured by software to become a special-purpose processor, the general-purpose processor may be configured as respectively different special-purpose processors (e.g., comprising different hardware components) at different times. Software accordingly configures a particular processor or processors, for example, to constitute a particular hardware component at one instance of time and to constitute a different hardware component at a different instance of time. Hardware components can provide information to, and receive information from, other hardware components. Accordingly, the described hardware components may be regarded as being communicatively coupled. Where multiple hardware components exist contemporaneously, communications may be achieved through signal transmission (e.g., over appropriate circuits and buses) between or among two or more of the hardware components. In examples in which multiple hardware components are configured or instantiated at different times, communications between such hardware components may be achieved, for example, through the storage and retrieval of information in memory structures to which the multiple hardware components have access. For example, one hardware component may perform an operation and store the output of that operation in a memory device to which it is communicatively coupled. A further hardware component may then, at a later time, access the memory device to retrieve and process the stored output. Hardware components may also initiate communications with input or output devices, and can operate on a resource (e.g., a collection of information). The various operations of example methods described herein may be performed, at least partially, by one or more processors that are temporarily configured (e.g., by software) or permanently configured to perform the relevant operations. Whether temporarily or permanently configured, such processors may constitute processor-implemented components that operate to perform

one or more operations or functions described herein. As used herein, "processor-implemented component" refers to a hardware component implemented using one or more processors. Similarly, the methods described herein may be at least partially processor-implemented, with a particular processor or processors being an example of hardware. For example, at least some of the operations of a method may be performed by one or more processors or processor-implemented components. Moreover, the one or more processors may also operate to support performance of the relevant operations in a "cloud computing" environment or as a "software as a service" (SaaS). For example, at least some of the operations may be performed by a group of computers (as examples of machines including processors), with these operations being accessible via a network (e.g., the Internet) and via one or more appropriate interfaces (e.g., an API). The performance of certain of the operations may be distributed among the processors, not only residing within a single machine, but deployed across a number of machines. In some examples, the processors or processor-implemented components may be located in a single geographic location (e.g., within a home environment, an office environment, or a server farm). In other examples, the processors or processor-implemented components may be distributed across a number of geographic locations.

[0369] "Machine-readable storage medium" refers to both machine-storage media and transmission media. Thus, the terms include both storage devices/media and carrier waves/modulated data signals. The terms "computer-readable medium," "machine-readable medium" and "device-readable medium" mean the same thing and may be used interchangeably in this disclosure.

[0370] "Ephemeral message" refers to a message that is accessible for a time-limited duration. An ephemeral message may be a text, an image, a video and the like. The access time for the ephemeral message may be set by the message sender. Alternatively, the access time may be a default setting or a setting specified by the recipient. Regardless of the setting technique, the message is transitory.

[0371] "Machine storage medium" refers to a single or multiple storage devices and media (e.g., a centralized or distributed database, and associated caches and servers) that store executable instructions, routines and data. The term shall accordingly be taken to include, but not be limited to, solid-state memories, and optical and magnetic media, including memory internal or external to processors. Specific examples of machine-storage media, computer-storage media and device-storage media include non-volatile

memory, including by way of example semiconductor memory devices, e.g., erasable programmable read-only memory (EPROM), electrically erasable programmable read-only memory (EEPROM), FPGA, and flash memory devices; magnetic disks such as internal hard disks and removable disks; magneto-optical disks; and CD-ROM and DVD-ROM disks. The terms "machine-storage medium," "device-storage medium," "computer-storage medium" mean the same thing and may be used interchangeably in this disclosure. The terms "machine-storage media," "computer-storage media," and "device-storage media" specifically exclude carrier waves, modulated data signals, and other such media, at least some of which are covered under the term "signal medium."

[0372] "Non-transitory machine-readable storage medium" refers to a tangible medium that is capable of storing, encoding, or carrying the instructions for execution by a machine.

[0373] "Signal medium" refers to any intangible medium that is capable of storing, encoding, or carrying the instructions for execution by a machine and includes digital or analog communications signals or other intangible media to facilitate communication of software or data. The term "signal medium" shall be taken to include any form of a modulated data signal, carrier wave, and so forth. The term "modulated data signal" means a signal that has one or more of its characteristics set or changed in such a manner as to encode information in the signal. The terms "transmission medium" and "signal medium" mean the same thing and may be used interchangeably in this disclosure.

