



- (51) International Patent Classification:
G06Q 30/02 (2012.01) G06F 17/30 (2006.01)
G06Q 50/00 (2012.01)
- (21) International Application Number:
PCT/US2015/064354
- (22) International Filing Date:
8 December 2015 (08.12.2015)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
14/570,520 15 December 2014 (15.12.2014) US
- (71) Applicant: MICROSOFT TECHNOLOGY LICENSING, LLC [US/US]; Attn: Patent Group Docketing (Bldg. 8/1000), One Microsoft Way, Redmond, Washington 98052-6399 (US).

- (72) Inventors: KICIMAN, Emre Mehmet; Microsoft Technology Licensing, LLC, Attn: Patent Group Docketing (Bldg. 8/1000), One Microsoft Way, Redmond, Washington 98052-6399 (US). BENNETT, Paul Nathan; Microsoft Technology Licensing, LLC, Attn: Patent Group Docketing (Bldg. 8/1000), One Microsoft Way, Redmond, Washington 98052-6399 (US). TEEVAN, Jaime Brooks; Microsoft Technology Licensing, LLC, Attn: Patent Group Docketing (Bldg. 8/1000), One Microsoft Way, Redmond, Washington 98052-6399 (US). DUMAIS, Susan Theresa; Microsoft Technology Licensing, LLC, Attn: Patent Group Docketing (Bldg. 8/1000), One Microsoft Way, Redmond, Washington 98052-6399 (US).
- (74) Agents: MINHAS, Sandip et al.; Microsoft Corporation, Attn: Patent Group Docketing (Bldg. 8/1000), One Microsoft Way, Redmond, Washington 98052-6399 (US).
- (81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.
- (84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH,

[Continued on next page]

(54) Title: MODELING ACTIONS, CONSEQUENCES AND GOAL ACHIEVEMENT FROM SOCIAL MEDIA AND OTHER DIGITAL TRACES

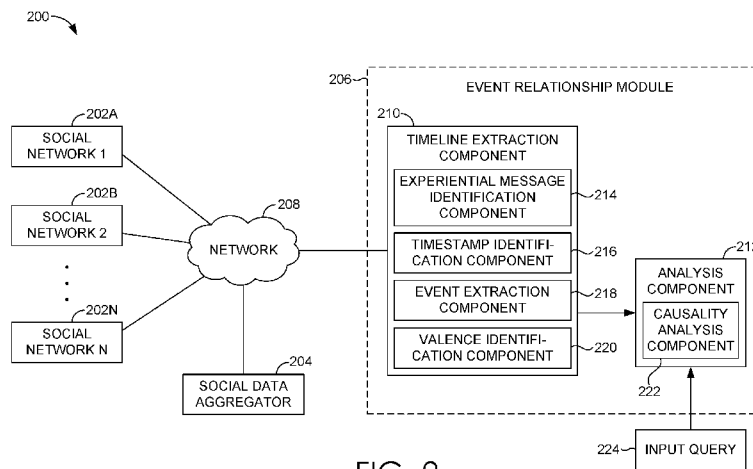


FIG. 2

(57) Abstract: Systems, methods, and computer storage media are provided for analyzing a large amount of social media data from a large population of social media users and constructing correlational data models between one or more events that occur within each user's timeline. Social media posts directed to personal experiences of a large number of social media users are extracted. Event timelines are generated for each of the social media users, based on their personal experiences. The event timelines are analyzed with a particular event of interest to measure correlations between events occurring within the timelines and the particular event of interest. Using the measured correlations, a correlational data model is thereby constructed. The correlational data model may be used for application to decision-making calculations by one or more systems in an active or passive manner.

WO 2016/099983 A1

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).

Declarations under Rule 4.17:

- *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*
- *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

Published:

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

MODELING ACTIONS, CONSEQUENCES AND GOAL ACHIEVEMENT FROM SOCIAL MEDIA AND OTHER DIGITAL TRACES

BACKGROUND OF THE INVENTION

5 [0001] Technology-assisted decision making is becoming more commonplace in the lives of modern society. Users rely on technologies to research product and/or service reviews on-demand, harnessing data from sources such as social media, blogs, mobile applications, etc. These existing technologies may provide very specific empirical data to users for making better-informed decisions. Many of these technologies, however, are
10 limited to provide information to the user when the users know exactly what they are looking for. Social media, particularly, serves as an enormous resource of publicly available data, whereby users share their personal experiences via social media posts. The data gathered from social media may cover a broad scope of personal experiences related to a particular event, action, goal, product, service, or other subject of interest. Users,
15 however, are generally limited to parsing social media data using simple search queries that narrow the scope of data available for review based on keywords of interest.

SUMMARY OF THE INVENTION

[0002] This Summary is provided to introduce a selection of concepts in a simplified form that are further described below in the Detailed Description. This
20 Summary is not intended to identify key features or essential features of the claimed subject matter, nor is it intended to be used as an aid in determining the scope of the claimed subject matter.

[0003] Embodiments of the present invention generally relate to analyzing social media experiential data to influence decision-making calculations. Experiential social
25 media posts for a plurality of social media users are extracted. Event timelines are generated for each social media user. Based on a particular event of interest, a correlational data model is constructed for the social media users having experienced the particular event. The correlational data model is applied to a decision-making calculation corresponding to the particular event of interest.

