

(51) International Patent Classification:
H04L 51/52 (2022.01)(21) International Application Number:
PCT/US2022/032039(22) International Filing Date:
03 June 2022 (03.06.2022)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
63/196,625 03 June 2021 (03.06.2021) US(71) Applicant: **TWITTER, INC.** [US/US]; 1355 Market St.,
Suite 900, San Francisco, CA 94103 (US).

(72) Inventors; and

(71) Applicants (for US only): **MEHTA, Abhijit, C.** [US/US];
c/o Twitter, Inc., 1355 Market St., Suite 900, San Francisco, CA 94103 (US). **DZWONEK, Mateusz** [US/US];
c/o Twitter, Inc., 1355 Market St., Suite 900, San Francisco, CA 94103 (US). **LIU, Gloria** [US/US]; c/o Twitter, Inc., 1355 Market St., Suite 900, San Francisco, CA 94103 (US). **CARRER, David** [US/US]; c/o Twitter, Inc., 1355 Market St., Suite 900, San Francisco, CA 94103 (US). **CARR, David** [US/US]; c/o Twitter, Inc., 1355 Market St., Suite 900, San Francisco, CA 94103 (US). **TARDIO, Sabrina** [US/US]; c/o Twitter, Inc., 1355 Market St., Suite 900, San Francisco, CA 94103 (US). **VARGHESE, Arpith** [US/US]; c/o Twitter, Inc., 1355 Market St., Suite 900, San Francisco, CA 94103 (US). **BEYKPOUR, Sara** [US/US]; c/o Twitter, Inc., 1355 Market St., Suite 900, San Francisco, CA 94103 (US). **GUPTA, Smita Mittal** [US/US]; c/o Twitter, Inc., 1355 Market St., Suite 900, San Francisco, CA 94103 (US).

(54) Title: LABELING MESSAGES ON A SOCIAL MESSAGING PLATFORM USING MESSAGE RESPONSE INFORMATION




FIG. 1A

(57) **Abstract:** A social messaging platform includes a labeling module to label a base message according to an aggregate message response parameter, which represents the sentiment of users towards the content of the base message. The labels provide information that can be used to distinguish more nuanced sentiments and the degree of the sentiment users may have towards the base message. The aggregate message response parameter and corresponding labels are determined, in part, by identifying and evaluating icons (e.g., emojis, emoticons) present in one or more response messages posted in response to a base message. The labels, in turn, can be used in a variety of applications including recommending new content to users based on their mood, identifying messages potentially containing toxic content for review, or providing a way for businesses to evaluate public sentiment towards an advertisement and facilitate targeted advertisements to users.

94103 (US). **MARKS, Timothy** [US/US]; c/o Twitter, Inc.,
1355 Market St., Suite 900, San Francisco, CA 94103 (US).

(74) **Agent: SUD, Dhruv** et al.; SMITH BALUCH LLP, 376
Boylston St., Suite 401, Boston, MA 02116 (US).

(81) **Designated States** (*unless otherwise indicated, for every kind of national protection available*): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IQ, IR, IS, IT, JM, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, 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, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

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))*

LABELING MESSAGES ON A SOCIAL MESSAGING PLATFORM USING MESSAGE RESPONSE INFORMATION

CROSS-REFERENCE TO RELATED APPLICATION(S)

[0001] The present application claims priority to U.S. Provisional Application No. 63/196,625, filed June 3, 2021, entitled “Audience Response Signals,” which is incorporated by reference herein in its entirety.

BACKGROUND

[0002] Social messaging platforms and network-connected personal computing devices allow users to create and share content across multiple devices in real-time. Sophisticated mobile computing devices, such as smartphones and tablets, further make it easy and convenient for people, companies, and other entities to use social networking messaging platforms and applications. Examples of social messaging platforms include Facebook®, Instagram®, Pinterest®, and Twitter®. (“Facebook” and “Instagram” are trademarks of Facebook, Inc., “Pinterest” is a trademark of Pinterest, Inc., and “Twitter” is a trademark of Twitter, Inc.).

[0003] A social messaging platform generally provides functionality for users to draft and post (or send) messages both synchronously and asynchronously to other users. The message can generally include textual content, graphical content, video content, and/or audio content. Other common features include the ability to post messages that are visible to one or more identified other users of the platform, to other users by virtue of a connection to the authoring user on the platform, or even publicly to any user of the platform without specific designation by the authoring user.

[0004] The social messaging platform also typically provides a user interface to display multiple messages with content for users to view and consume. In one example, the user interface can display messages in a stream (also referred to herein as a “timeline” or a “feed”) where the messages are arranged, for example, in chronological order or reverse chronological order based on the respective dates the messages were posted or according to a computationally predicted relevance. The stream can further display a first message and a second message posted by a user in response to the first message.

[0005] Herein, the first message is referred to as a “base message” and the second message can be referred to as a “response message.” The response message can be a “reply message,” which is a

message posted in direct response to the base message and sent to at least the authoring user of the base message. The response message can also be a “share message,” which is a message intended for other users that may or may not include the authoring user of the base message and includes the content of the base message without any additional content from the responding user. The response message can also be a “quote message,” which is a message intended for other users that may or may not include the authoring user of the base message and includes the content of the base message along with additional content from the responding user. A base message and one or more response messages posted in response to the base message constitute a “message branch.” It should be appreciated a base message can be a “root message” (i.e., a message that is not posted in response to another message) or a “response message” (e.g., a message posted in response to the root message or another response message). A “message thread” includes a single root message and response messages posted in response to the root message or another response message. The message thread can generally include multiple message branches.

SUMMARY

[0006] The Inventors have recognized and appreciated social messaging platforms that provide users different ways of personalizing the content displayed for consumption generally lead to a better user experience, which in turn results in greater user growth, user retention, and engagement between the users. For example, social messaging platforms often allow users to preferentially display messages posted by other users of the social messaging platform that have an established connection with the user (e.g., a first user is following multiple second users, the second users are following the first user). In another example, some social messaging platforms can display messages with content that corresponds to one or more topics of interest selected by a user.

[0007] However, the Inventors have also recognized social messaging platforms are still limited in terms of evaluating and understanding user engagement with different content. This, in turn, makes it more difficult for social messaging platforms to identify a trend or a pattern to recommend new content to users.

[0008] In particular, social messaging platforms often provide several ways for users to show positive engagement with a particular message, such as by posting a reply message, using the interface to “like” the message, indicate the message as their “favorite”, or sharing the message with other users (e.g., as a share message or a quote message). However, these mechanisms

generally do not provide a way to evaluate the degree of positive engagement with a particular message, which can be different between users. For example, social messaging platforms are typically unable to differentiate one user who only “likes” a message if it causes them to laugh from another user that “likes” that message because they find it thoughtful, and from yet another user that “likes” a large portion of the messages that they read without much discernment. Additionally, social messaging platforms are generally unable to determine the underlying reasons why a user positively engages a particular message. For example, one user may “like” a message because they think the message is funny while another user may “like” the message because they think the message is cute. This is, in part, due to the information that can be extracted from a message being limited to the topic of the message, the username of the authoring user, the location of the authoring user, or the age of the authoring user.

[0009] Conversely, social messaging platforms can also sometimes limit the ability of users to show negative engagement with a particular message. For example, some social messaging platforms may not provide users a user interface element to “dislike” a message. Instead, users may only use the user interface to “report” messages for review. Again, if two users “report” a message, the social messaging platform is typically unable to differentiate which user merely dislikes a message and wants to avoid seeing similar messages in the future and the user who is personally offended by a message and wants the message deleted. Users can also express their negative engagement by posting a response message (e.g., a reply message, or a quote message), but these types of messages are generally not standardized making their interpretation more challenging.

[0010] To the extent social messaging platforms provide users recommendations on new content despite the limited information on user engagement described above, the social messaging platforms are also slow or, in some instances, unable to adapt to changes in the user’s preferences. For example, a user’s mood can influence the type of messages and content they want to see. Moreover, the user’s mood can change over time (e.g., over the course of a day, a week, a month). However, social messaging platforms are generally unable to adjust the content delivered to a user based on their mood. As an illustrative example, a user may generally prefer their stream to display messages related to the news. However, if the user’s stream shows messages with predominantly negative news, the user may become fatigued over time. Rather than provide new messages to help improve the user’s mood (e.g., messages with happy or funny content), social messaging platforms

typically continue recommending the same content worsening user fatigue and causing the user to temporarily or, in some instances, permanently stop using the social messaging platform.

[0011] In view of the foregoing, the present disclosure is directed to various inventive implementations of a social messaging platform and methods for using the social messaging platform to label a base message of a message branch according to the sentiment of users who post a response message to the base message and thereafter use the labels to selectively deliver content to users of the messaging platform. The sentiment generally represents a reaction or attitude of a user towards a message (e.g., happy, sad, anger, fear, disgust), which can be evaluated by analyzing the content of one or more response messages posted in response to the base message. By evaluating the message response parameters of a first group of users towards a base message, the social messaging platforms disclosed herein can obtain an aggregate message response parameter (e.g., a digital representation of the sentiment towards the base message). The aggregate message response parameter, in turn, can be used to determine an appropriate label to assign the base message. The label indicates the expected reaction or sentiment when the base message is viewed by a second user (e.g., a user who did not previously view or respond to the base message, a new user of the social messaging platform).

[0012] The labels disclosed herein can provide more information on user engagement compared to previous social messaging platforms. First, the labels allow the social messaging platform to distinguish between specific types of engagement (e.g., the user finds the content funny, inspiring, sad, or offensive) compared to previous social messaging platforms, which are typically only able to distinguish between positive and negative engagement. Second, a confidence score associated with each label can be calculated representing the confidence in which the label assigned to a base message represents the sentiment of the users who posted a response message. The confidence score can be used to distinguish different degrees of engagement (e.g., one user finds a message moderately funny while another finds the same message very funny). It should be appreciated the label may be used in conjunction with or in replacement of previous mechanisms used to evaluate user engagement including, but not limited to, “liking” a message, indicating a message as a “favorite,” sharing the message with other users with (e.g., “quoting” the message as part of their own message) or without (e.g., “re-sharing” the message without additional annotation) further comment, and “reporting” a message.

[0013] The additional information provided by the labels can be used to provide better recommendations of new content for each user. This is due, in part, to the labels providing a more objective assessment of the base message compared to analyzing the base message alone as is typical in previous social messaging platforms. It not only provides insight into how users react to the base message, but also the nature of the base message itself without necessarily relying upon an interpretation of the content in the base message.

[0014] By utilizing the labels to acquire more information on user engagement, the sentiment to a message by one group of users can influence the recommendation of the message to other users of the social messaging platform. Although each user may have different preferences from other users and may further react and respond to a message differently from other users, the effects of these differences on the recommendations of new content are reduced by aggregating the message response parameters from multiple users. In other words, if the aggregate message response parameter includes responses from a sufficiently large number of users, the effect each user may have on labeling a base message can be reduced. This allows for a more objective assessment of user sentiment towards a base message, which in turn provides a way for the social messaging platform to determine whether a user is likely to have a similar sentiment to a base message as other users.

[0015] Additionally, the user's mood may generally influence the type of content they want to consume. For example, a user who is sad or depressed may want to view happy and/or uplifting content. This can be achieved herein by utilizing the labels associated with various base messages to recommend new content to user's based on their mood. Returning to the above example of the user who is sad or depressed, the user may provide an indication that they want to view happy and/or uplifting content. The social messaging platform, in turn, can recommend new content with labels of "funny", "uplifting," and/or "happy" to the user. The social messaging platform can further dynamically adjust the content delivered to each user as the user's mood changes over time. For example, the social messaging platform may provide a user interface that allows the user to input their mood or preferences for the type content they want to consume at any time. If the user's mood or preferences change, the social messaging platform can change the recommendations of new content accordingly.

[0016] The sentiment of the users who post a response message to a base message can be

evaluated, in part, by identifying icons present in the response messages (e.g., a reply message, a quote message). The icons can generally be any visual representation intended to convey an expression or emotion. For example, the icons can be emojis (e.g., a pictogram). In another example, the icons can be emoticons (e.g., a combination of punctuation marks, letters, numbers, and/or the like arranged to resemble a face or an object). Additionally, textual content in the response messages can be mapped to one or more icons using a natural language processing (NLP) model in order to obtain a set of icons representing the sentiment of the users who responded to the base message. Each icon can be associated with at least one message response parameter (e.g., an indexed value that associates a laughing icon with a “funny” label). Generally, the set of icons can include different icons representing different sentiments.

[0017] The set of icons are evaluated to determine an aggregate message response parameter and the base message is thereafter labeled according to the aggregate message response parameter. The label can be the representative icon, a textual description, or any other suitable visual indicator conveying the sentiment of the users to the base message. The evaluation process can be accomplished in several ways. In one example, the aggregate message response parameter and the label are determined by counting the number of times each unique icon is present in the response messages and identifying the icon with the highest count. As an illustrative example, a majority of the icons in the set of icons correspond to a message response parameter corresponding to a “funny” label, thus the base message is labeled as “funny.” In another example, the determination of the aggregate message response parameter and the label can be viewed as a classification problem. Specifically, the set of icons can be viewed as the input and the aggregate message response parameter and the label viewed as the output of a classification model. Various classification models can be used to evaluate the aggregate message response parameter and the label including, but not limited to, a logistic regression model, a Naive Bayes’s classifier model, a decision tree model, a random forest model, and a support vector machine model.

[0018] The label(s) for a base message are generally stored in memory on a platform server or in a database of the social message platform. The label(s) can thereafter be retrieved and used by different applications. For example, the social messaging platform can receive input from the user via the user interface on the user’s personal computing device indicating the user’s mood or a preference for specific types of content. This user input can be represented as a preferred label corresponding to the type of content the user wants to see. Thus, the social messaging platform

can retrieve messages with the same label for transmission to the user's personal computing device. In another example, the social messaging platform can evaluate and notify an agent (e.g., a human operator of the social messaging platform) to review a message that potentially contains toxic content based on the label(s) associated with the message (e.g., "angry," "disgust" labels). In another example, the social messaging platform can provide a way for advertisers to evaluate the response to particular advertisements and target other users with the same advertisement if the response receives a positive sentiment.