What is claimed is:

1. A method comprising:

receiving, by one or more processors, from a client system, a prompt of a user during an interactive session;

filtering, by the one or more processors, the prompt of the user based on a set of platform policies;

generating, by the one or more processors, a response based on the filtering of the prompt of the user; and

communicating, by the one or more processors, to the client system, the response.

2. The method of claim 1, wherein filtering the prompt of the user comprises:

filtering the user prompt for keywords based on a set of platform policies.

3. The method of claim 2, wherein the keywords are keywords of sensitive content.

4. The method of claim 1, wherein generating the response based on the filtering comprises:

generating a filtered prompt based on the user prompt and a set of platform policies;

communicating the filtered prompt to a chatbot component; and

receiving the response from the chatbot component, the response generated by the chatbot component based on the filtered prompt.

5. The method of claim 1, wherein generating the response based on the filtering comprises:

determining a set of platform resources based on the prompt of the user; and

generating the response based on the set of platform resources.

6. The method of claim 1, further comprising:

determining, by the one or more processors, that the response comprises sensitive content; and

in response to determining that the response comprises sensitive content, generating, by the one or more processors, an appropriate response that does not comprise the sensitive content.

7. The method of claim 1, further comprising:

detecting, by the one or more processors, that the sensitive content comprises harmful conduct; and

in response to detecting the harmful conduct, redirecting, by the one or more processors, the user to seek appropriate help.

8. A machine comprising:

one or more processors; and

a memory storing instructions that, when executed by the one or more processors, cause the machine to perform operations comprising:

receiving from a client system, a prompt of a user during an interactive session;

filtering the prompt of the user based on a set of platform policies;

generating a response based on the filtering of the prompt of the user; and

communicating, to the client system, the response.

9. The computing apparatus of claim 8, wherein filtering the prompt of the user comprises:

filtering the user prompt for keywords based on a set of platform policies.

10. The computing apparatus of claim 9, wherein the keywords are keywords of sensitive content.

11. The computing apparatus of claim 8, wherein generating the response based on the filtering comprises:

generating a filtered prompt based on the user prompt and a set of platform policies;

communicating the filtered prompt to a chatbot component; and

receiving the response from the chatbot component, the response generated by the chatbot component based on the filtered prompt.

12. The computing apparatus of claim 8, wherein generating the response based on the filtering comprises:

determining a set of platform resources based on the prompt of the user; and
generating the response based on the set of platform resources.

13. The computing apparatus of claim 8, wherein the operations further comprise:

determining that the response comprises sensitive content; and
in response to determining that the response comprises sensitive content,
generating an appropriate response that does not comprise the sensitive content.

14. The computing apparatus of claim 8, wherein the operations further comprise:

detecting that the sensitive content comprises harmful conduct; and
in response to detecting the harmful conduct, redirecting the user to seek
appropriate help.

15. A machine-readable medium storing executable instructions that, when executed by
one or more processors, cause the machine to perform operations comprising:

receiving from a client system, a prompt of a user during an interactive session;
filtering the prompt of the user based on a set of platform policies;
generating a response based on the filtering of the prompt of the user; and
communicating, to the client system, the response.

16. The machine-readable medium of claim 15, wherein filtering the prompt of the user
comprises:

filtering the user prompt for keywords based on a set of platform policies.

17. The machine-readable medium of claim 16, wherein the keywords are keywords of
sensitive content.

18. The machine-readable medium of claim 15, wherein generating the response based
on the filtering comprises:

generating a filtered prompt based on the user prompt and a set of platform
policies;

communicating the filtered prompt to a chatbot component; and

receiving the response from the chatbot component, the response generated by the
chatbot component based on the filtered prompt.

19. The machine-readable medium of claim 15, wherein generating the response based
on the filtering comprises:

determining a set of platform resources based on the prompt of the user; and
generating the response based on the set of platform resources.

20. The machine-readable medium of claim 15, the operations further comprise:

determining, by the one or more processors, that the response comprises sensitive content; and

in response to determining that the response comprises sensitive content,
generating, by the one or more processors, an appropriate response that does not
comprise the sensitive content.