30 BRIEF DESCRIPTION OF THE DRAWING

[0004] The present invention is illustrated by way of example and not limitation in the accompanying figures in which like reference numerals indicate similar elements and in which:

[0005] FIG. 1 is a block diagram of an exemplary computing environment suitable

for use in implementing embodiments of the present invention;

[0006] FIG. 2 is a block diagram of an exemplary system architecture in which embodiments of the invention may be employed;

[0007] FIG. 3 is a flow diagram showing a method for analyzing a large amount of social media data from a large population of social media users and constructing correlational data models between one or more events that occur within each user's timelines; and

[0008] FIG. 4 is a flow diagram showing a method for analyzing a large amount of social media data from a large population of social media users and constructing correlational data models between one or more events that occur within each user's timelines.

DETAILED DESCRIPTION OF THE INVENTION

[0009] The subject matter of the present invention is described with specificity herein to meet statutory requirements. However, the description itself is not intended to limit the scope of this patent. Rather, the inventors have contemplated that the claimed subject matter might also be embodied in other ways, to include different steps or combinations of steps similar to the ones described in this document, in conjunction with other present or future technologies. Moreover, although the terms "step" and/or "block" may be used herein to connote different elements of methods employed, the terms should not be interpreted as implying any particular order among or between various steps herein disclosed unless and except when the order of individual steps is explicitly described.

[0010] To assist individuals with their decision-making, some technologies have been developed that provide empirical data to curious individuals based on user opinions and reviews. Review sites, blogs, product and service webpages, etc., are typically referenced by individuals when seeking decision-making assistance based on others' experiences with a particular event or item. Generally speaking, these resources are limited to the specific event, service, or item itself. Social media, however, is a continuously increasing resource comprising an enormous amount of uncategorized, user-specific, experiential data for every day experiences, actions taken, and consequences resulting therefrom. Many social media users typically publish or "post" first-hand experiences nearly every day, sometimes multiple times each day. Typically, these social media users continue to post ongoing, seemingly related or unrelated events to a particular event or action previously or to eventually be taken.

[0011] Embodiments of the present invention are generally directed to analyzing a

large amount of social media data from a large population of social media users and constructing correlational data models between one or more events that occur within each user's timelines. By analyzing first-person events experienced by a single user and correlating events to other events that occur in that single person's life, as posted on social media, such data collected over a large sample of social media users can be used to construct decision-making models used for assisting a user potentially interested in making similar decisions or experiencing similar experiences. In other words, social media experiential data can be used to assist one to make better decisions and take better actions to reach their goals by harnessing the experiential data published by hundreds of millions of social media users posting about the actions they've taken and consequences experienced in their lives thereafter.

[0012] Accordingly, in one aspect, an embodiment of the present invention is directed to a non-transitory computer storage medium storing computer-usable instructions that, when used by one or more computing devices, cause the one or more computing devices to perform operations. The operations include extracting a plurality of experiential social media posts associated to a plurality of social media users. The operations further include generating a plurality of event timelines based on at least the metadata of the plurality of experiential social media posts, each event timeline being unique to each of the plurality of social media users. The operations also include constructing a correlational data model by measuring correlations between a particular event of interest and at least some of the plurality of experiential social media posts in at least some of the plurality of event timelines. The operations further include applying the correlational data model to at least one decision-making calculation corresponding to the particular event of interest.

[0013] In another embodiment of the invention, an aspect is directed to a computer-implemented method. The method includes extracting a plurality of experiential social media posts associated to a plurality of social media users. The method also includes generating a plurality of event timelines based on at least the metadata of the plurality of experiential social media posts, each event timeline being unique to each of the plurality of social media users. The method also includes constructing, using at least one processor, a correlational data model by measuring correlations between a particular event of interest and at least some of the plurality of experiential social media posts in at least some of the plurality of event timelines. The method further includes applying the correlational data model to at least one decision-making calculation corresponding to the particular event of interest.

[0014] A further embodiment is directed to a computerized system comprising one or more processors; and one or more computer storage media storing computer-useable instructions that, when used by the one or more processors, cause the one or more processors to: extract a plurality of experiential social media posts associated to a plurality of social media users; generate a plurality of event timelines based on at least the metadata and some content of the plurality of experiential social media posts, each event timeline being unique to each of the plurality of social media users, and wherein the content includes at least some temporal description; construct, using the one or more processors, a correlational data model by measuring correlations between a particular event of interest and at least some of the plurality of experiential social media posts in at least some of the plurality of event timelines, wherein the correlations measured includes identifying a positive and/or negative valence of the at least some of the plurality of experiential social media posts to the particular event of interest; and apply the correlational data model to at least one decision-making calculation corresponding to the particular event of interest.

[0015] Having briefly described an overview of embodiments of the present invention, an exemplary operating environment in which embodiments of the present invention may be implemented is described below in order to provide a general context for various aspects of the present invention. Referring to the figures in general and initially to FIG. 1 in particular, an exemplary operating environment for implementing embodiments of the present invention is shown and designated generally as computing device 100. The computing device 100 is but one example of a suitable computing environment and is not intended to suggest any limitation as to the scope of use or functionality of embodiments of the invention. Neither should the computing device 100 be interpreted as having any dependency or requirement relating to any one component nor any combination of components illustrated.

[0016] Embodiments of the invention may be described in the general context of computer code or machine-useable instructions, including computer-useable or computer-executable instructions such as program modules, being executed by a computer or other machine, such as a personal data assistant or other handheld device. Generally, program modules include routines, programs, objects, components, data structures, and the like, and/or refer to code that performs particular tasks or implements particular abstract data types. Embodiments of the invention may be practiced in a variety of system configurations, including hand-held devices, consumer electronics, general-purpose computers, data centers, more specialty computing devices, and the like. Embodiments of

the invention may also be practiced in distributed computing environments where tasks are performed by remote-processing devices that are linked through a communications network.