[0019] It should be appreciated that all combinations of the foregoing concepts and additional concepts discussed in greater detail below (provided such concepts are not mutually inconsistent) are contemplated as being part of the inventive subject matter disclosed herein. In particular, all combinations of claimed subject matter appearing at the end of this disclosure are contemplated as being part of the inventive subject matter disclosed herein. It should also be appreciated that terminology explicitly employed herein that also may appear in any disclosure incorporated by reference should be accorded a meaning most consistent with the particular concepts disclosed herein.

BRIEF DESCRIPTION OF THE DRAWINGS

[0020] The skilled artisan will understand that the drawings primarily are for illustrative purposes and are not intended to limit the scope of the inventive subject matter described herein. The drawings are not necessarily to scale; in some instances, various aspects of the inventive subject matter disclosed herein may be shown exaggerated or enlarged in the drawings to facilitate an understanding of different features. In the drawings, like reference characters generally refer to like features (e.g., functionally similar and/or structurally similar elements).

[0021] FIG. 1A shows an example message branch displayed via a user interface of a social messaging platform that includes a base message and multiple response messages posted in response to the base message.

[0022] FIG. 1B is a table that includes icons extracted from the response messages of FIG. 1A, an effective icon associated with each response message, and a label associated with each response message.

[0023] FIG. 1C shows the base message of FIG. 1A with a label selected based on the labels of

the response messages.

[0024] FIG. 2 shows the base message of FIG. 1A with multiple labels selected based on labels associated with response messages.

[0025] FIG. 3A shows another example of a base message posted on a social messaging platform.

[0026] FIG. 3B shows example response messages posted in response to the base message of FIG. 3A where one of the response messages is a share message.

[0027] FIG. 4 shows a flow chart of an example method for labeling a base message based on icons extracted from a plurality of response messages posted in response to the base message.

[0028] FIG. 5 shows a flow chart of an example method for retrieving and displaying messages on a computing device of a user based on a preferred label selected by the user.

[0029] FIG. 6 shows a flow chart of an example method for notifying an agent of a base message that potentially contains toxic content based on one or more labels associated with the base message.

[0030] FIG. 7 shows a flow chart of an example method for transmitting and displaying a base message with an advertisement targeted to one group of users based on the response to the base message by another group of users.

[0031] FIG. 8A shows an example social messaging platform communicatively coupled to multiple user devices via a network.

[0032] FIG. 8B shows a magnified view of the social messaging platform of FIG. 8A with a labeling module to facilitate labeling of base messages based on response messages posted in response to the base messages.

DETAILED DESCRIPTION

[0033] Following below are more detailed descriptions of various concepts related to, and implementations of, a social messaging platform configured to evaluate an aggregate message response parameter based on response messages (e.g., a reply message, a quote message) posted in response to a base message, label the base message based on the aggregate message response parameter, and use the label(s) associated with the base message to recommend content to users. Various methods of processing response messages, labeling messages, and using the labels are

also disclosed herein. It should be appreciated that various concepts introduced above and discussed in greater detail below may be implemented in multiple ways. Examples of specific implementations and applications are provided primarily for illustrative purposes so as to enable those skilled in the art to practice the implementations and alternatives apparent to those skilled in the art.

[0034] The figures and example implementations described below are not meant to limit the scope of the present implementations to a single embodiment. Other implementations are possible by way of interchange of some or all of the described or illustrated elements. Moreover, where certain elements of the disclosed example implementations may be partially or fully implemented using known components, in some instances only those portions of such known components that are necessary for an understanding of the present implementations are described, and detailed descriptions of other portions of such known components are omitted so as not to obscure the present implementations.

[0035] In the discussion below, various examples of inventive social messaging platforms are provided, wherein a given example or set of examples showcases one or more particular features of a labeling module and a user interface. It should be appreciated that one or more features discussed in connection with a given example of a social messaging platform may be employed in other examples of social messaging platforms according to the present disclosure, such that the various features disclosed herein may be readily combined in a given social messaging platform according to the present disclosure (provided that respective features are not mutually inconsistent).

1. Labeling a Base Message based on an Aggregate Message Response Parameter

[0036] The social messaging platforms disclosed herein generally include a labeling module (e.g., the labeling module 230 in FIGS. 8A and 8B) executed on one or more servers to label base messages based on an aggregate message response parameter. Specifically, the response messages posted in response to a base message are used to determine the sentiment of the users who posted the response messages. Once an aggregate message response parameter is determined, the base message can be labeled with one or more labels based on the aggregate message response parameter.

[0037] The aggregate message response parameter can be a digital representation of user

sentiment. By comparison, the label(s) can provide a visual, human-readable representation of user sentiment. For example, the social messaging platform can store in memory an array where each element of the array contains a label corresponding to a particular user sentiment. Here, the index of the array can be the aggregate message response parameter. In this example, the aggregate message response parameter may not have any meaning in isolation, but can nevertheless be used to distinguish different user sentiments by virtue of referencing different elements of the array.

[0038] In this manner, the label(s) associated with a base message represent the expected response of a user who consumes the base message. For example, if a group of users reacts to a base message with laughter, the social messaging platform can label the base message as “funny” to indicate most users are likely to laugh when reading the base message. Generally, various labels can be considered including, but not limited to, “funny,” “happy,” “uplifting,” “inspirational,” “angry,” “disgust,” “controversial,” “triumphant,” and “sad.”

[0039] Herein, the aggregate message response parameter is determined by evaluating the response messages for the presence of any icons (also referred to in the art as “emojis” or “emoticons”) and/or evaluating any other content in the response messages that can be mapped to one or more icons. For example, FIG. 1A shows a user interface on a user’s personal computing device displaying an example message branch 100. The message branch 100 includes a base message 110a posted by User 1 and a response section 130a with multiple response messages 132a, 132b, 132c, and 132d posted directly in response to the base message 110a by Users 2, 3, 4, and 5, respectively. The response messages 132a-132d are shown as reply messages (e.g., messages posted in direct response to the base message 110a, but visible to any users viewing the message branch 100. However, it should be appreciated any of the response messages 132a-132d can also be quote messages (e.g., the content of the response message includes the content 118 of the base message 110a in quotes).

[0040] Each user can be identified by a username 112, a profile picture 114, and a link 116 directed to the user’s account page. Additionally, each message can include user interface elements to facilitate user engagement with the message. For example, FIG. 1A shows the base message 110a and the response messages 132a, 132b, 132c, and 132d each include a reply element 120 for users to post reply messages directly to the base message 110a or the response messages 132a, 132b, 132c, and 132d, a share element 122 to share the base message 110a with other users of the social

messaging platform (e.g., by posting a share message or a quote message), a favorite element 124 for the user to “like” the base message 110a or indicate the base message 110a as a “favorite”, and an external share element 126 for users to share the base message 110a with individuals who are not users of the social messaging platform, e.g., via email, a text-messaging system, or another social messaging platform.

[0041] As shown in FIG. 1A, the base message 110a includes content 118 and the response messages 132a, 132b, 132c, and 132d respectively include content 134a, 134b, 134c, and 134d. In this example, the content 118, 134a, 134b, 134c, and 134d includes only textual content and/or one or more icons. More generally, the content of each message can include other types of content including, but not limited to an image (e.g., a Joint Photographic Experts Group (JPEG) image, a Graphical Interchange Format (GIF) image), a video (e.g., a Moving Picture Experts Group Layer-4 (MPEG-4) video, a QuickTime (MOV) video), and an audio file (e.g., a Moving Picture Experts Group Layer-3 (MPEG-3) audio file, a Waveform Audio File Format (WAV) audio file).

[0042] Herein, the base message can be a root message (e.g., a message that isn’t posted in response to any other message) or a response message. The response message considered when evaluating the aggregate message response parameter can be in direct response to the base message, such as the response messages 132a-132d in the message branch 100. However, it should be appreciated the response messages can also include messages that are not posted in direct response to the base message. For example, a first response message posted in response to any of the response messages 132a-132d can also be considered in the evaluation of the aggregate message response parameter. In another example, a second response message posted in response to the first response message can also be considered. More generally, any response messages that are linked to the base message via one or more message branches can be included in the evaluation of the aggregate message response parameter.

[0043] The response messages are generally visible to at least one other user that is different from the user who authored the base message or the user who authored the response message. For example, the response messages can be public messages visible to all users of the social messaging platform. In another example, the response messages can be messages that are visible only to a group of users defined, for example, by the user who authored the base message or the user who authored the response message. The response messages can generally be any message associated

with the base message being labeled including, but not limited to, a reply message, a share message, or a quote message.

1.1 Identifying Icons and Determining Message Response Parameters

[0044] To determine an aggregate message response parameter, the labeling module can begin by identifying any icons present in each response message. In general, the icons can generally be any visual representation intended to convey an expression or emotion. For example, the icons can be emojis, which are pictograms that often include a picture of a face with an expression (e.g., a smiley face, a crying face) or a picture of an object and are often used to convey a reaction or sentiment in line with textual content. In another example, the icons can be emoticons (e.g., a combination of punctuation marks, letters, numbers, and/or the like arranged to resemble a face or an object). The icons can also include static images, animated images, or video. Herein, icons provide a convenient way to evaluate the sentiment of the users towards a base message (e.g., the sentiment of the Users 2-5 towards the base message 110a). For example, a “laughing” emoji used in a response message is often used to show the user thinks the content of the base message is funny.

[0045] In some implementations, the icons used in messages can conform to a standard set of icons. For example, the Unicode Standard includes many emojis represented as pictograms with indexed codes and descriptions of the emojis. In the Unicode Standard, the emojis are treated in the same manner as any alphanumeric character. Thus, the process of identifying an icons can include, for example, comparing the textual content in a response message against an array containing a reference set of icons (e.g., the standard emojis in the Unicode Standard) to determine which individual characters in each response message are icons and to identify the icons amongst the reference set of icons. The reference set of icons and the identified icons in each response message can be stored in memory (e.g., onboard memory of a platform server, or a database). For example, the icons can be stored in a tuple where each element of the tuple includes the icons identified for a particular response message and the corresponding index of the element corresponds to a particular response message.

[0046] In addition to identifying icons in a response message, other types of content in the message can also be mapped onto one or more icons. For example, a response message containing the statement “that’s hilarious” can be mapped to a “laughing” icon. This can be accomplished using,

for example, a natural language processing (NLP) model trained to convert textual statements to one or more icons. The icons that are mapped onto the textual statements can be the same as the reference set of icons described above. The mapped icons can also be stored in memory in the same manner as described above.

[0047] In one example, the MIT DeepMoji model can be used to map one or more icons onto textual content in a response message. The DeepMoji model utilizes a Long Short-Term Memory (LSTM) neural network model trained to predict one or more emojis associated with an input sentence or statement provided as text. The LSTM model is pretrained to predict emojis associated with text using millions of social media messages that include both text and emojis. To predict new icons and/or interpret new words or statements over time, a transfer learning approach is adopted where the pretrained LSTM model is used to initialize fitting parameters in a new model (e.g., a new LSTM model). Specifically, a chain-thaw approach is applied where each layer of the neural network is iteratively trained with new data while keeping all remaining layers frozen (i.e., the fitting parameters for the layers are fixed). Once each layer is trained, all the layers of the model are trained with the new data together.

[0048] It should be appreciated other types of content, such as an image, a video, or an audio file can also be mapped onto one or more icons. For image and video content, this can be accomplished, for example, using various Optical Character Recognition (OCR) techniques. For example, a matrix matching technique can be used to compare an image or a still frame of a video in a response message against a library of images for the reference set of icons on a pixel-by-pixel basis. The comparison can entail, for example, determining differences in grayscale values between pixels in an image of a response image and pixels in a reference image having the same pixel coordinates. This approach may entail including different reference images commonly used to represent the same icon as well as including functions to automatically rescale images and/or frames of a video in a response message. In another example, a feature extraction technique can be used where a machine model is trained to identify icons by detecting and matching representative features associated with each icon in an image or still frame of a video to pretrained icons. The representative features can include, for example, the overall shape of the icon, the color of the icon, the features that define a facial expression.

[0049] For audio content, various techniques can be utilized to identify sounds (e.g., laughter,

crying, anger, joy) that convey specific emotions and map the identified sounds to one or more icons. For example, a model can be calibrated to detect and identify variations in the pitch of a person's voice and match the variations to specific emotions. In another example, a machine learning model can be trained using training data that includes a sample of different sounds and corresponding labels indicating an emotion or sentiment conveyed by the sounds.

[0050] Returning to the example of the message branch 100 of FIG. 1A, the content 134a of the response message 132a includes multiple different “laughing” icons 190a and a “crying” icon 190b. The content 134b of the response message 132b includes several identical “laughing” icons 190c. The content 134c of the response message 132c does not include any icons, but includes the statement “That’s too funny,” which can be mapped onto a “laughing” icon 190d. Lastly, the content 134d of the response message 132d includes an “angry” icon 190e. These icons are tabulated in FIG. 1B.

[0051] In some implementations, once one or more icons are identified in a response message, the labeling module can optionally evaluate each response message for the user's sentiment expressed in the response message and assign one or more labels accordingly. By evaluating each response message, the labeling module can identify more nuanced expressions conveyed in a response message especially for response messages that include multiple icons. This can be accomplished, in part, by each icon in the reference set of icons having one or more defined labels where each label has a corresponding message response parameter.

[0052] The reference set of icons can be stored in a tuple in memory where each row corresponds to a particular icon and at least one column includes the associated labels. When the labeling module assigns one or more labels to a response message, the label(s) associated with the response message can also be stored in memory. For example, a multi-dimensional tuple can be used where each row corresponds to a response message and respective columns correspond to the identified icons and the labels associated with a particular response message. This multi-dimensional tuple thus stores a set of icons identified in the response messages.

[0053] For response messages that have only one identifiable icon, the labeling module can assign a label to the response message based on the label associated with the icon. For example, a “smiley face” icon having a Unicode index of 1F600 can have a label of “happy.” When a user uses the “smiley face” icon in a message, the message can be interpreted as conveying a happy sentiment.