1/21




FIG. 1




FIG. 2

User Chat with Chatbot	Intent Extraction from Chatbot Chat
User: Show me a list of good hotels in Cancun ChatBot: Here is a list of good hotels in Cancun: • Hotel A • Hotel B • Hotel C • Hotel D User: Tell me which is the best month to get a good deal for Hotel B in spring? ChatBot: April is the least expensive month for Hotel B User: How is the weather in April? ChatBot: The weather in Cancun in April is typically warm and sunny with daytime temperatures around 30°C (86°F) and high humidity. The chance of rain is low, but occasional thunderstorms can occur. The sea temperature is also warm, making it a good time for swimming and other water activities. Overall, April is a great time to visit Cancun for a beach vacation, but be prepared for heat and humidity.	Direct Keywords = {Cancun (high intent), Hotel A, Hotel B (high intent), Hotel C, Hotel D, April (high intent), Spring (high intent), Least Expensive (high intent)} Expanded Concepts = {Warm Weather, Beach, Caribbean Sea, All-Inclusive Resorts, Scuba Diving, Snorkeling, Water Sports, Mexican Culture, Archaeological Sites, Yucatan, Peninsula, Mexican Caribbean, Luxury Resorts, Ecotourism, Adventure Activities, Beach Clubs, Spas, Sun and Fun, White Sand Beaches, Clear Blue Waters, Tropical Climate, Culture and History, Relaxation, Nature and Adventure}

FIG. 3A




FIG. 3B

5/21




FIG. 3C

6/21

300




FIG. 3D




FIG. 4A

8/21




FIG. 4B




FIG. 5A




FIG. 5B

10/21




FIG. 6




FIG. 7A

12/21




FIG. 7B




FIG. 7C

14/21




FIG. 7D

15/21




FIG. 8

16/21




FIG. 9

17/21




FIG. 10

18/21





FIG. 11

**FIG. 12**

20/21





FIG. 13

**FIG. 14**

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2024/014450

A. CLASSIFICATION OF SUBJECT MATTER INV. H04L51/02 H04L51/04 H04L51/214 ADD.		
According to International Patent Classification (IPC) or to both national classification and IPC		
B. FIELDS SEARCHED Minimum documentation searched (classification system followed by classification symbols) H04L Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched Electronic data base consulted during the international search (name of data base and, where practicable, search terms used) EPO-Internal		
C. DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2017/250930 A1 (BEN-ITZHAK YUVAL [CZ]) 31 August 2017 (2017-08-31)	1, 8, 15
Y	abstract paragraph [0002] - paragraph [0107] figures 1-8 -----	2-7, 9-14, 16-20
Y	US 2019/171845 A1 (DOTAN-COHEN DIKLA [IL] ET AL) 6 June 2019 (2019-06-06) abstract paragraph [0014] - paragraph [0098] figures 1-5 -----	2-7, 9-14, 16-20
A	US 2022/199079 A1 (HANSON MICHAEL ROBERT [US] ET AL) 23 June 2022 (2022-06-23) abstract paragraph [0056] - paragraph [0187] figures 2-9 -----	1-20
☐ Further documents are listed in the continuation of Box C. ☒ See patent family annex.		
* Special categories of cited documents : "A" document defining the general state of the art which is not considered to be of particular relevance "E" earlier application or patent but published on or after the international filing date "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified) "O" document referring to an oral disclosure, use, exhibition or other means "P" document published prior to the international filing date but later than the priority date claimed "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention "X" document of particular relevance;; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone "Y" document of particular relevance;; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art "&" document member of the same patent family		
Date of the actual completion of the international search 22 May 2024		Date of mailing of the international search report 03/06/2024
Name and mailing address of the ISA/ European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016		Authorized officer Poggio, Francesca

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No
PCT/US2024/014450

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2017250930 A1	31-08-2017	NONE	
US 2019171845 A1	06-06-2019	US 2019171845 A1	06-06-2019
		US 2021026984 A1	28-01-2021
		WO 2019112841 A1	13-06-2019
US 2022199079 A1	23-06-2022	NONE	