[0017] With continued reference to FIG. 1, the computing device 100 includes a bus 110 that directly or indirectly couples the following devices: a memory 112, one or more processors 114, one or more presentation components 116, one or more input/output (I/O) ports 118, one or more I/O components 120, and an illustrative power supply 122. The bus 110 represents what may be one or more busses (such as an address bus, data bus, or combination thereof). Although the various blocks of FIG. 1 are shown with lines for the sake of clarity, in reality, these blocks represent logical, not necessarily actual, components. For example, one may consider a presentation component such as a display device to be an I/O component. Also, processors have memory. The inventors hereof recognize that such is the nature of the art, and reiterate that the diagram of FIG. 1 is merely illustrative of an exemplary computing device that can be used in connection with one or more embodiments of the present invention. Distinction is not made between such categories as “workstation,” “server,” “laptop,” “hand-held device,” etc., as all are contemplated within the scope of FIG. 1 and reference to “computing device.”

[0018] The computing device 100 typically includes a variety of computer-readable media. Computer-readable media may be any available media that is accessible by the computing device 100 and includes both volatile and nonvolatile media, removable and non-removable media. Computer-readable media comprises computer storage media and communication media; computer storage media excluding signals per se. Computer storage media includes volatile and nonvolatile, removable and non-removable media implemented in any method or technology for storage of information such as computer-readable instructions, data structures, program modules or other data. Computer storage media includes, but is not limited to, RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM, digital versatile disks (DVD) or other optical disk storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other medium which can be used to store the desired information and which can be accessed by computing device 100.

[0019] Communication media, on the other hand, embodies computer-readable instructions, data structures, program modules or other data in a modulated data signal such as a carrier wave or other transport mechanism and includes any information delivery media. The term “modulated data signal” means a signal that has one or more of its

characteristics set or changed in such a manner as to encode information in the signal. By way of example, and not limitation, communication media includes wired media such as a wired network or direct-wired connection, and wireless media such as acoustic, RF, infrared and other wireless media. Combinations of any of the above should also be included within the scope of computer-readable media.

[0020] The memory 112 includes computer-storage media in the form of volatile and/or nonvolatile memory. The memory may be removable, non-removable, or a combination thereof. Exemplary hardware devices include solid-state memory, hard drives, optical-disc drives, and the like. The computing device 100 includes one or more processors that read data from various entities such as the memory 112 or the I/O components 120. The presentation component(s) 116 present data indications to a user or other device. Exemplary presentation components include a display device, speaker, printing component, vibrating component, and the like.

[0021] The I/O ports 118 allow the computing device 100 to be logically coupled to other devices including the I/O components 120, some of which may be built in. Illustrative I/O components include a microphone, joystick, game pad, satellite dish, scanner, printer, wireless device, a controller, such as a stylus, a keyboard and a mouse, a natural user interface (NUI), and the like.

[0022] A NUI processes air gestures (*i.e.*, motion or movements associated with a user's hand or hands or other parts of the user's body), voice, or other physiological inputs generated by a user. A NUI implements any combination of speech recognition, touch and stylus recognition, facial recognition, biometric recognition, gesture recognition both on screen and adjacent to the screen, air gestures, head and eye tracking, and touch recognition associated with displays on the computing device 100. The computing device 100 may be equipped with one or more touch digitizers and/or depth cameras, such as, stereoscopic camera systems, infrared camera systems, RGB camera systems, and combinations of these for direct and/or hover gesture detection and recognition. Additionally, the computing device 100 may be equipped with accelerometers or gyroscopes that enable detection of motion. The output of the accelerometers or gyroscopes is provided to the display of the computing device 100 to render immersive augmented reality or virtual reality.

[0023] Aspects of the subject matter described herein may be described in the general context of computer-executable instructions, such as program modules, being executed by a computing device. Generally, program modules include routines, programs,

objects, components, data structures, and so forth, which perform particular tasks or implement particular abstract data types. Aspects of the subject matter described herein may also be practiced in distributed computing environments where tasks are performed by remote processing devices that are linked through a communications network. In a distributed computing environment, program modules may be located in both local and remote computer storage media including memory storage devices.

[0024] Turning now to FIG. 2, a block diagram is provided illustrating an exemplary system 200 in which some embodiments of the present invention may be employed. It should be understood that this and other arrangements described herein are set forth only as examples. Other arrangements and elements (e.g., machines, interfaces, functions, orders, and groupings of functions, etc.) can be used in addition to or instead of those shown, and some elements may be omitted altogether. Further, many of the elements described herein are functional entities that may be implemented as discrete or distributed components or in conjunction with other components, and in any suitable combination and location. Various functions described herein as being performed by one or more entities may be carried out by hardware, firmware, and/or software. For instance, various functions may be carried out by a processor executing instructions stored in memory.

[0025] Among other components not shown, the system 200 may include a number of social networking services 202A, 202B, 202N, a social data aggregator 204, and an event relationship module 206. It should be understood that the system 200 shown in FIG. 2 is an example of one suitable computing system architecture. Each of the components shown in FIG. 2 may be implemented via any type of computing device, such as computing device 200 described with reference to FIG. 2, for example. The components may communicate with each other via a network 208, which may include, without limitation, one or more local area networks (LANs) and/or wide area networks (WANs). Such networking environments are commonplace in offices, enterprise-wide computer networks, intranets, and the Internet. It should be understood that any number of social networking services, social data aggregators, and social analysis tools may be employed within the system 200 within the scope of the present invention. Each may comprise a single device or multiple devices cooperating in a distributed environment. For instance, the event relationship module 206 may be provided via multiple devices arranged in a distributed environment that collectively provide the functionality described herein. Additionally, other components not shown may also be included within the

network environment.