In another example, the “laughing” icon 190d mapped from the response message 132c can be assigned a label of “funny.”

[0054] Multiple labels can also be associated with a particular icon. For example, a “crying face” icon having a Unicode index of 1F62D can have labels of “sad” and “funny.” To determine which message response parameter and, hence, which label to associate with the icon, the labeling module can evaluate other content in the message for context. In one scenario, a message can include a “laughing” icon together with the “crying face” icon. Thus, the “crying face” icon is more likely to convey laughter and should thus be labeled as “funny.” In some implementations, the combination of the “laughing” icon together with the “crying face” icon can be stored as specific combination of icons with a particular meaning in the reference set of icons as described below. In another scenario, a NLP model can identify the statement “I’m really depressed” together with the “crying face” icon in a response message. Here, the “crying face” icon is more likely to convey grief and should thus be labeled as “sad.” In another example, the “angry” icon 190e in the response message 132d can have the labels of “angry” or “triumphant.” The accompanying statement in the content 134d of “Cat is just being rigorous” can be interpreted by a NLP as not conveying anger, thus the “angry” icon 190e is more likely to convey triumph and should thus be labeled as “triumphant.” Various pre-trained NLP models can be utilized to perform semantic analysis on textual content including, but not limited to, Bidirectional Encoder Representations from Transformers (BERT), Robustly Optimized BERT Pretraining Approach (RoBERTa), OpenAI’s GPT-3 model, A Lite BERT (ALBERT), StructBERT, XLNet, T5, Efficiently Learning an Encoder that Classifies Token Replacements Accurately (ELECTRA), and Decoding-enhanced BERT with disentangled attention (DeBERTa).

[0055] Additionally, multiple icons can share the same label. For example, multiple icons that convey laughter, such as the “laughing” icons 190a and can each be labeled as “funny.” In this manner, response messages that convey a similar sentiment using different icons are accommodated by the labeling module and can thus contribute similarly towards determining the aggregate message response parameter.

[0056] For response messages that have multiple identifiable icons, the labeling module can also be configured to evaluate the overall sentiment expressed by the combination of icons. For instance, certain combinations of icons can have a specific meaning that cannot be readily

determined by evaluating each icon in isolation. For example, a “face with open mouth” icon (Unicode index of 1F62E) is often used to convey surprise and a “dashing away” icon (Unicode index of 1F4A8) is often used to show a person or an object is running away. However, the combination of the “face with open mouth” icon and the “dashing away” icon can convey the sentiment that a user is shocked and either wishes to have nothing to do with the subject of the base message, or is extremely interested. To account for these combinations of icons, the reference set of icons can include both individual icons as described above and combinations of icons with respective labels. The label for a combination of icons can differ from the labels used for each icon in the combination.

[0057] In some implementations, the labeling module can be configured to search and label combinations in order of the number of icons in the combinations before lastly searching and labeling individual icons. For example, the labeling module can search and label combinations with three icons first, combinations with two icons second, and individual icons last. If a combination of icons is identified and assigned a label by the labeling module, the labeling module can skip these icons when searching for individual icons to label.

[0058] Additionally, response messages often include multiple icons that don’t have a distinct meaning separate from each individual icon, but are instead used to emphasize a user’s reaction to a base message. For example, the response message 132b includes multiple identical “laughter” icons 190c. The presence of multiple icons that share the same message response parameter and/or label can be treated in several ways.

[0059] In some implementations, the labeling module can treat multiple icons having the same message response parameter or label as effectively one icon, thus effectively ignoring the repeated usage of the icon. Returning to the example of the response message 132b, the labeling module would only consider one of the “laughter” icons 190c when assigning the label “funny” to the response message 132b. Thus, the response messages 132b and 132c contribute equally when the labeling module is determining the aggregate message response parameter.

[0060] In some implementations, the labeling module can consider multiple icons in the response message when determining the aggregate message response parameter. For example, the labeling module can treat messages with multiple icons having the same label as conveying a stronger sentiment than messages that only have a single icon with the same label. When the labeling

module determines the aggregate message response parameter, a response message with four “laughter” icons can have stronger effect on determining the aggregate message response parameter than a response message with a single “laughter” icon. In one example, if the labeling module determines the aggregate message response parameter by counting the number of icons present in the response messages, response messages with more “laughter” icons contribute more towards the total number of “laughter” icons present amongst the response messages.

[0061] This approach can also be used when the response message includes multiple icons conveying different sentiments. For example, the labeling module can count the number of icons having the same label in the response message and choose the label with the highest count to assign to the response message. Returning to the example of a response message that contains four “laughing” icons and one “angry” icon, the labeling module can determine there are more “laughing” icons than “angry” icons and thus assign the response message as “funny” based on the higher number of “laughing” icons present in the response messages.

[0062] In some implementations, the labeling module can assign multiple labels to a response message to indicate the response message includes multiple sentiments. That way, different sentiments present in a response message are considered when the labeling module is determining the aggregate message response parameter. The labeling module can further apply a weight to each label to indicate the relative strength of the sentiments present in the response message. Returning to the above example of a response message having four “laughing” icons and one “angry” icon, the labeling module can label the response message as “funny” and “angry.” Additionally, the “funny” label can have a normalized weight of 0.8 (4 icons out of 5 total icons) and the “angry” label can have a normalized weight of 0.2 (1 icon out of 5 total icons).

[0063] It should be appreciated not all response messages contain identifiable icons. FIG. 3A shows another example base message 110b and FIG. 3B shows an example response section 130b with response messages 132e and 132f. As shown, the response message 132e is a share message where the content 134e only quotes the base message 110b and does not include any additional content or icons from the authoring user. In comparison, the content 134f of the response message 132f includes multiple “laughing” icons. When evaluating the aggregate message response parameter, the labeling module can skip the response message 132e and consider only the response message 132f.

1.2 Determining an Aggregate Message Response Parameter and Label

[0064] After the labeling module identifies one or more icons in the response messages and, optionally, assigns one or more labels to each response message in a message branch, the labeling module can store the set of icons identified in memory. Thereafter, the labeling module can evaluate the aggregate message response parameter and assign one or more labels to the base message accordingly. This can be accomplished in several ways.

[0065] In one example, the labeling module can count the number of times each unique icon or label is present in the set of icons and assign a label to the base message based on the label with the highest count. Returning to the example of the message branch 100, FIG. 1B shows that the response messages 132a-132c are labeled as “funny” and the response message 132d is labeled as “triumphant.” Since there are three “funny” labels and one “triumphant” label, the labeling module determines the aggregate message response parameter to the base message 110a corresponds to the “funny” label and assigns the “funny” label to the base message 110a accordingly.

[0066] In some implementations, the user interface of the social messaging platform can update the appearance of the base message 110a to include a label 150 once the labeling process is completed, as shown in FIG. 1C. As shown, the label 150 includes textual content describing the user sentiment. More generally, the label 150 can include a representative icon (e.g., a “funny” icon), textual content (e.g., “funny”), or any other suitable visual indicator conveying the sentiment of the users who posted a response message.

[0067] In another example, the labeling module can utilize a classification model to predict the aggregate message response parameter and the corresponding label using the set of icons identified in the response messages as input. Various classification models can be used to evaluate the aggregate message response parameter and the label including, but not limited to, a logistic regression model, a Naive Baye’s classifier model, a decision tree model, a random forest model, and a support vector machine model.

[0068] As an illustrative example, a logistic regression model can be used to predict the aggregate message response parameter. The logistic regression model can utilize a sigmoid function to evaluate the probability of a particular aggregate message response parameter based on the set of icons identified in the response messages. Generally, the sigmoid function is used to evaluate classification problems with a binary output. However, the logistic regression model can be

adapted for use with multiple discrete outputs (e.g., different labels corresponding to different aggregate message response parameters) using a one-versus-all approach. For example, the probability of one aggregate message response parameter (e.g., “funny”) is evaluated with respect to all remaining aggregate message response parameters (e.g., “angry,” “happy,” “triumphant,” “sad”) where the remaining aggregate message response parameters are collectively treated as one output. By iteratively evaluating the probability of each possible aggregate message response parameter, the aggregate message response parameter with the highest probability can be determined and the label corresponding to the aggregate message response parameter can thereafter be assigned to the base message. The logistic regression model can generally be trained using multiple labeled message branches that each include response messages with different aggregate message response parameters and labels. In some implementations, the labeled message branches can be obtained from existing content posted on the social messaging platform.

[0069] In some implementations, the labeling module can assign multiple labels to a base message. This can be accomplished, in part, by the aggregate message response parameter being configured to be an array that contains multiple values corresponding to different labels. Alternatively, multiple aggregate message response parameters can be determined where each aggregate message response parameter corresponds to a particular label. For example, FIG. 2 shows the base message 110a can be displayed with multiple labels 150a and 150b to indicate different sentiments from the users who replied to the base message 150b. The process of choosing which labels to include can be achieved in several ways.

[0070] In one example approach, the labeling module can assign multiple labels to the base message based on the labels that have the highest count or highest probability. For example, a collection of response messages can have 10 “funny” icons, 5 “happy” icons, and 2 “angry” icons. If the labeling module is configured to label the base message based on the top two icons with the highest count, the aggregate response parameter can be an array that includes a pair of indexed values corresponding to the “funny” and “happy” labels. In another example approach, the labeling module can assign a label to the base message according to a predefined threshold. For example, the labeling module can assign a label to the base message if the label (or icon) exceeds a certain fraction of the total number of labels (or icons) in the response messages (e.g., more than 20%, more than 30%, more than 40%). In another example, the labeling module can assign a label to the base message when the probability of the label (or icon) exceeds a threshold (e.g., more than 20%,

more than 30%, more than 40%).

[0071] As described above, a response message can have multiple labels assigned with different weights. The labeling module can include the weights when determining the aggregate message response parameter. For example, if the message branch 100 includes a fifth response message with a “funny” label having a normalized weight of 0.4 and a “triumphant” label having a normalized weight of 0.6, the number of “funny” labels is evaluated to be 3.4 and the number of “triumphant” labels is evaluated to be 1.6. In another example, if a logistic regression model is used, the weights can influence the predicted probability of different aggregate message response parameters.

[0072] In some situations, one user can post multiple response messages in response to a base message. If each response message is considered equally, then some users may have more influence on the aggregate message response parameter than other users (e.g., users who post more response messages than other users). To reduce the influence of users who post multiple response messages, the labeling module can treat multiple response messages posted by one user as effectively one single response message. For example, a user who posts one response message of “That’s hilarious” and another response message of “I’m still laughing” can be grouped together and treated as one response message when determining the aggregate message response parameter. This can be accomplished, for example, by each response message having metadata that includes information on the authoring user (e.g., the username). The labeling module can identify and group together response messages based on the authoring user before evaluating the aggregate message response parameter.

[0073] In addition to assigning one or more labels to a base message, the labeling module can also determine a confidence score associated with each label(s). The confidence score can generally represent the confidence in assigning a label to the base message. For example, if multiple response messages only include “laughter” icons, the labeling module can assign a “funny” label to the base message and a confidence score of 1. The confidence score of 1 indicates the determined aggregate message response parameter corresponding to a “funny” label is universal amongst the response messages since only “laughter” icons are present. However, if one response message includes a “angry” icon, the labeling module may still assign a “funny” label to the base message, but with a confidence score less than 1 (and greater than 0). The confidence score can also be stored in

memory (e.g., onboard memory of the platform server, or a database) together with the labels associated with the base message.

[0074] The confidence score can be calculated in several ways. For example, if the number of each unique icon is counted when determining the aggregate message response parameter, the confidence score can be calculated as the fraction of the icon corresponding to the label ultimately assigned to the base message and the total number of icons identified and/or mapped in the response message. Returning to the example of a collection of response messages having 10 “funny” labels, 5 “happy” labels, and 2 “angry” labels, if the labeling module assigns the “funny” label to the base message, the confidence score can be equal to 0.59 (i.e., 10 “funny” labels divided by 17 labels in total). In another example, the confidence score can correspond to the predicted probability of an aggregate message response parameter determined using a logistic regression model.

[0075] The confidence score can be used, in part, to determine which label to assign the base message in accordance with the various approaches described above. In some implementations, the confidence score can also be used to rank and compare base messages with the same label. This can be used, for example, to determine which base messages to recommend to users. If a user wants to view more content with a “happy” label, the confidence score can be used to identify base messages with the “happy” label and the highest confidence score for the “happy” label. The confidence score can also be used to identify messages that potentially have toxic content (e.g., base messages assigned negative labels with high confidence scores). The confidence score can also provide a quantitative metric for companies to evaluate the sentiment users have towards different advertisements.

[0076] In some implementations, certain requirements may be imposed to label a base message. In other words, not all base messages posted to a social messaging platform are labeled by the labeling module. For example, the labeling module can check the base message and the response messages to determine whether they meet the requirements. In another example, another software module of the social messaging platform can monitor messages posted to the social messaging platform and identify base messages with response messages that meet the requirements for labeling. If the response messages associated with a base message meet the requirements, the software module can trigger the labeling module to begin the process of labeling the base message.

For example, the social messaging platform can include a software module to detect and monitor popular social media posts, which can be defined for the purposes of the present disclosure as base messages that receive a certain number of response messages over a set period of time (e.g., 100 response messages per hour, 1000 response messages per hour, 10,000 response messages per hour).

[0077] In one example, the labeling module only labels a base message if the number of response messages posted in direct response to the base message exceed a predefined threshold. The threshold number of response messages can be chosen to ensure the response messages provide a representative sample of the users of the social messaging platform when determining the aggregate message response parameter and the corresponding label. The threshold number of response messages can range between 100 to 1000 response messages, though any suitable threshold can be employed.

[0078] In another example, the labeling module only labels a base message if the number of response messages having identifiable icons posted in direct response to the base message exceeds a predefined threshold. For example, the response message 132e is not considered as having an identifiable icon due to the lack of content. This requirement ensures a sufficient number of response messages can be considered by the labeling module when determining the aggregate message response parameter. For this requirement, the labeling module can identify icons present in the response messages and/or map content in the messages to one or more icons to determine whether a response message has identifiable icons. Similar to the preceding example, the threshold number of response messages having identifiable icons can range between 100 to 1000 response messages, though any suitable threshold can be employed.

[0079] In another example, the labeling module only labels a base message if the number of unique users who post a response messages exceeds a predefined threshold. This requirement is intended to prevent one user who posts multiple response messages to a base message from having an appreciable effect on the determination of the aggregate message response parameter. Similar to the preceding examples, the threshold number of unique users who post a response message can range between 100 to 1000 users, though any suitable threshold can be employed.

[0080] It should be appreciated that the example requirements described above can be used in isolation or in combination with one another. For example, the labeling module can first determine

whether the number of response messages exceeds a first threshold and then determine the number of unique users who posted the response messages exceeds a second threshold.

[0081] It should also be appreciated users typically post response messages over an extended period of time. Thus, if the response messages for a base message initially fails to meet the requirements for labeling, it's possible they may later meet the requirements once more users post response messages. In some implementations, the labeling module (or another software module) can be configured to periodically check whether the response messages posted in response to a base message meet the requirements for labeling. For example, the response messages can be checked once a minute, once an hour, once a day, or once a week after the base message is posted. In another example, the response messages can be checked after the labeling module (or another software module) detects a threshold number of response messages are added to the previous set of response messages. In some implementations, the threshold number of additional response messages can be 5 response messages, 10 response messages, 100 response messages, or the like.

[0082] Additionally, even if the response messages meets the various requirements imposed by the labeling module (or another software module), the labeling module is not guaranteed to label the base message. The labeling module can determine there is no clear aggregate message response parameter for a base message. For example, the labeling module can be configured to assign a label to a base message if the label (or icon) exceeds a threshold fraction of the total number of labels (or icons) of 20%. If, however, the labeling module determines no unique labels meets this threshold, the labeling module does not assign any label to the base message.

1.3 Example Applications using Labels

[0083] The labels assigned to the base message can be stored in memory (e.g., onboard memory of a platform server, a database) together with the base messages. Thereafter, a software application supported by the social messaging platform can access the memory and retrieve the base messages together with the labels. Various software applications can utilize the labels assigned to the base messages to improve the user's experience with the social messaging platform, personalize the content provided to each user, and help businesses to evaluate public sentiment on their products and/or advertisements.

[0084] In the following, several example methods are disclosed representing various applications in which the labels assigned to the base message can be used. It should be appreciated that one or

more steps discussed in connection with one method can be readily incorporated and added to other methods contemplated herein.

1.3.1 An Example Method for Labeling a Base Message

[0085] FIG. 4 shows an example method 300 for assigning one or more labels to a base message and thereafter displaying the labels together with the base message to users who posted the base message and the response messages. At step 302, a base message is received from a first computing device associated with a first user. This can be accomplished by the first computing device receiving user input via a user interface of the first computing device (e.g., a touchscreen, a keyboard). For example, the first user can type a base message via a touchscreen of their mobile device and post the base message via the user interface. Thereafter, the first computing device can transmit the base message to a platform server of the social messaging platform via a network (e.g., see network 220 in FIG. 8A) and the platform server can thereafter store the base message for retrieval by other users in onboard memory or a database.

[0086] At step 304, the base message is transmitted to a plurality of second computing devices associated with a plurality of second users. This can be accomplished by the platform server retrieving the base message from memory and thereafter transmitting the base message to at least some of the users of the social messaging platform (e.g., users following the first user, users following a topic where the base message includes content related to the topic, all the users of the social messaging platform). It should be appreciated the plurality of second users is generally different from the first user. The base message is then displayed via respective user interfaces on each of the plurality of second computing devices in step 306. Thereafter, a plurality of response messages posted in response to the base message is received from the plurality of second computing devices in step 308. This can be accomplished in a similar manner as the receipt of the base message in step 302. For example, the second computing device can receive the response message via the user interface, transmit the response message to the platform server, and the platform server can thereafter store the response message in onboard memory or a database. It should also be appreciated that not all of the second users will post a response message. Thus, in step 308, the response messages are more generally received from at least a subset of the plurality of second computing devices.

[0087] Once a plurality of response messages is received, the response messages are then

evaluated to determine whether they meet the requirements to label the base message in step 310. This step can be executed by the labeling module via one or more platform servers (e.g., the labeling module 230 is executed by the processor(s) 212a in the platform server 210a of FIG. 8B) or another software module (e.g., a module configured to identify popular messages posted to the social messaging platform). The requirements correspond to the requirements discussed above in Section 1.2. For example, the labeling module can determine whether the number of response messages or the number of response messages containing identifiable icons exceed a predefined threshold.

[0088] If the plurality of response messages fails to meet the requirements, the labeling module or another software module can reevaluate the plurality of response messages after a set period of time to assess any new response messages posted in response to the base message (e.g., after a day, or a week). Otherwise, if the plurality of response messages meets the requirements, the method 300 can proceed to the next step of identifying icons in the response messages. At step 312, the labeling module can identify icons present in the plurality of response messages using the various processes described above in Section 1.1. Additionally, the labeling module can also map other types of content (e.g., textual content) to one or more icons using a natural language processing model in step 314. Steps 312 and 314 can be executed in parallel as shown in FIG. 4 or in a serial manner. The method can also optionally include steps to evaluate and assign label(s) to each of the response messages. The icons identified in the response messages can be consolidated into a set of icons that are stored in memory (e.g., as a tuple).

[0089] At step 316, the labeling module then determines the aggregate message response parameter and the corresponding label(s) to assign the base message based on the set of icons. Specifically, the identified icons in the response messages and/or any assigned label(s) to the response messages can be used as input. As described above, the labeling module can make this determination in several ways, for example, by counting the number of times each unique icon or label appears in the response messages or determining the probability of an aggregate message response parameter using a machine learning model (e.g., a logistic regression model). Once the labeling module determines the aggregate message response parameter, the labeling module can then assign a label to the base message. It should be appreciated multiple labels can be assigned to the base message.

[0090] Thereafter, the label(s) assigned to the base message can be stored in memory (e.g., the onboard memory of the platform server, or a database) in step 318. The label(s) can further include metadata identifying the base message associated with the label(s). Alternatively, the base message itself can include metadata that includes the label(s). At step 320, the label(s) can be transmitted to the first computing device and the plurality of second computing devices. Alternatively, the base message with updated metadata that includes the label(s) can be transmitted to the first computing device and the plurality of second computing devices. Thereafter, the base message can be displayed via respective user interfaces of the first computing device and the plurality of second computing devices with the label(s). That way, the first user and the plurality of second users can see the label(s) assigned to the base message if, for example, the users revisit the base message.

1.3.2 An Example Method for Recommending New Content

[0091] Once a base message is labeled, the social messaging platform can utilize the labels to recommend new content to users based on their mood. This can be accomplished by the user providing a user input to indicate their mood or a preference for content that elicits a particular sentiment. FIG. 5 shows an example method 400 for transmitting and displaying messages to a user based on a user input indicating a preferred label for the messages to be viewed. It should be appreciated the method 400 can follow some of the steps in the method 300 (e.g., step 318). At step 402, a user input is received from a computing device associated with a user that indicates a preference for messages to be displayed having a preferred label. This can be accomplished, for example, by the user interface of the computing device displaying a selection of labels that the user can select. Once the user input is received, the computing device can then transmit the user input to a platform server.

[0092] The labeling module or another software module (e.g., the applications 232a or 232b in FIG. 8B) can then be executed to retrieve one or more base messages from memory (e.g., onboard memory 214a or a database 216) with labels that match the preferred label indicated in the user input in step 404. For example, if the labels and the preferred label each include text (e.g., a string) describing the aggregated message response parameter, the preferred label can be compared with the label of a base message to determine whether the respective text of the preferred label matches with the text of the labels corresponding to the base messages stored in memory. This can be accomplished using, for example, a string compare function (e.g., string 1 == string 2). In another

example, if the labels and the preferred label each include an emoji from the Unicode Standard, the Unicode index of the emoji in the preferred label can be compared against the respective Unicode indices of the labels corresponding to the base messages stored in memory to determine whether the respective Unicode indices match.

[0093] As described above, in some implementations, the labeling module can include a confidence score for each label assigned to the base message. The confidence score can be used during step 404 to identify messages that not only have a label that match the preferred label, but also a high confidence score for that label. For example, only base messages having a matching label and a confidence score greater than 0.9 are retrieved from memory. The retrieved base messages can thereafter be transmitted by the platform server to the computing device in step 406 and displayed via the user interface in step 408 for the user to view and consume. In some implementations, the base message and the response messages can be transmitted and displayed together.

[0094] The labels can also be used to provide content to new users of the platform where data on the user's preferences for content is limited or, in some instances, non-existent. In previous social messaging platforms, the lack of data on new users can result in the social messaging platform being unable to provide sufficient content to fill the user's stream. Herein, a new user can be defined as a user having an account for the social messaging platform that is less than or equal to a threshold number of days (e.g., 5 days, 10 days, 30 days) and/or an account that follows less than a threshold number of users (e.g., 1 user, 5 users, 10 users).

[0095] The social messaging platform can be configured to provide content with certain labels without any data on the new user to fill the user's stream. For example, the social messaging platform can be configured to provide messages with "funny," "uplifting," or "insightful" labels by default. The social messaging platform can also provide a way for the user to identify one or more labels when, for example, the user is creating their account with the social messaging platform. For example, the social messaging platform may display a message asking the user to select one or more labels from a list of labels. The label(s) can then be used to identify and deliver content to the new user using the method 400 described above.

[0096] It should be appreciated the recommendation of new content based solely on the labels assigned to the base message in the method 400 is a non-limiting example. In some

implementations, the labels can be used in conjunction with other information associated with user engagement to determine which content to recommend to a user. For example, user engagement mechanisms used to indicate positive engagement (e.g., “liking” a message, indicating a message as a “favorite,” replying to a message) or negative engagement (e.g., “reporting” a message) with previous messages can also be used to determine whether the user is likely to react to a base message in the same manner as indicated by the labels assigned to the base message.

[0097] In another example, the labels disclosed herein can be used together with other labels that provide additional information associated with the base message. For example, the social messaging platform can include labels to indicate the topic of the base message (e.g., news, sports, fishing, tabloids). By combining labels representing user sentiment and topic labels representing a topic of interest, more personalized content can be delivered to each user.

1.3.3 An Example Method for Identifying Base Messages for Health Assessment

[0098] The labels assigned to a base message can also be used to assess the health of the content in the base message. Specifically, the labels can be used identify content with potentially toxic content (e.g., abusive, disturbing, or violent content). Messages with potentially toxic content can be sent to an agent of the social messaging platform for manual review. For example, FIG. 6 shows a method 600 for identifying base messages with potentially toxic content and notifying an agent to review the base message. It should be appreciated the method 600 can follow some of the steps in the method 300 (e.g., step 318). At step 602, an application (e.g., the application 232a or 232b in FIG. 8B) can be executed on a platform server to retrieve a base message with one or more labels from memory (e.g., the memory 214a or the database 216 in FIG. 8B). The application can thereafter be executed to identify the base message has potentially negative toxic content in step 604 by determining whether the labels associated with the base message correspond to labels indicative of a sentiment to toxic content. For example, if the base message has labels of “angry,” “disgust,” or “controversial,” the application determines the base message potentially has toxic content.

[0099] If the application determines the base message does not include labels indicative of toxic content, the application can proceed to retrieve another message from memory for evaluation in step 610 if another message is available. Otherwise, if the application determines the base message does include labels indicative of toxic content, the application can thereafter generate a notification

for an agent to review the base message in step 606. The notification can include, for example, a copy of the base message and/or the response messages and any labels assigned to the base message and/or the response messages. Lastly, at step 608, the notification can be transmitted from the platform server to a computing device associated with the agent to notify the agent to review the base message.

[0100] The method 600 can be triggered in several ways. For example, each time the labeling module labels a base message, the method 600 is automatically executed to evaluate the health of the base message. In another example, if a user “reports” a base message for review, the method 600 can first be executed to verify the base message may contain toxic content before being passed onto an agent for review. Furthermore, a confidence score associated with each base message can also be used to readily identify base messages that have the highest likelihood of containing toxic content. For example, a base message having an “angry” label with a confidence score of 1 is likely to include toxic content due to the response messages universally having a negative sentiment.

1.3.4 An Example Method for Generating Targeted Advertisements using Labels

[0101] The labels can also be used by businesses to evaluate the response of users to certain products, advertisements, and/or promotions posted to the social messaging platform. For example, a company representative can post a base message promoting a new product. In another example, the company can sponsor a user of the social messaging platform to promote their product in a base message posted to the social messaging platform. If the message receives multiple response messages from one group of users, the labels assigned to the base message can be used to quickly evaluate user sentiment towards the base message instead of manually reviewing each response message. This can provide insight whether the base message received positive engagement or negative engagement and can inform the business of other users they can target with same message. For example, a base message with an advertisement that receives a positive sentiment can be sent to users who indicate a preference for content with the same labels as the base message (e.g., using the method 400 described above) or users who respond to base messages with the same labels.

[0102] For example, FIG. 7 shows a method 700 for determining a message with an advertisement receives positive engagement from one group of users and transmitting the message to a second group of users having similar characteristics for the purposes of targeted advertising. This method can be executed, for example, by an application supported by the social messaging platform (e.g.,

the application 232a or 232b in FIG. 8B) or a third-party application that has access to data stored on the social messaging platform (e.g., an application developed by the company responsible for posting the base message). It should be appreciated the method 700 can follow some of the steps in the method 300 (e.g., step 318).

[0103] At step 702, a base message having content containing an advertisement is retrieved from memory along with one or more labels associated with the base message. The advertisement can generally be for a product or a service. The application can thereafter be executed to identify the base message has a positive sentiment in step 704 by determining whether the labels associated with the base message correspond to labels indicative of a positive sentiment. For example, if the base message has labels of “inspiring,” “happy,” or “intrigued,” the application determines the base message received positive engagement from the users who posted the response messages.