[0026] The event relationship module 206 may be employed by search engine technologies in conjunction with, by way of example only, personal assistant/advisor applications. Among other things, the event relationship module 206 operates to collect
5 social data from social networking services 202A, 202B, 202N. As represented in FIG. 2, social data may be collected from any number of social networking services. These services generally include any online presence at which users may share messages with other users within a social network of users. In some instances, the event relationship module 206 may access social data directly from a social networking service or an entity
10 providing the event relationship module 206 may access the data from a social networking service and provide the data to the module 206. For instance, a social networking service may provide APIs that expose the data. In other instances, the event relationship module 206 may access social data from a third-party social data aggregator 204, which may operate to access data from one or more social networking services, standardize the data,
15 and provide the standardized data. Any and all such variations are contemplated to be within the scope of embodiments of the present invention.

[0027] As shown in FIG. 2, the event relationship module 206 may include, among other things not shown, a timeline extraction component 210 and an analysis component 212. The timeline extraction component may include an experiential message
20 identification component 214, a timestamp identification component 216, an event extraction component 218, and a valence identification component 220. The analysis component 212 may also include a causality analysis component 222. In some cases, an input query 224, corresponding to an action, event, or item of particular interest, may be communicated to the analysis component 212 for processing, as will be described.

[0028] The social data communicated to the event relationship module 206 may include a corpus of social media messages or “posts” (i.e., messages comprising original
25 microblog texts posted by individuals). Such messages are comprised of original text communications posted by individuals. In some instances, the messages may comprise unstructured text data, typically seen with original text communications as described. In
30 some other instances, the messages may comprise structured data. For example, a fitness tracker may automatically post, for a user, a non-textual structured representation of an exercise and data associated therewith. In another example, a music streaming service may post, for the user, a structured representation of a song that the user is currently listening to. Such structured representations may avoid simple textual descriptions of an activity,

and instead use some predetermined representation that describes the key facets of a user action without specifying how these facets should be combined for presentation to other users. In some embodiments, messages may generally include a combination of simple textual representations in addition to structured representations. In some other
5 embodiments, these messages may also include at least a user identifier and a timestamp, but may also include other metadata (e.g., location information, biographical information, popularity statistics, social network connections, etc.).

[0029] The event relationship module 206 performs an analysis on the corpus of messages using the timeline extraction component 210. The timeline extraction component
10 210 analyzes the corpus of social media messages to generate a set of timelines of event occurrences. Each timeline in the set comprises a timeline of events occurring in a single individual's life (i.e., a social media user). In some instances, the events may include actions explicitly taken by the individual. In other instances, the events may describe outcomes that occurred as a consequence of a particular action, or even background events
15 that occur as a result of unrelated causes. The events, as described, may be directly extracted from each individual social media message, or inferred from the corpus as a whole.

[0030] In more detail, the timeline extraction component 210 may include subcomponents for facilitating event timeline extraction. The experiential message
20 identification component 214, for example, may be utilized for identifying experiential messages that describe personal experiences of the author (i.e., the social media user), whether written by the author in past, current or (expected) future form. Social messages may include an amount of social "noise," including conversational texts, hearsay, pointers to news articles and current events, among others. The experiential message identification
25 component 214 is configured to distinguish experiential messages from other social media content using a machine learning technique implementing any learning technology including, by way of example only, linear regressions, support vector machines, and/or deep neural networks.

[0031] The timeline extraction component 210 may further include a timestamp
30 identification component 216 for identifying a time period referred to in an experiential message, or a time period inferred from a plurality of experiential messages. In essence, while many social media messages provide *in situ* reports of an individual's experiences, it is not uncommon for authors to also report on past and/or anticipated experiences. As such, the timestamp identification component 216 analyzes experiential messages using a

rule-based system that can recognize and resolve basic expressions of relative offsets (i.e., “yesterday” or “next weekend”), as well as references to nearby days and dates (i.e., “Tuesday” and “February 10th”). By way of example only, an experiential message that recites “I can’t believe my marathon is coming up in two days,” can be interpreted by the timestamp identification component 216 as an event described in the experiential message (e.g., the “marathon”) to occur two days after the timestamp stored within metadata of the message.

[0032] The timeline extraction component 210 may further include an event extraction component 218 for the recognition and canonicalization of events (i.e., identifying both actions and outcomes). In more detail, to facilitate the generation of a timeline of “events” from the timeline of messages, the timeline extraction component 210 identifies and extracts events from the social media message text. These events may be extracted directly from the textual representation of a message, or inferred from a plurality of messages. The task of extracting such events may be analogous to the task of named entity recognition, while sharing many of its challenges (i.e., including candidate identification, disambiguation, and canonicalization) as one skilled in the art may appreciate. In one embodiment, all phrases of the message are extracted as potential events, without attempting to classify them as actions, outcomes, or neither. Statistical phrase segmentation is used to identify candidates, while distribution of context is utilized as a comparison to identify canonicalization.

[0033] In some embodiments, a statistical modeling approach is utilized by the timeline extraction component 210 to infer hidden phrase boundaries in a text for phrase segmentation. To efficiently locate phrases, a phrase unigram language model is utilized. Briefly, each token in a phrase unigram language model includes one or more white-space separated words. By encoding multiple words within a single unigram, the phrase language model is able to capture long distance relationships without requiring high Markov order statistics and concomitant large models. The phrase unigram language model itself is trained from a large corpus of text using an expectation-maximization (EM) process that iteratively segments a corpus into likely phrases and then retrains a new phrase unigram language model. As one skilled in the art may appreciate, provided a phrase unigram language model, identifying phrase segmentations in a message is a matter of searching for the most probable combination of component unigrams. To facilitate canonicalization, some embodiments may build a distribution of single-word token probabilities based on the co-occurrences of single-word tokens that co-occur with each of

the segmented phrases. Agglomerative hierarchical clustering may be utilized to group together all tokens that are within a distance threshold of one another, where the distances between two tokens is measured as the cosine similarity between the two feature vectors.