[0104] If the application determines the base message did not receive a positive sentiment, the application can proceed to retrieve another message containing an advertisement for evaluation if available in step 712. Otherwise, if the application determines the message did receive a positive sentiment, the application can proceed to identify another group of users different from the users who posted response messages to the base message in step 706.

[0105] This can be accomplished, for example, by first identifying one or more characteristics associated with the users who posted a response message. The characteristics can include, but is not limited to, a geographic location, an age, an ethnicity, an education level, an occupation, one or more topics of interest identified by the user, one or more hobbies identified by the user, groups that the user either created or is a member of on the social messaging platform, and other users that are being followed by the user. These characteristics, in turn, can be used by the application to identify other users of the social messaging platform that share one or more of the same characteristics. Additionally, other data related to the base message can be used to identify users for targeted advertisements, such as users who indicate a preference for content having the same label as the base message or users who respond to other base messages with the same label. Once a group of users are identified having similar characteristics as the users who posted response messages to the base message, the base message can then be transmitted to the respective computing devices of the group of users in step 708 and thereafter displayed on respective user interfaces of the computing devices in step 710.

[0106] In some implementations, the confidence score associated with the base messages can also be used to compare base messages with different advertisements to determine which advertisement to send to users. For example, if a company posts two base messages containing advertisements A and B, both advertisements A and B may receive a positive sentiment (e.g., a “happy” label) from the users who replied to the respective messages. However, the confidence scores for the “happy” label may be higher for advertisement A than advertisement B. Thus, the application evaluating the advertisement can thereafter instruct the platform server to send only messages with advertisement A to users.

2. An Example Social Messaging Platform with a Labeling Module

[0107] FIG. 8A illustrates an example online social messaging platform 200 and example user devices 204a-204n configured to interact with the platform over one or more wired or wireless data communication networks 220. Users 202a-202n of the platform use their user devices, on which client software 206a-206n is installed, to use the platform. A user device can be any Internet-connected computing device, e.g., a laptop or desktop computer, a smartphone, or an electronic tablet. The user device can be connected to the Internet through a mobile network, through an Internet service provider (ISP), or otherwise.

[0108] Each user device is configured with software, which will be referred to as a client or as client software 206a-206n, that in operation can access the platform 200 so that a user can post and receive messages, view, interact with, and create streams of messages and other content items, and otherwise use the service provided by the platform. Generally, the client software 206a-206n can be adapted for operation on different user devices and/or different operating systems. For example, the client software 206a-206n can run on various operating systems including, but not limited to, Google Android™, Apple iOS®, Google Chrome OS™, Apple MacOS®, Microsoft Windows®, and Linux®. The client software 206a-206n can further include web applications (e.g., the client software 206a isn't installed directly onto the user's device, but is rather accessible through a web browser on the user's device).

[0109] Generally, a message posted on the platform contains data representing content provided or selected by the author of the message. The message may be an instance of a container data type storing the content data. The types of data that may be stored in a message include text, graphics, images, video, and computer code, e.g., uniform resource locators (URLs), for example. Messages

can also include key phrases, e.g., hashtags, that can aid in categorizing messages or in linking messages to topics. Messages can also include metadata that may or may not be editable by the message author or account holder, depending on what the platform permits. Examples of message metadata include a time and date of authorship and a geographical location of the user device when the client posted the message. In some implementations, what metadata is provided to the platform by a client is determined by privacy settings controlled by the user or the account holder.

[0110] Messages composed by one account holder may include references to other accounts, other messages, or both. For example, a message may be composed in reply to another message posted by another account or by the user. Messages may also be republications of messages received from another account. Generally, an account referenced in a message may appear as visible content in the message, e.g., as the name of the account, and may also appear as metadata in the message. As a result, the referenced accounts can be interactive in the platform. For example, users may interact with account names that appear in their message stream to navigate to message streams of those accounts. The platform also allows users to designate particular messages as private; a private message will only appear in the message streams of the composing and recipient accounts.

[0111] Generally, messages on the platform are microblog posts, which differ from email messages in a number of ways, for example, in that an author of the post does not necessarily need to specify, or even know, which accounts the platform will select to provide the message to.

[0112] The platform is implemented on one or more servers 210a-210m in one or more locations. Each server is implemented on one or more computers, e.g., on a cluster of computers. The platform, the user devices, or both are configured, as will be described, to implement or perform one or more of the innovative technologies described in this specification. Further information about user devices, clients, servers, and the platform is provided later in this specification.

[0113] One or more of the servers implement a labeling module 230, directed to identifying one or more icons in response messages posted in response to a base message, determining an aggregate message response parameter representing the sentiment of the users who posted the response messages, and labeling the base message based on the aggregate message response parameter. In some implementations, the labeling module 230 can also identify base messages having specific labels for the purposes of recommending new content to users. For example, the platform server 210a can transmit new content to the users 102a-102n via user devices 204a-204n

and the client software 204a-204n can thereafter display the new content via respective user interfaces.

[0114] Aspects related to the use of the labeling module 230 and/or the processing of base messages with assigned labels may be executable by any suitable components of the platform 200 such as, for example, by one or more of the platform servers 210a-210m, and/or by any suitable components of the user devices 204a-204n. For example, FIG. 8B shows an expanded view of the example platform 200. As shown, the platform 200 can include the platform servers 210a and 210m along with a database 216. The database 216 can store one or more messages 240 (e.g., base messages and response messages) and labels 242 assigned to respective messages 240. Unless indicated otherwise, all components of the platform 200 herein can be in communication with each other.

[0115] The platform server 210a can include one or more processors 212a and onboard memory 214a. The onboard memory 214a can store, for example, the labeling module 230 and one or more applications 232a that utilize the labels 242 to perform a function (e.g., identify new content, evaluate the health of a message, identify users for targeted advertising). The processor(s) 212a can execute various processes and/or instructions according to the labeling module 230 and the application(s) 232a. The platform server 210m can also include one or more processors 212b and onboard memory 214b. As shown, the platform server 210m can store one or more applications 232b as well. The applications 232b represent software installed on other servers and/or systems of the platform 200 that can access and use the labels 242.

[0116] The memory 214a, 214b, and/or the database 216 can encompass, for example, a random-access memory (RAM), a memory buffer, a hard drive, a database, an erasable programmable read-only memory (EPROM), an electrically erasable read-only memory (EEPROM), a read-only memory (ROM), Flash memory, and/or so forth. The memory 214a, 214b, and/or the database 216 can store instructions to cause the one or more processors (e.g., the processors 212a or 212b) to execute processes and/or functions associated with the labeling module 230, the application(s) 232a, or the application(s) 232b, or, more generally, the platform 200. The database 216 can further store any suitable content for use with, or generated by, the platform 200 including, but not limited to, one or more connection graphs, a rule repository, and/or the like.

[0117] The one or more processors 212a and 212b can each (independently) be any suitable

processing device configured to run and/or execute a set of instructions or code associated with respective platform servers 210a and 210m and/or the platform 200. Each processor can be, for example, a general-purpose processor, a Field Programmable Gate Array (FPGA), an Application Specific Integrated Circuit (ASIC), a Digital Signal Processor (DSP), and/or the like. The one or more processors 212a and 212b can execute a labeling module 230 directed to identifying icons in one or more response messages posted in response to a base message, determining an aggregate message response parameter based on the identified icons, and assigning one or more labels to the base message representing the aggregate message response parameter.

2.1 User Devices and Clients

[0118] On any particular user device, the client may be a web browser and an HTML (hypertext markup language) document rendered by the web browser. The client may be or include JavaScript code or Java code. Or the client may be dedicated software, e.g., an installed app or installed application, that is designed to work specifically with the platform. The client may include, for example, a Short Messaging Service (SMS) interface, an instant messaging interface, an email-based interface, and HTML-based interface, or an API function-based interface for interacting with the platform.

[0119] A user device can include a camera, microphone, or both, and the client can include or be coupled to software to record pictures, audio, and video. The user device can include both a front-facing, i.e., a user-facing, camera, and a rear-facing camera.

2.2 Platform

[0120] The platform may have many millions of accounts, and anywhere from hundreds of thousands to millions of connections may be established or in use between clients and the platform at any given moment. The accounts may be accounts of individuals, businesses, or other entities, including, e.g., pseudonym accounts, novelty accounts, and so on.

[0121] The platform and client software are configured to enable users to draft messages and to use the platform, over data communication networks, to post messages to the platform and to receive messages posted by other users. The platform and client software are configured to enable users to post other kinds of content, e.g., image, video, or audio content, or a combination of kinds of content, either separately or combined with text messages.

[0122] Optionally, the platform is configured to enable users to define immediate or scheduled sessions with individual or groups of users for audio or audio and video interactions. The platform enables users to specify participation in such sessions using the relationships defined, i.e., in the connection graphs maintained by the platform.

[0123] The platform is configured to deliver content, generally messages, to users in their home feed stream. The messages will generally include messages from accounts the user is following, meaning that the recipient account has registered to receive messages posted by the followed account. The platform generally also includes in the stream messages that the platform determines are likely to be of interest to the recipient, e.g., messages on topics of particular current interest, as represented by the number of messages on the topics posted by platform users, or messages posted on topics of apparent interest to the recipient, as represented by messages the recipient has posted or engaged with, or messages on topics the recipient has expressly identified to the platform as being of interest to the recipient, as well as selected advertisements, public service announcements, promoted content, or the like.

[0124] The platform enables users to send messages directly to one or more other users of the platform, allowing the sender and recipients to have a private exchange of messages. The platform is configured with interfaces through which a client can post messages directed to other users, both synchronously and asynchronously. Thus, users are able to exchange messages in real-time, i.e., with a minimal delay, creating what are essentially live conversations, or to respond to messages posted earlier, on the order of hours or days or even longer.

[0125] The platform also indexes content items and access data that characterizes users' access to content. The platform provides interfaces that enable users to use their clients to search for users, content items, and other entities on the platform.

2.3 Messages

[0126] Generally, a message posted on the platform contains data representing content provided or selected by the author of the message. The message may be an instance of a container data type storing the content data. The types of data that may be stored in a message include text, graphics, images, video, and computer code, e.g., uniform resource locators (URLs), for example. Messages can also include key phrases, e.g., hashtags, that can aid in categorizing messages or in linking messages to topics. Messages can also include metadata that may or may not be editable by the

message author or account holder, depending on what the platform permits. Examples of message metadata include a time and date of authorship and a geographical location of the user device when the client posted the message. In some implementations, what metadata is provided to the platform by a client is determined by privacy settings controlled by the user or the account holder.

[0127] Messages composed by one account holder may include references to other accounts, other messages, or both. For example, a message may be composed in reply to another message posted by another account or by the user. Messages may also be republications of messages received from another account. Generally, an account referenced in a message may appear as visible content in the message, e.g., as the name of the account, and may also appear as metadata in the message. As a result, the referenced accounts can be interactive in the platform. For example, users may interact with account names that appear in their message stream to navigate to message streams of those accounts. The platform also allows users to designate particular messages as private; a private message will only appear in the message streams of the composing and recipient accounts.

[0128] Generally, messages on the platform are microblog posts, which differ from email messages in a number of ways, for example, in that an author of the post does not necessarily need to specify, or even know, which accounts the platform will select to provide the message to.

2.4 Streams

[0129] A stream is an ordered collection of messages on the platform that meet one or more stream criteria. A stream can be defined by the stream criteria to include messages posted by one or more accounts. For example, the contents of a stream for a requesting account holder may include one or more of (i) messages composed by that account holder, (ii) messages composed by the other accounts that the requested account holder follows, (iii) messages authored by other accounts that reference the requested account holder, or (iv) messages sponsored by third parties for inclusion in the account holder's message stream. The messages of a stream may be ordered chronologically by time and date of authorship, or reverse chronologically. Streams may also be ordered in other ways, e.g., according to a computationally predicted relevance to the requesting account holder, or according to some combination of time and relevance score.

[0130] A stream may potentially include a large number of messages. For both processing efficiency and the requesting account holder's viewing convenience, the platform generally identifies a subset of messages meeting the stream criteria to send to a requesting client once the

stream is generated. The remainder of the messages in the stream are maintained in a stream repository and some or all of the remainder are sent in response to subsequent client requests.

2.5 Relationships

[0131] Accounts will generally have relationships with other accounts on the platform. Relationships between accounts of the platform are represented by connection data maintained by the platform, e.g., in the form of data representing one or more connection graphs. The connection data can be maintained in a connection repository. Data repositories of the platform are generally stored in distributed replicas for high throughput and reliability. A connection graph includes nodes representing accounts of the platform and edges connecting the nodes according to the respective relationships between the entities represented by the nodes. A relationship may be any kind of association between accounts, e.g., a following, friending, subscribing, tracking, liking, tagging, or other relationship. The edges of the connection graph may be directed or undirected based on the type of relationship.

[0132] The platform can also represent relationships between accounts and entities other than accounts. For example, when an account belongs to a company, a team, a government, or other organized group, a relationship with that account can also be, for example, a relationship of being a member of the group, having a particular role in the group, or being an expert about the group. The platform can also represent abstract entities, e.g., topics, activities, or philosophies, as entities that can have relationships with accounts and, in some implementations, other entities. Such relationships can also be represented in a common connection graph or in one or more separate connection graphs, as described above.

2.6 Engagements

[0133] The platform records user engagements with messages and maintains, in a message repository, data that describes and represents at least a collection of recent messages as well as the engagements with the messages.

[0134] Engagement data relative to messages includes data representing user activity with respect to messages. Examples of engagement by a user with a message include reposting the message, marking the message to indicate it is a favorite of, liked by, or endorsed by the user, responding to the message, responding to a message with a response having a sentiment determined by the

platform to be positive or negative, quoting the message with further comments, and mentioning or referencing the message.

[0135] Engagement data relative to accounts includes data representing connections between accounts. Examples of engagements by a user with respect to an account include aggregate measures of engagement with messages authored by the account. Other examples include how many followers and followees the account has, i.e., how many other accounts are following the account and how many other accounts the account is following. Other examples include measures of similarity between the groups of followers, the groups of followees, or both, of two accounts, including non-account followees.

[0136] Data about engagements can be represented on the platform as graphs with connections between accounts and messages, and stored in a graph repository.

2.7 Services provided by Platform Servers

[0137] The servers of the platform perform a number of different services that are implemented by software installed and running on the servers. The services will be described as being performed by software modules. In some cases, particular servers may be dedicated to performing one or a few particular services and only have installed those components of the software modules needed for the particular services. Some modules will generally be installed on most or all of the non-special-purpose servers of the platform. In some cases, multiple instances of a module may operate in parallel so as to complete a request for service within a short period of time, so that the platform can respond to users with low latency. The software of each module may be implemented in any convenient form, and parts of a module may be distributed across multiple computers in one or more locations so that the operations of the module are performed by multiple computers running software performing the operations in cooperation with each other. In some implementations, some of the operations of a module are performed by special-purpose hardware.

2.8 Front End Services

[0138] In some implementations, the platform includes numerous different but functionally equivalent front end servers, which are dedicated to managing network connections with remote clients.

[0139] The front end servers provide a variety of interfaces for interacting with different types of

clients. For example, when a web browser accesses the platform, a web interface module in the front end module provides the client access. Similarly, when a client calls an API made available by the platform for such a purpose, an API interface provides the client access.

[0140] The front end servers are configured to communicate with other servers of the platform, which carry out the bulk of the computational processing performed by the platform as a whole.

2.9 Routing Services

[0141] A routing module stores newly composed messages in a message repository. The routing module also stores an identifier for each message. The identifier is used to identify a message that is to be included in a stream. This allows the message to be stored only once and accessed for a variety of different streams without needing to store more than one copy of the message.

2.10 Relationship Graph Services

[0142] A graph module manages connections between accounts, between accounts and entities, and between entities. Connections determine which streams include messages from which accounts. In some implementations, the platform uses unidirectional connections between accounts and streams to allow account holders to subscribe to the message streams of other accounts. A unidirectional connection does not imply any sort of reciprocal relationship. An account holder who establishes a unidirectional connection to receive another account's message stream may be referred to as a "follower," and the act of creating the unidirectional connection is referred to as "following" another account.

[0143] The graph module receives client requests to create and delete unidirectional connections between accounts and updates the connection graph or graphs accordingly. Similarly, for entities that are represented by the platform as entities with which accounts can have relationships, the graph module can also receive client requests to create and delete connections representing account-to-entity relationships.

2.11 Recommendation services

[0144] A recommendation module of the platform can recommend content, accounts, topics, or entities to a user. The recommendation module can be, for example, one of the applications 232a or 232b in FIG. 8B. The recommendation module specifically tailors recommendations to the user.

[0145] A user or a client can generate a request for a recommendation, or another module of the

platform can generate a request for a recommendation on its own, e.g., in order to include a recommendation in a stream being generated for a user. A recommendation can be a call to action, i.e., a suggestion that the user take a particular action, or the recommendation can be the recommended content itself, e.g., a message to include in the stream. The recommendation module can also provide a recommendation in response to a user action that does not explicitly request a recommendation, e.g., interacting with content on the platform in a way that indicates interest.

[0146] The recommendation module makes recommendations using, for example, information users provided about themselves and other data found in the users' profiles, and data about the users' engagements and relationships stored in graph data and otherwise in the platform's repositories.

[0147] To make a recommendation for a user, that user's behavior and other users' behaviors are taken into account. Thus, the relationships and interactions between (i) a user, on the one hand, and (ii) content or users or other entities, on the other hand, are used to make personalized content recommendations for the user. In addition to being presented to the user by a client, recommendations can be provided without going through the client, e.g., in an email, text message, or push notification. Recommendations can also identify, in a personalized way, content popular near a certain geographic location, or real-time trending topics.

[0148] The platform maintains data, especially about live events, with a high degree of currency and for quick access, so that the platform can provide recommendations of current interest, especially during live events.

[0149] In some implementations, the platform presents with a recommendation user-related reasons for the recommendation, e.g., because a message relates to a topic followed by the user or to a trending topic or to a topic trending in the user's location, or because a message had strong engagement among the user's followees, or because the message was endorsed by other users sharing common interests or sharing common followed topics with the user. In some implementations, the platform ranks recommendations according to the reasons for the recommendations, giving preference to recommendations based on endorsements from followees, experts, or celebrities.

2.12 Delivery Services

[0150] A delivery module constructs message streams and provides them to requesting clients, for example, through a front end server. Responding to a request for a stream, the delivery module either generates the stream in real time, or accesses from a stream repository some or all of a stream that has already been generated. The delivery module stores generated streams in the stream repository. An account holder may request any of their own streams, or the streams of any other account that they are permitted to access based on privacy and security settings. If a stream includes a large number of messages, the delivery module generally identifies a subset of the messages to send to a requesting client, in which case the remaining messages are maintained in a stream repository from which more messages are sent upon client request.

2.13 Health and Safety

[0151] The platform includes modules that enable users to filter the content they receive from the platform. For example, users may select settings that cause the platform to filter out sensitive content. The platform also enables a user to control how the user is visible on the platform. For example, the platform enables a user to prevent particular users from following the user, from viewing the user's messages on the platform, from sending messages directed to the user, or from tagging the user in a photo. The platform also enables a user to mute particular users to prevent messages from particular users from being included in any incoming streams, or to block incoming push or SMS notifications from particular users. The platform enables a user to mute another user while continuing to be a follower of the other user.

[0152] In addition, the platform itself can filter out content that is identified by the platform as toxic or abusive, or that originates from accounts identified by the platform as toxic or abusive, with or without a user request to do so. For example, the platform can include a module that executes the method 600 of FIG. 6.

2.14 Account Services

[0153] An account module enables account holders to manage their platform accounts. The account module allows an account holder to manage privacy and security settings, and their connections to other account holders. In particular, a user can choose to be anonymous on the platform. Data about each account is stored in an account repository.

2.15 Engagement Services

[0154] Client software allows account holders receiving a stream to engage, e.g., interact with, comment on, or repost, the messages in the stream. An engagement module receives these engagements and stores them in an engagement repository. Types of engagement include selecting a message for more information regarding the message, selecting a URI (universal resource identifier) or hashtag in a message, reposting the message, or making a message a favorite. Other example engagement types include opening a card included in a message, which presents additional content, e.g., an image, that represents a target of a link in the message, or that links to an application installed on the user device. Account holders may engage further with the additional content, e.g., by playing a video or audio file or by voting in a poll.

[0155] In addition to recording active interactions with messages through explicitly received user input, the engagement module may also record passive interactions with messages. An impression occurs when a client presents the content of a message on a user device. Impression engagements include the mere fact that an impression occurred, as well as other information, e.g., whether a message in a stream appeared on a display of the user device, and how long the message appeared on the display.

[0156] Any engagement stored in the engagement repository may reference the messages, accounts, or streams involved in the engagement.

[0157] In addition to the sentiments disclosed herein, engagements may also include engagements expressing a positive sentiment about a message (“positive engagements”), engagements expressing a negative sentiment about a message (“negative engagements”), engagements that allow an account to receive monetary compensation (“monetizable engagements”), engagements that are expected to result in additional future engagements (“performance engagements”), or engagements that are likely to result in one account holder following another account (“connection engagements”). The negative engagements category includes, for example, engagements dismissing a message or reporting a message as offensive, while the positive engagements category typically includes engagements not in the negative engagements category. Example performance engagements include selecting a URL in a message or expanding a card. Example monetizable engagements include, for example, engagements that result in an eventual purchase or a software application installation on a user device. Generally, categories and types are not coextensive, and

a given type of engagement may fall into more than one category and vice versa.

3. Conclusion

[0158] While various inventive embodiments have been described and illustrated herein, those of ordinary skill in the art will readily envision a variety of other means and/or structures for performing the function and/or obtaining the results and/or one or more of the advantages described herein, and each of such variations and/or modifications is deemed to be within the scope of the inventive embodiments described herein. More generally, those skilled in the art will readily appreciate that all parameters, dimensions, materials, and configurations described herein are meant to be exemplary and that the actual parameters, dimensions, materials, and/or configurations will depend upon the specific application or applications for which the inventive teachings is/are used. Those skilled in the art will recognize or be able to ascertain, using no more than routine experimentation, many equivalents to the specific inventive embodiments described herein. It is, therefore, to be understood that the foregoing embodiments are presented by way of example only and that, within the scope of the appended claims and equivalents thereto, inventive embodiments may be practiced otherwise than as specifically described and claimed. Inventive embodiments of the present disclosure are directed to each individual feature, system, article, material, kit, and/or method described herein. In addition, any combination of two or more such features, systems, articles, materials, kits, and/or methods, if such features, systems, articles, materials, kits, and/or methods are not mutually inconsistent, is included within the inventive scope of the present disclosure.

[0159] Also, various inventive concepts may be embodied as one or more methods, of which an example has been provided. The acts performed as part of the method may be ordered in any suitable way. Accordingly, embodiments may be constructed in which acts are performed in an order different than illustrated, which may include performing some acts simultaneously, even though shown as sequential acts in illustrative embodiments.

[0160] All definitions, as defined and used herein, should be understood to control over dictionary definitions, definitions in documents incorporated by reference, and/or ordinary meanings of the defined terms.

[0161] The indefinite articles “a” and “an,” as used herein in the specification and in the claims, unless clearly indicated to the contrary, should be understood to mean “at least one.”

[0162] The phrase “and/or,” as used herein in the specification and in the claims, should be understood to mean “either or both” of the elements so conjoined, i.e., elements that are conjunctively present in some cases and disjunctively present in other cases. Multiple elements listed with “and/or” should be construed in the same fashion, i.e., “one or more” of the elements so conjoined. Other elements may optionally be present other than the elements specifically identified by the “and/or” clause, whether related or unrelated to those elements specifically identified. Thus, as a non-limiting example, a reference to “A and/or B”, when used in conjunction with open-ended language such as “comprising” can refer, in one embodiment, to A only (optionally including elements other than B); in another embodiment, to B only (optionally including elements other than A); in yet another embodiment, to both A and B (optionally including other elements); etc.

[0163] As used herein in the specification and in the claims, “or” should be understood to have the same meaning as “and/or” as defined above. For example, when separating items in a list, “or” or “and/or” shall be interpreted as being inclusive, i.e., the inclusion of at least one, but also including more than one, of a number or list of elements, and, optionally, additional unlisted items. Only terms clearly indicated to the contrary, such as “only one of” or “exactly one of,” or, when used in the claims, “consisting of,” will refer to the inclusion of exactly one element of a number or list of elements. In general, the term “or” as used herein shall only be interpreted as indicating exclusive alternatives (i.e. “one or the other but not both”) when preceded by terms of exclusivity, such as “either,” “one of,” “only one of,” or “exactly one of.” “Consisting essentially of,” when used in the claims, shall have its ordinary meaning as used in the field of patent law.

[0164] As used herein in the specification and in the claims, the phrase “at least one,” in reference to a list of one or more elements, should be understood to mean at least one element selected from any one or more of the elements in the list of elements, but not necessarily including at least one of each and every element specifically listed within the list of elements and not excluding any combinations of elements in the list of elements. This definition also allows that elements may optionally be present other than the elements specifically identified within the list of elements to which the phrase “at least one” refers, whether related or unrelated to those elements specifically identified. Thus, as a non-limiting example, “at least one of A and B” (or, equivalently, “at least one of A or B,” or, equivalently “at least one of A and/or B”) can refer, in one embodiment, to at least one, optionally including more than one, A, with no B present (and optionally including

elements other than B); in another embodiment, to at least one, optionally including more than one, B, with no A present (and optionally including elements other than A); in yet another embodiment, to at least one, optionally including more than one, A, and at least one, optionally including more than one, B (and optionally including other elements); etc.

[0165] This specification uses the term “configured to” in connection with systems, apparatus, and computer program components. That a system of one or more computers is configured to perform particular operations or actions means that the system has installed on it software, firmware, hardware, or a combination of them that in operation cause the system to perform those operations or actions. That one or more computer programs is configured to perform particular operations or actions means that the one or more programs include instructions that, when executed by data processing apparatus, cause the apparatus to perform those operations or actions. That special-purpose logic circuitry is configured to perform particular operations or actions means that the circuitry has electronic logic that performs those operations or actions.

[0166] In any claims, as well as in the specification above, all transitional phrases such as “comprising,” “including,” “carrying,” “having,” “containing,” “involving,” “holding,” “composed of,” and the like are to be understood to be open-ended, i.e., to mean including but not limited to. Only the transitional phrases “consisting of” and “consisting essentially of” shall be closed or semi-closed transitional phrases, respectively, as set forth in the United States Patent Office Manual of Patent Examining Procedures, Section 2111.03.

CLAIMS

1. A social messaging system for labeling messages, the system comprising:
 - a first computing device associated with a first user of a plurality of users;
 - a set of second computing devices associated with a set of second users of the plurality of users;
 - a platform server comprising at least one processor and memory;
 - a database; and
 - a computer-readable storage medium, stored in the memory, storing instructions that, when executed, cause the at least one processor to:
 - receive, from the first computing device, a first message for broadcast to the set of second computing devices;
 - display, in a timeline of each second computing device of the set of second computing devices, the first message;
 - receive, from the set of second computing devices, a set of second messages in response to the first message; and
 - analyze the set of second messages by:
 - identifying a set of icons in the set of second messages, the set of icons being selected from a predetermined list of icons; and
 - determining a first aggregate message response parameter based on the set of icons, the first aggregate message response parameter being a digital representation of a first sentiment of the set of second users towards the first message;
 - assign a first label from a set of labels to the first message based on the first aggregate message response parameter, the first label providing one of a visual or a textual representation of the first sentiment of the second users; and
 - store the first label as associated with the first message in the database.
2. The system of claim 1, wherein:
 - the first message includes first content; and
 - at least one second message in the set of second messages includes second content that quotes the first content.