[0034] The timeline extraction component 210 may further include a valence identification component 220 for identifying whether an event is associated to a positive or negative valence. In other words, the valence identification component 220 may be operable to detect moods or sentiments associated with words in social media messages. Detection of these moods or sentiments and associating them with outcomes may help with reasoning about their significance. In some embodiments, a domain-agnostic affect extractor may be utilized to extract the author's levels of associated emotions to an event (e.g., joviality, sadness, fatigue, hostility, etc.).

[0035] The analysis component 212 receives, as briefly described above, an input query 224 for processing. The input query 224 may be defined by two events, a first event ($E+$) and a second event ($E-$). The causality analysis component 222 may be configured to identify the precedent and subsequent relationships of an event $E+$ that distinguish the social media timelines containing $E+$ from timelines comparing some event $E-$. Semantically, $E+$ and $E-$ may be thought of as identifying either positive and negative outcomes, or control and treatment classes. The query 224, defined as two classes of events, $E+$ and $E-$, is used to identify and align two sets of timelines received from timeline extraction component 210. A class of events E is specified as, for example, some specific observation, or a complex matching function. As will be further described herein, different forms of high-level questions may be analyzed via analysis component 212 depending on the specific query chosen. For example, if a query is chosen such that $E+$ selects the performance of some specific action (and $E-$ selects the non-performance of that action), then the results generated by analysis component 212 will identify what is likely to happen after taking the specified action. On the other hand, if a query is chosen such that $E+$ selects the achievement of a specific goal (and $E-$ selects the non-achievement of that goal), then the precedents identified by analysis component 212 will identify what is likely done and differentiates between people achieving the goal and not achieving the goal. In some embodiments, regardless of the specification of $E+$ or $E-$, in either the context of a goal or action, both precedent and subsequent events may be calculated. As such, precedent events leading to taking a specific action may be analyzed, as well as subsequent events following from having achieved a specific goal.

[0036] The causality analysis component 222 may identify the distinguishing

precedent and subsequent events when comparing timelines containing an event E^+ to those containing an event E^- . In one embodiment, the events may be identified using a simple correlational analysis while, in another embodiment, the events may be identified using a correlational analysis with semantic restrictions. In a first example, simple correlations are determined between a target event and the events that precede or follow it. The goal is to find events that are more correlated with occurring before or after E^+ (but not both before and after) than occurring before or after E^- . In another example, semantic correlations add restrictions that only consider those events that are believed to be semantically closely related to the domain of interest. Semantic correlation makes an assumption that if an event of interest is related to target events E^+ and E^- , then at least one person would have clearly mentioned the event of interest in the recognizable context of the target domain. As such, the ranked events of interest should be more robust to noise and confounds and, although the query model is extended to include a specification of a domain of interest, it is expected that events found to be correlated are more likely to be easily interpretable by humans.

[0037] Turning now to FIG. 3, a flow diagram is provided that illustrates a method 300 for analyzing a large amount of social media data from a large population of social media users and constructing correlational data models between one or more events that occur within each user's timeline. Each block of the method 300 and other methods described herein comprises a computing process that may be performed using any combination of hardware, firmware, and/or software. For instance, various functions may be carried out by a processor executing instructions stored in memory. The methods may also be embodied as computer-usable instructions stored on computer storage media. The methods may be provided by a standalone application, a service or hosted service (standalone or in combination with another hosted service), or a plug-in to another product, to name a few.

[0038] As shown at block 302, a plurality of experiential social media posts associated to a plurality of social media users are extracted (for instance, by the timeline extraction component 210 of FIG. 2). The experiential social media posts are identified and extracted, as described herein, from a plurality of uncategorized social media posts received, for instance, by social data aggregator 204 of FIG. 2. Each social media post, regardless of type or categorization, may include metadata that includes a variety of information about the post and/or authoring user, such as authoring user's name or username, geographic location, and timestamp of the post submission, among other things.

[0039] An experiential social media post is defined as a social media post that includes content that communicates an authoring user's first person experience. In order to identify experiential social media posts among a plurality of uncategorized social media posts, methods using various machine learning techniques may be utilized. In some
5 embodiments, simple statistical language modeling may be used to identify experiential social media posts from uncategorized social media posts. Statistical language modeling may include a broad spectrum of machine learning techniques ranging from simple learning algorithms, such as linear regressions, to more complex processes, such as support vector machines (SVMs) and deep neural networks. The foregoing are merely
10 examples of machine learning techniques and are not intended to be limiting in scope, as any machine learning technique for implementing statistical language modeling may be considered within the scope of the present invention.

[0040] A plurality of event timelines are generated (for instance, by timeline extraction component 210 of FIG. 2), as shown at block 304. An event timeline is
15 generated for each unique user, such that each of the plurality of social media users is associated with a unique event timeline. Event timelines may include a timeline of events for each unique individual, wherein the events are described using experiential social media posts by the individual user. The event timelines may be generated based on, among
20 other things, the metadata of each experiential social media post. In some embodiments, when an experiential social media post includes a temporal description of an event, the experiential social media post placement on the event timeline may be adjusted to align with the described temporal description (for instance, using timestamp identification component 216 of FIG. 2). Event timelines may also be generated based on a combination
25 of social media data and other non-personal data sources. For example, if someone says, in a post, that they ate a hamburger and fries, some embodiments may seek and retrieve nutritional information from the government and/or nutritional resource and create an "ate an unhealthy meal" event and/or specifically add nutritional information into the event itself. In another example, if someone posts that they were at a particular location on a particular day, some embodiments may add contextual information about the location. For
30 example, if someone posts that they "jogged in Seattle today," embodiments may seek and retrieve weather information for the location and day to add context to the post that the author "jogged in the rain."

[0041] In some embodiments, the experiential social media posts may each be analyzed (for instance, by valence identification component 220) to identify a positive or

negative valence of each experiential social media post. In some embodiments, valence may be identified by analyzing multiple social media posts in aggregation. For example, if someone posts “today was a great day,” some embodiments may assign a positive valence to all events that occurred in that particular day in that person’s timeline. Positive
5 or negative valence may also be multi-dimensional. For example, someone may express that they are, to some degree, both happy and tired. This multi-faceted representation may yield more insights and more appropriately represent the complexity of an event than a single positive/negative value.

[0042] A correlational data model is constructed (for instance, by analysis
10 component 212 of FIG. 2), as shown at block 306, by measuring correlations between a particular event of interest received (for example, through input query 224 of FIG. 2) and at least some of the plurality of experiential social media posts in at least some of the plurality of event timelines. For instance, events defined in an input query may also occur or be relevant to one or more posts found in none, some, or all of the plurality of event
15 timelines. As such, the analysis component 212 may analyze the event timelines to determine whether the timelines include events that are relevant to the input query. The correlational data model may include models of the correlational analysis or causal analysis variety. Correlational analysis, such as Bayesian analysis, or more sophisticated techniques motivated to estimate causality, such as propensity score matching, structural
20 equation modeling, granger causality, and path analysis, among others, may be considered for use as a correlational data model within the scope of the present invention. As such, the event timelines are analyzed to generate or construct a correlational data model that includes correlational measurements of any or all events that occur within the plurality of event timelines. In some embodiments, the correlational data model including
25 correlational measurements can be generated on-the-fly upon receiving input queries, pre-generated and cached using a predetermined set of input queries, or generated and cached over time, such that each input query generates correlational measurements on-the-fly, which are cached or stored in memory immediately thereafter.

[0043] In some embodiments, the correlational data model can be made available
30 as a service to the system or a 3rd party system (not shown) for querying relationships between one or more particular events of interest. In further embodiments, a decision-making calculation may request, from the correlational data model, data including and or all correlations between the event timelines and a particular event of interest having, for example, a magnitude above a particular threshold. As such, the correlational data model

can be applied to the decision-making calculation corresponding to the particular event of interest, to ultimately provide correlational data for decision-making assistance. One of ordinary skill in the art may appreciate that correlational data provided between a particular event and a large sample of event data can deliver empirical support for decision-making in various applications. In some embodiments, the magnitude of correlational measurements may indicate strength and/or potential importance of a particular relationship. Causal relationships, in particular, may be inferred when all possible causal factors have been observed and one or more combinations of causal factors are observed to lead deterministically to some outcome. In some instances, when one or more combinations of causal factors probabilistically lead to a particular outcome, causal relationships may be established. As such, algorithms that perform causal analysis attempt to factor out all alternative explanations for the occurrence of a particular event. In essence, algorithms may determine whether a particular combination of factors will lead to a specific outcome, and whether the lack of that combination of factors otherwise lead to failure of the specific outcome. In some embodiments, algorithms for determining causal relationships attempt to determine, when all other events are substantially equal (thereby factoring out alternative explanations for the occurrence of a particular event), whether a causal relationship exists. In such instances, the event relationship module 206 may be operable to determine a causal relationship between the particular event of interest and various events in the event timelines, further utilizing such determinations for more active applications, as will be described.

[0044] Application of the correlational data model may be conducted by a system in an active or passive manner. When a particular event of interest is submitted or noted to be of importance in a particular application, an active application of the model can be designed to present alerts upon, for example, detecting timeline events having a threshold correlation being reached, or a causal relationship being established. In some embodiments, an active application of the data model may include a computing device or software application that autonomously takes a predefined action without the user's intervention. For instance, a software application may be operable to interrupt the user with an alert, automatically purchase tickets to and/or schedule events, or even actively prevent presentation of some alerts. In other instances, a passive application of the model can be designed to provide correlational data only when queried by a particular application or user. In embodiments, the correlational data model may further provide additional information with the correlational data used for facilitating interpretation of the

correlational data. In such embodiments, the additional information may include content from the experiential social media posts having reached the threshold correlation.

[0045] Turning now to FIG. 4, a flow diagram is provided that illustrates a method 400 for analyzing a large amount of social media data from a large population of social media users and constructing correlational data models between one or more events that occur within each user's timeline. Initially, as shown at block 402, a plurality of social media posts are analyzed (for instance, using timeline extraction component 210 of FIG. 2) to identify experiential social media posts. The experiential social media posts are identified and extracted, as shown at block 404, as communicating a first-person experience of the author, as opposed to a non-experiential post. Identification of the experiential social media posts can be conducted by, for example, experiential message identification component 214 of FIG. 2. As shown at block 406, all experiential events are identified and extracted (for instance, using event extraction component 218 of FIG. 2) for each social media user having authored posts in the plurality of social media posts. Further, a positive or negative valence or mood of the events are determined (for instance, using valence identification component 220 of FIG. 2) for each event in the timelines, as shown at block 408. As shown at block 410, the timestamp of events are determined by analyzing, among other things, the metadata of each social media post from which the events were extracted and/or the content of the social media post from which the events were extracted (i.e., temporal descriptions).