3. The system of claim 1, wherein the predetermined list of icons includes at least one of an emoji or an emoticon.
4. The system of claim 1, wherein:
 - each label of the set of labels corresponds to at least one icon of the predetermined list of icons; and
 - the predetermined list of icons includes at least one combination of icons having two or more icons, the combination associated with a corresponding label in the set of labels.
5. The system of claim 1, wherein:
 - the set of labels is arranged into an indexed array; and
 - the first aggregate message response parameter is an index value that references the first label in the set of labels.
6. The system of claim 1, wherein the instructions of the computer-readable storage medium further cause the at least one processor to:
 - before identifying the set of icons, check a number of second messages of the set of second messages is above a predetermined threshold,
 - wherein identifying the set of icons is executed only when the number of second messages is above the predetermined threshold.
7. The system of claim 6, wherein the predetermined threshold is equal to 500.
8. The system of claim 6, wherein:
 - the number of second messages increases over time; and
 - the at least one processor executes the instructions of the computer-readable storage medium in runtime such that the at least one processor checks the number of second messages of the set of second messages every time after at least one of a predetermined period of time or a predetermined increment in the number of second messages until the number of second messages is above the predetermined threshold.

9. The system of claim 8, wherein the at least one processor checks the number of second messages of the set of second messages every time after the predetermined period of time, the predetermined period of time being equal to one of 1 hour, 1 day, 1 week, or 1 month.

10. The system of claim 8, wherein the at least one processor checks the number of second messages of the set of second messages every time after the predetermined increment in the number of second message, the predetermined increment being equal to one of 5, 10, or 100.

11. The system of claim 1, wherein identifying the set of icons in the set of second messages comprises, for each second message of the set of second messages:

identifying one or more icons present in a content portion of that second message, the one or more icons being selected from the predetermined list of icons; and
adding the one or more icons to the set of icons.

12. The system of claim 1, wherein identifying the set of icons the set of second messages comprises, for each second message of the set of second messages:

analyzing, by a natural language processing model, textual content from a content portion of that second message to determine one or more sentiments expressed in the textual content of that second message;
mapping each sentiment of the one or more sentiments to one or more icons selected from the predetermined list of icons; and
adding the one or more icons to the set of icons.

13. The system of claim 1, wherein determining the first aggregate message response parameter based on the set of icons comprises:

determining a count for each unique icon present in the set of icons;
identifying a first icon of the set of icons as having a first count that is highest of all icons of the set of icons, the first icon corresponding to the first sentiment of the set of second users;
and
determining the first aggregate message response parameter corresponding to the first

icon.

14. The system of claim 13, wherein the instructions of the computer-readable storage medium further cause the at least one processor to:

analyze the set of second messages by:

identifying a second icon of the set of icons as having a second count that is second highest count of all icons of the set of icons, the second icon corresponding to a second sentiment of the set of second users; and

determining a second aggregate message response parameter corresponding to the second icon, the second aggregate message response parameter being a digital representation of the second sentiment of the set of second users towards the first message;

assign a second label from the set of labels to the first message based on the second aggregate message response parameter, the second label providing one of a visual or a textual representation of the second sentiment of the second users; and

store the second label as associated with the first message in the database.

15. The system of claim 13, wherein the instructions of the computer-readable storage medium further cause the at least one processor to:

before assigning the first label, determine a first confidence score associated with the first aggregate message response parameter, the first confidence score being defined as the first count divided by a sum of the count for all icons.

16. The system of claim 15, wherein the instructions of the computer-readable storage medium cause the at least one processor to assign the first label occurs when the first confidence score is above a predetermined threshold.

17. The system of claim 16, wherein the instructions of the computer-readable storage medium further cause the at least one processor to:

analyze the set of second messages by:

identifying a second icon of the set of icons as having a second count that is

second highest count of all icons of the set of icons, the second icon corresponding to a second sentiment of the set of second users;

determining a second aggregate message response parameter corresponding to the second icon, the second aggregate message response parameter being a digital representation of the second sentiment of the set of second users towards the first message; and

determining a second confidence score associated with the second aggregate message response parameter, the second confidence score being defined as the second count divided by the sum of the count for all icons;

assign a second label from the set of labels to the first message based on the second aggregate message response parameter when the second confidence score is above the predetermined threshold, the second label providing one of a visual or a textual representation of the second sentiment of the second users; and

store the second label as associated with the first message in the database.

18. The system of claim 16, wherein the predetermined threshold is equal to 30%.

19. The system of claim 1, wherein the instructions of the computer-readable storage medium further cause the at least one processor to:

update a display of the first message for the first user and the set of second users to include an indication of the first label.

20. The system of claim 1, further comprising:

a third computing device associated with a third user of the plurality of users, wherein the instructions of the computer-readable storage medium further cause the at least one processor to:

receive, from the third computing device, a user input indicating a second label;

retrieve the first message and the first label from the database;

determine the first label matches the second label; and

display, in a timeline of the third user, the first message.

21. The system of claim 20, wherein:
the first label is a first string;
the second label is a second string; and
determining the first label matches the second label comprises:
determining the first string matches the second string by comparing the first string to the second string.
22. The system of claim 20, wherein:
the first label is a first icon having a first Unicode index;
the second label is a second icon having a second Unicode index; and
determining the first label matches the second label comprises:
determining the first Unicode index is the same as the second Unicode index.
23. A method for labeling messages of a social messaging system, the method comprising:
receiving, from a first computing device associated with a first user of a plurality of users of the social messaging system, a first message for broadcast to a set of second computing devices associated with a set of second users of the plurality of users;
displaying, in a timeline of each second computing device of the set of second computing devices, the first message;
receiving, from the set of second computing devices, a set of second messages in response to the first message; and
analyzing the set of second messages by:
identifying a set of icons in the set of second messages, the set of icons being selected from a predetermined list of icons; and
determining a first aggregate message response parameter based on the set of icons, the first aggregate message response parameter being a digital representation of a first sentiment of the set of second users towards the first message;
assigning a first label from a set of labels to the first message based on the first aggregate message response parameter, the first label providing one of a visual or a textual representation of the first sentiment of the second users; and
storing the first label as associated with the first message in a database.

24. The method of claim 23, wherein:
the first message includes first content; and
at least one second message in the set of second messages includes second content that quotes the first content.
25. The method of claim 23, wherein the predetermined list of icons includes at least one of an emoji or an emoticon.
26. The method of claim 23, wherein:
each label of the set of labels corresponds to at least one icon of the predetermined list of icons; and
the predetermined list of icons includes at least one combination of icons having two or more icons, the combination associated with a corresponding label in the set of labels.
27. The method of claim 23, wherein:
the set of labels is arranged into an indexed array; and
the first aggregate message response parameter is an index value that references the first label in the set of labels.
28. The method of claim 23, wherein:
the analyzing further includes, before identifying the set of icons, checking a number of second messages of the set of second messages is above a predetermined threshold; and
the identifying the set of icons is executed only when the number of second messages is above the predetermined threshold.
29. The method of claim 28, wherein the predetermined threshold is equal to 500.
30. The method of claim 28, wherein:
the number of second messages increases over time;
the method is executed in runtime; and

the checking the number of second messages of the set of second messages repeatedly occurs after at least one of a predetermined period of time or a predetermined increment in the number of second messages until the number of second messages is above the predetermined threshold.

31. The method of claim 30, wherein the checking the number of second messages of the set of second messages repeatedly occurs after the predetermined period of time, the predetermined period of time being equal to one of 1 hour, 1 day, 1 week, or 1 month.

32. The method of claim 30, wherein the checking the number of second messages of the set of second messages repeatedly occurs after the predetermined increment in the number of second message, the predetermined increment being equal to one of 5, 10, or 100.

33. The method of claim 23, wherein identifying the set of icons in the set of second messages comprises, for each second message of the set of second messages:

identifying one or more icons present in a content portion of that second message, the one or more icons being selected from the predetermined list of icons; and
adding the one or more icons to the set of icons.

34. The method of claim 23, wherein identifying the set of icons the set of second messages comprises, for each second message of the set of second messages:

analyzing, by a natural language processing model, textual content from a content portion of that second message to determine one or more sentiments expressed in the textual content of that second message;

mapping each sentiment of the one or more sentiments to one or more icons selected from the predetermined list of icons; and

adding the one or more icons to the set of icons.

35. The method of claim 23, wherein determining the first aggregate message response parameter based on the set of icons comprises:

determining a count for each unique icon present in the set of icons;

identifying a first icon of the set of icons as having a first count that is highest of all icons of the set of icons, the first icon corresponding to the first sentiment of the set of second users; and

determining the first aggregate message response parameter corresponding to the first icon.

36. The method of claim 35, wherein:

the analyzing the set of second messages further comprises:

identifying a second icon of the set of icons as having a second count that is second highest count of all icons of the set of icons, the second icon corresponding to a second sentiment of the set of second users; and

determining a second aggregate message response parameter corresponding to the second icon, the second aggregate message response parameter being a digital representation of the second sentiment of the set of second users towards the first message; and

the method further comprises:

assigning a second label from the set of labels to the first message based on the second aggregate message response parameter, the second label providing one of a visual or a textual representation of the second sentiment of the second users; and

storing the second label as associated with the first message in the database.

37. The method of claim 35, further comprising, before assigning the first label:

determining a first confidence score associated with the first aggregate message response parameter, the first confidence score being defined as the first count divided by a sum of the count for all icons.

38. The method of claim 37, wherein assigning the first label occurs when the first confidence score is above a predetermined threshold.

39. The method of claim 38, wherein:

the analyzing the set of second messages further comprises:

identifying a second icon of the set of icons as having a second count that is second highest count of all icons of the set of icons, the second icon corresponding to a second sentiment of the set of second users;

determining a second aggregate message response parameter corresponding to the second icon, the second aggregate message response parameter being a digital representation of the second sentiment of the set of second users towards the first message; and

determining a second confidence score associated with the second aggregate message response parameter, the second confidence score being defined as the second count divided by the sum of the count for all icons; and
the method further comprises:

assigning a second label from the set of labels to the first message based on the second aggregate message response parameter when the second confidence score is above the predetermined threshold, the second label providing one of a visual or a textual representation of the second sentiment of the second users; and

storing the second label as associated with the first message in the database.

40. The method of claim 38, wherein the predetermined threshold is equal to 30%.

41. The method of claim 23, further comprising:

updating a display of the first message for the first user and the set of second users to include an indication of the first label.

42. The method of claim 23, further comprising:

receiving, from a third computing device associated with a third user of the plurality of users of the social messaging system, a user input indicating a second label;

retrieving the first message and the first label from the database;

determining the first label matches the second label; and

displaying, in a timeline of the third user, the first message.

43. The method of claim 42, wherein:

the first label is a first string;

the second label is a second string; and

determining the first label matches the second label comprises:

determining the first string matches the second string by comparing the first string to the second string.

44. The method of claim 42, wherein:

the first label is a first icon having a first Unicode index;

the second label is a second icon having a second Unicode index; and

determining the first label matches the second label comprises:

determining the first Unicode index is the same as the second Unicode index.

45. A non-transitory computer-readable storage medium storing instructions for labeling messages of a social messaging system in a software application executing on a social messaging platform, the instructions, when executed cause, at least one processor to:

receive, from a first computing device associated with a first user of a plurality of users of the social messaging system, a first message for broadcast to a set of second computing devices associated with a set of second users of the plurality of users;

display, in a timeline of each second computing device of the set of second computing devices, the first message;

receive, from the set of second computing devices, a set of second messages in response to the first message; and

analyze the set of second messages by:

identifying a set of icons in the set of second messages, the set of icons being selected from a predetermined list of icons; and

determining a first aggregate message response parameter based on the set of icons, the first aggregate message response parameter being a digital representation of a first sentiment of the set of second users towards the first message;

assign a first label from a set of labels to the first message based on the first aggregate message response parameter, the first label providing one of a visual or a textual representation of the first sentiment of the second users; and

store the first label as associated with the first message in a database.

46. The non-transitory computer-readable storage medium of claim 45, wherein:
the first message includes first content; and
at least one second message in the set of second messages includes second content that quotes the first content.
47. The non-transitory computer-readable storage medium of claim 45, wherein the predetermined list of icons includes at least one of an emoji or an emoticon.
48. The non-transitory computer-readable storage medium of claim 45, wherein:
each label of the set of labels corresponds to at least one icon of the predetermined list of icons; and
the predetermined list of icons includes at least one combination of icons having two or more icons, the combination associated with a corresponding label in the set of labels.
49. The non-transitory computer-readable storage medium of claim 45, wherein:
the set of labels is arranged into an indexed array; and
the first aggregate message response parameter is an index value that references the first label in the set of labels.
50. The non-transitory computer-readable storage medium of claim 45, further comprising instructions that, when executed, cause the at least one processor to:
before identifying the set of icons, check a number of second messages of the set of second messages is above a predetermined threshold,
wherein identifying the set of icons is executed only when the number of second messages is above the predetermined threshold.
51. The non-transitory computer-readable storage medium of claim 50, wherein the predetermined threshold is equal to 500.

52. The non-transitory computer-readable storage medium of claim 50, wherein:
the number of second messages increases over time; and
the instructions of the non-transitory computer-readable storage medium, when executed, are executed in runtime causing the at least one processor to check the number of second messages of the set of second messages every time after at least one of a predetermined period of time or a predetermined increment in the number of second messages until the number of second messages is above the predetermined threshold.

53. The non-transitory computer-readable storage medium of claim 52, wherein the instructions of the non-transitory computer-readable storage medium, when executed, cause the at least one processor to check the number of second messages of the set of second messages every time after the predetermined period of time, the predetermined period of time being equal to one of 1 hour, 1 day, 1 week, or 1 month.

54. The non-transitory computer-readable storage medium of claim 52, wherein the instructions of the non-transitory computer-readable storage medium, when executed, cause the at least one processor to check the number of second messages of the set of second messages every time after the predetermined increment in the number of second message, the predetermined increment being equal to one of 5, 10, or 100.