[0046] Although the method 400 illustrates event extraction 406, valence identification 408, and determination of appropriate timestamps 410 in a sequential order, the foregoing events are not limited to a particular order and any order of the foregoing are considered within the scope of the present invention. Following the foregoing processes, an event timeline for each social media user responsible for authoring communications corresponding to the events is generated. As such, a plurality of event timelines is generated, each timeline being unique to a particular social media user.

[0047] After the event timelines are generated, they are subsequently analyzed according to a received input query, as shown at block 414, 416, and 418. More specifically, an input query is received (for instance, by analysis component 212 of FIG. 2) and analyzed in conjunction with the event timelines generated (for instance, by timeline extraction component 210 of FIG. 2). The analysis conducts a plurality of correlation measurements in the event timelines based on the input query, as shown at block 416. For example, the input query may include a particular event of interest, such as a specific

action or goal. The analysis broadly identifies precedents and subsequent relationships to the particular event of interest that distinguishes the event timelines containing the particular event of interest from timelines comparing, for example, an inverse action, a null action of the particular event, or the non-achievement of a goal. As such, a set of
5 timelines matching the particular event of interest may be extracted, in addition to a set of timelines that match the inverse, null, or non-achievement events, as described. The foregoing sets may provide a framework of generating a correlational data model for more flexibly querying the data. The correlational data model may be constructed (for instance, using analysis component 212 of FIG. 2). The data model is now operable to receive
10 specific queries such as, by way of example only: (1) inquiring about a specific action to identify what is likely to happen after taking the specified action; or (2) inquiring about planning for a particular goal achievement to identify what is likely done and differentiates between people actually achieve the goal and not achieving it.

[0048] The correlational data model may be utilized by any component of the
15 system 200 or made available as a service to a 3rd party, as shown at block 418. The application of the correlational model may be considered in various use cases. By way of example only, electronic personal assistants, search engine technologies, social researchers, market researchers, product development, any application for which an inquiry is made regarding action/consequence relationships, and others, are considered
20 within the scope of the present invention.

[0049] The present invention has been described in relation to particular embodiments, which are intended in all respects to be illustrative rather than restrictive. Alternative embodiments will become apparent to those of ordinary skill in the art to which the present invention pertains without departing from its scope.

[0050] While the invention is susceptible to various modifications and alternative constructions, certain illustrated embodiments thereof are shown in the drawings and have been described above in detail. It should be understood, however, that there is no intention to limit the invention to the specific forms disclosed, but on the contrary, the intention is to cover all modifications, alternative constructions, and equivalents falling within the spirit
25 and scope of the invention.
30

CLAIMS

1. A computer-implemented method for generating correlational data models from experiential social media posts to facilitate computer-assisted decision-making calculations, the method comprising:

extracting a plurality of experiential social media posts associated to a plurality of social media users;

generating a plurality of event timelines based on at least the metadata of the plurality of experiential social media posts, each event timeline being unique to each of the plurality of social media users;

constructing, using at least one processor, a correlational data model by measuring correlations between a particular event of interest and at least some of the plurality of experiential social media posts in at least some of the plurality of event timelines;

employing the correlational data model to conduct, via an analysis component, at least one decision-making calculation corresponding to the particular event of interest.

2. The computer-implemented method of claim 1, wherein an experiential social media post is a social media post comprising communications about a first-person experience.

3. The computer-implemented method of any of claims 1 and 2, wherein generating the plurality of event timelines is further based on the content of the plurality of experiential social media posts, the content including at least some temporal description.

4. The computer-implemented method of any of claims 1-3, wherein generating the plurality of event timelines further includes identifying a positive or negative valence of the at least some of the plurality of experiential social media posts.

5. The computer-implemented method of any of claims 1-4, wherein the metadata includes at least timestamps of the experiential social media posts.

6. The computer-implemented method of any of claims 1-5, further comprising:

applying the correlational data model to at least one decision-making calculation corresponding to the particular event of interest, wherein application of the correlational data model is conducted in one of an active or passive manner.

7. The computer-implemented method of any of claims 1-6, wherein the active manner includes producing active alerts to indicate that the correlational data model applied to the at least one decision-making calculation corresponding to the particular event of interest indicates a relevant relationship, and wherein the passive manner includes presenting data representative of the causal relationship when the correlational data model is applied to the at least one decision-making calculation corresponding to the particular event of interest.

8. The computer-implemented method of any of claims 1-7, wherein extracting the plurality of experiential social media posts is performed by identifying the plurality of experiential social media posts from a plurality of uncategorized social media posts using machine learning techniques.

9. The computer-implemented method of any of claims 1-8, further comprising:

determining whether the correlational data model applied to the at least one decision-making calculation corresponding to the particular event of interest is one of a causal relationship or a merely correlative relationship.

10. The computer-implemented method of claim 9, wherein when the relationship is determined to be merely correlative, presenting at least some content from the at least some of the plurality of experiential social media posts to provide some context in the at least one decision-making calculation.