55. The non-transitory computer-readable storage medium of claim 45, wherein identifying the set of icons in the set of second messages comprises, for each second message of the set of second messages:

identifying one or more icons present in a content portion of that second message, the one or more icons being selected from the predetermined list of icons; and
adding the one or more icons to the set of icons.

56. The non-transitory computer-readable storage medium of claim 45, wherein identifying the set of icons the set of second messages comprises, for each second message of the set of second messages:

analyzing, by a natural language processing model, textual content from a content portion

of that second message to determine one or more sentiments expressed in the textual content of that second message;

mapping each sentiment of the one or more sentiments to an icon selected from the predetermined list of icons; and

adding the one or more icons to the set of icons.

57. The non-transitory computer-readable storage medium of claim 45, wherein determining the first aggregate message response parameter based on the set of icons comprises:

determining a count for each unique icon present in the set of icons;

identifying a first icon of the set of icons as having a first count that is highest of all icons of the set of icons, the first icon corresponding to the first sentiment of the set of second users; and

determining the first aggregate message response parameter corresponding to the first icon.

58. The non-transitory computer-readable storage medium of claim 57, further comprising instructions that, when executed, cause the at least one processor to:

analyze the set of second messages by:

identifying a second icon of the set of icons as having a second count that is second highest count of all icons of the set of icons, the second icon corresponding to a second sentiment of the set of second users; and

determining a second aggregate message response parameter corresponding to the second icon, the second aggregate message response parameter being a digital representation of the second sentiment of the set of second users towards the first message;

assign a second label from the set of labels to the first message based on the second aggregate message response parameter, the second label providing one of a visual or a textual representation of the second sentiment of the second users; and

store the second label as associated with the first message in the database.

59. The non-transitory computer-readable storage medium of claim 57, further comprising

instructions that, when executed, cause the at least one processor to:

before assigning the first label, determine a first confidence score associated with the first aggregate message response parameter, the first confidence score being defined as the first count divided by a sum of the count for all icons.

60. The non-transitory computer-readable storage medium of claim 59, further comprising instructions that, when executed, cause the at least one processor to:

assign the first label occurs when the first confidence score is above a predetermined threshold.

61. The non-transitory computer-readable storage medium of claim 60, further comprising instructions that, when executed, cause the at least one processor to:

analyze the set of second messages by:

identifying a second icon of the set of icons as having a second count that is second highest count of all icons of the set of icons, the second icon corresponding to a second sentiment of the set of second users;

determining a second aggregate message response parameter corresponding to the second icon, the second aggregate message response parameter being a digital representation of the second sentiment of the set of second users towards the first message; and

determining a second confidence score associated with the second aggregate message response parameter, the second confidence score being defined as the second count divided by the sum of the count for all icons;

assign a second label from the set of labels to the first message based on the second aggregate message response parameter when the second confidence score is above the predetermined threshold, the second label providing one of a visual or a textual representation of the second sentiment of the second users; and

store the second label as associated with the first message in the database.

62. The non-transitory computer-readable storage medium of claim 60, wherein the predetermined threshold is equal to 30%.

63. The non-transitory computer-readable storage medium of claim 45, further comprising instructions that, when executed, cause the at least one processor to:

update a display of the first message for the first user and the set of second users to include an indication of the first label.

64. The non-transitory computer-readable storage medium of claim 45, further comprising instructions that, when executed, cause the at least one processor to:

receive, from the third computing device, a user input indicating a second label;
retrieve the first message and the first label from the database;
determine the first label matches the second label; and
display, in a timeline of the third user, the first message.

65. The non-transitory computer-readable storage medium of claim 64, wherein:

the first label is a first string;

the second label is a second string; and

determining the first label matches the second label comprises:

determining the first string matches the second string by comparing the first string to the second string.

66. The non-transitory computer-readable storage medium of claim 64, wherein:

the first label is a first icon having a first Unicode index;

the second label is a second icon having a second Unicode index; and

determining the first label matches the second label comprises:

determining the first Unicode index is the same as the second Unicode index.

67. A method for labeling messages of a social messaging system, the method comprising:

receiving, from a first computing device associated with a first user of a plurality of users of the social messaging system, a first message for broadcast to a set of second computing devices associated with a set of second users of the plurality of users;

displaying, in a timeline of each second computing device of the set of second computing

devices, the first message;

receiving, from the set of second computing devices, a set of second messages in response to the first message;

analyzing the set of second messages by, for each second message of the set of second messages:

identifying one or more first icons present in a content portion of that second message, the one or more first icons being selected from a predetermined list of icons;

analyzing, by a natural language processing model, textual content from the content portion of that second message to determine one or more sentiments expressed in the textual content of that second message;

mapping each sentiment of the one or more sentiments to one or more second icons selected from the predetermined list of icons, the one or more first icons and the one or more second icons forming a set of icons; and

determining a first aggregate message response parameter based on the set of icons, the first aggregate message response parameter being a digital representation of a first sentiment of the set of second users towards the first message;

assigning a first label from a set of labels to the first message based on the first aggregate message response parameter, the first label providing one of a visual or a textual representation of the first sentiment of the second users; and

storing the first label as associated with the first message in a database.

68. The method of claim 67, wherein:

the first message includes first content; and

at least one second message in the set of second messages includes second content that quotes the first content.

69. The method of claim 67, wherein the predetermined list of icons includes at least one of an emoji or an emoticon.

70. The method of claim 67, wherein:

each label of the set of labels corresponds to at least one icon of the predetermined list of

icons; and

the predetermined list of icons includes at least one combination of icons having two or more icons, the combination associated with a corresponding label in the set of labels.

71. The method of claim 67, wherein:

the set of labels is arranged into an indexed array; and

the first aggregate message response parameter is an index value that references the first label in the set of labels.

72. The method of claim 67, wherein:

the analyzing further includes, before identifying the set of icons, checking a number of second messages of the set of second messages is above a predetermined threshold; and

the identifying the set of icons is executed only when the number of second messages is above the predetermined threshold.

73. The method of claim 72, wherein the predetermined threshold is equal to 500.

74. The method of claim 72, wherein:

the number of second messages increases over time;

the method is executed in runtime; and

the checking the number of second messages of the set of second messages repeatedly occurs after at least one of a predetermined period of time or a predetermined increment in the number of second messages until the number of second messages is above the predetermined threshold.

75. The method of claim 74, wherein the checking the number of second messages of the set of second messages repeatedly occurs after the predetermined period of time, the predetermined period of time being equal to one of 1 hour, 1 day, 1 week, or 1 month.

76. The method of claim 74, wherein the checking the number of second messages of the set of second messages repeatedly occurs after the predetermined increment in the number of second

message, the predetermined increment being equal to one of 5, 10, or 100.

77. The method of claim 67, further comprising:

updating a display of the first message for the first user and the set of second users to include an indication of the first label.

78. The method of claim 67, further comprising:

receiving, from a third computing device associated with a third user of the plurality of users of the social messaging system, a user input indicating a second label;

retrieving the first message and the first label from the database;

determining the first label matches the second label; and

displaying, in a timeline of the third user, the first message.

79. The method of claim 78, wherein:

the first label is a first string;

the second label is a second string; and

determining the first label matches the second label comprises:

determining the first string matches the second string by comparing the first string to the second string.

80. The method of claim 78, wherein:

the first label is a first icon having a first Unicode index;

the second label is a second icon having a second Unicode index; and

determining the first label matches the second label comprises:

determining the first Unicode index is the same as the second Unicode index.

81. A method for labeling messages of a social messaging system, the method comprising:

receiving, from a first computing device associated with a first user of a plurality of users of the social messaging system, a first message for broadcast to a set of second computing devices associated with a set of second users of the plurality of users;

displaying, in a timeline of each second computing device of the set of second computing

devices, the first message;

receiving, from the set of second computing devices, a set of second messages in response to the first message;

analyzing the set of second messages by:

identifying a set of icons in the set of second messages, the set of icons being selected from a predetermined list of icons;

determining a count for each unique icon present in the set of icons;

identifying a first icon of the set of icons as having a first count that is highest of all icons of the set of icons, the first icon corresponding to the first sentiment of the set of second users; and

determining the first aggregate message response parameter corresponding to the first icon;

assigning a first label from a set of labels to the first message based on the first aggregate message response parameter, the first label providing one of a visual or a textual representation of the first sentiment of the second users; and

storing the first label as associated with the first message in a database,

wherein the predetermined list of icons includes at least one of an emoji or an emoticon.

82. The method of claim 81, wherein:

the analyzing the set of second messages further comprises:

identifying a second icon of the set of icons as having a second count that is second highest count of all icons of the set of icons, the second icon corresponding to a second sentiment of the set of second users; and

determining a second aggregate message response parameter corresponding to the second icon, the second aggregate message response parameter being a digital representation of the second sentiment of the set of second users towards the first message; and

the method further comprises:

assigning a second label from the set of labels to the first message based on the second aggregate message response parameter, the second label providing one of a visual or a textual representation of the second sentiment of the second users; and

storing the second label as associated with the first message in the database.

83. The method of claim 82, further comprising, before assigning the first label:
determining a first confidence score associated with the first aggregate message response parameter, the first confidence score being defined as the first count divided by a sum of the count for all icons.
84. The method of claim 83, wherein assigning the first label occurs when the first confidence score is above a predetermined threshold.
85. The method of claim 84, wherein:
the analyzing the set of second messages further comprises:
identifying a second icon of the set of icons as having a second count that is second highest count of all icons of the set of icons, the second icon corresponding to a second sentiment of the set of second users;
determining a second aggregate message response parameter corresponding to the second icon, the second aggregate message response parameter being a digital representation of the second sentiment of the set of second users towards the first message; and
determining a second confidence score associated with the second aggregate message response parameter, the second confidence score being defined as the second count divided by the sum of the count for all icons; and
the method further comprises:
assigning a second label from the set of labels to the first message based on the second aggregate message response parameter when the second confidence score is above the predetermined threshold, the second label providing one of a visual or a textual representation of the second sentiment of the second users; and
storing the second label as associated with the first message in the database.
86. The method of claim 84, wherein the predetermined threshold is equal to 30%.

87. The method of claim 81, wherein:
the first message includes first content; and
at least one second message in the set of second messages includes second content that quotes the first content.
88. The method of claim 81, wherein:
each label of the set of labels corresponds to at least one icon of the predetermined list of icons; and
the predetermined list of icons includes at least one combination of icons having two or more icons, the combination associated with a corresponding label in the set of labels.
89. The method of claim 81, wherein:
the set of labels is arranged into an indexed array; and
the first aggregate message response parameter is an index value that references the first label in the set of labels.
90. The method of claim 81, wherein:
the analyzing further includes, before identifying the set of icons, checking a number of second messages of the set of second messages is above a predetermined threshold; and
the identifying the set of icons is executed only when the number of second messages is above the predetermined threshold.
91. The method of claim 90, wherein the predetermined threshold is equal to 500.
92. The method of claim 90, wherein:
the number of second messages increases over time;
the method is executed in runtime; and
the checking the number of second messages of the set of second messages repeatedly occurs after at least one of a predetermined period of time or a predetermined increment in the number of second messages until the number of second messages is above the predetermined threshold.

93. The method of claim 92, wherein the checking the number of second messages of the set of second messages repeatedly occurs after the predetermined period of time, the predetermined period of time being equal to one of 1 hour, 1 day, 1 week, or 1 month.

94. The method of claim 92, wherein the checking the number of second messages of the set of second messages repeatedly occurs after the predetermined increment in the number of second message, the predetermined increment being equal to one of 5, 10, or 100.

95. The method of claim 81, further comprising:
updating a display of the first message for the first user and the set of second users to include an indication of the first label.

96. The method of claim 81, further comprising:
receiving, from a third computing device associated with a third user of the plurality of users of the social messaging system, a user input indicating a second label;
retrieving the first message and the first label from the database;
determining the first label matches the second label; and
displaying, in a timeline of the third user, the first message.

97. The method of claim 96, wherein:
the first label is a first string;
the second label is a second string; and
determining the first label matches the second label comprises:
determining the first string matches the second string by comparing the first string to the second string.

98. The method of claim 96, wherein:
the first label is a first icon having a first Unicode index;
the second label is a second icon having a second Unicode index; and
determining the first label matches the second label comprises:

determining the first Unicode index is the same as the second Unicode index.




FIG. 1A

User	Icons in Reply Message 190a	Mapped Icons in Reply Message	Representative Icon	Representative Label
2		None		Funny
3		None		Funny
4	None 190e			Funny
5		None		Triumphant

FIG. 1B




FIG. 1C




FIG. 2




FIG. 3A




FIG. 3B




FIG. 4




FIG. 5




FIG. 6




FIG. 7




FIG. 8A




FIG. 8B

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US 22/32039

A. CLASSIFICATION OF SUBJECT MATTER

IPC - INV. H04L 51/52 (2022.01)

CPC - INV. H04L 51/52, H04L 51/21; ADD. G06Q 10/107, H04L 51/216

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)

See Search History document

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

See Search History document

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

See Search History document

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2017/0220677 A1 (Facebook, Inc.), 03 August 2017 (03.08.2017), entire document, especially Abstract and para [0027], [0055]-[0057], [0067], [0071]-[0072], [0081], [0126] and [0132].	1-98
Y	US 2017/0185581 A1 (Machine Zone, Inc.), 29 June 2017 (29.06.2017), entire document, especially Abstract and para [0039], [0055], [0127].	1-98
Y	Novak et al. "Sentiment of Emojis", 07 December 2015 (07.12.2015) [online], [retrieved 2022-09-19], Retrieved from the Internet <URL:https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0144296>, entire document, especially pg. 3, para 5-pg.3, para 6.	13-18, 35-40, 57-62, 81-98

☐ 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

"D" document cited by the applicant in the international application

"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

20 September 202 (20.09.2022)

Date of mailing of the international search report

OCT 11 2022

Name and mailing address of the ISA/US

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents
P.O. Box 1450, Alexandria, Virginia 22313-1450

Facsimile No. 571-273-8300

Authorized officer

Kari Rodriguez

Telephone No. PCT Helpdesk: 571-272-4300