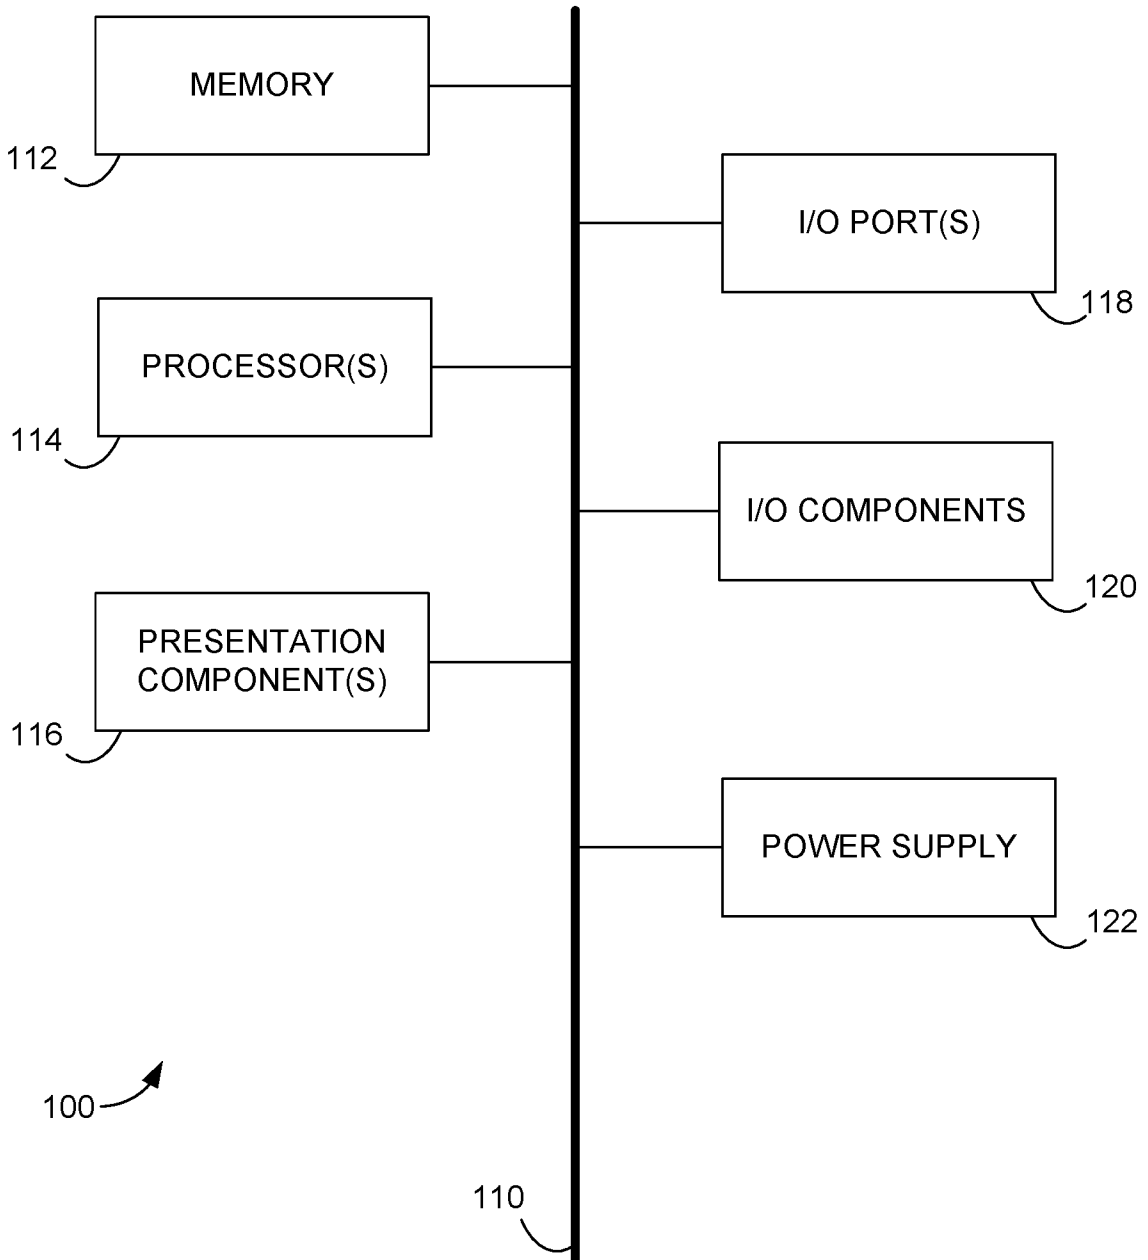


FIG. 1




FIG. 2

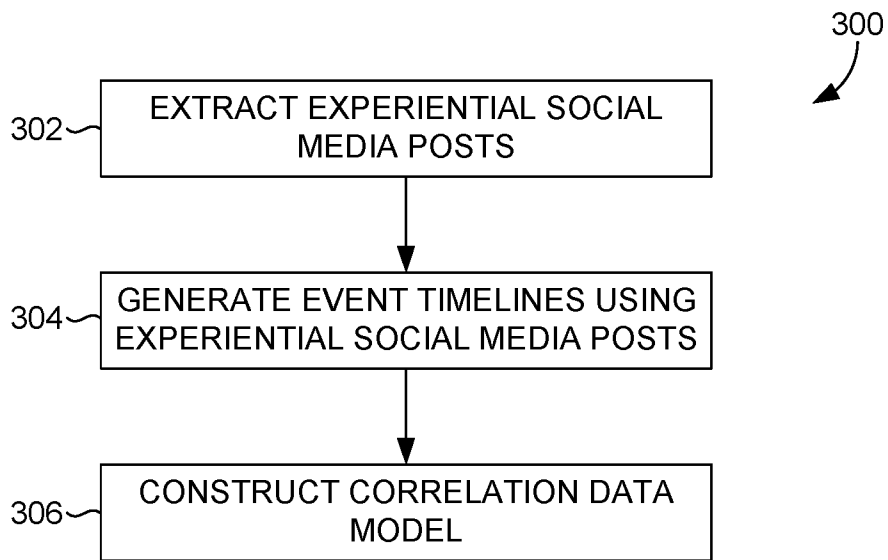


FIG. 3

4/4

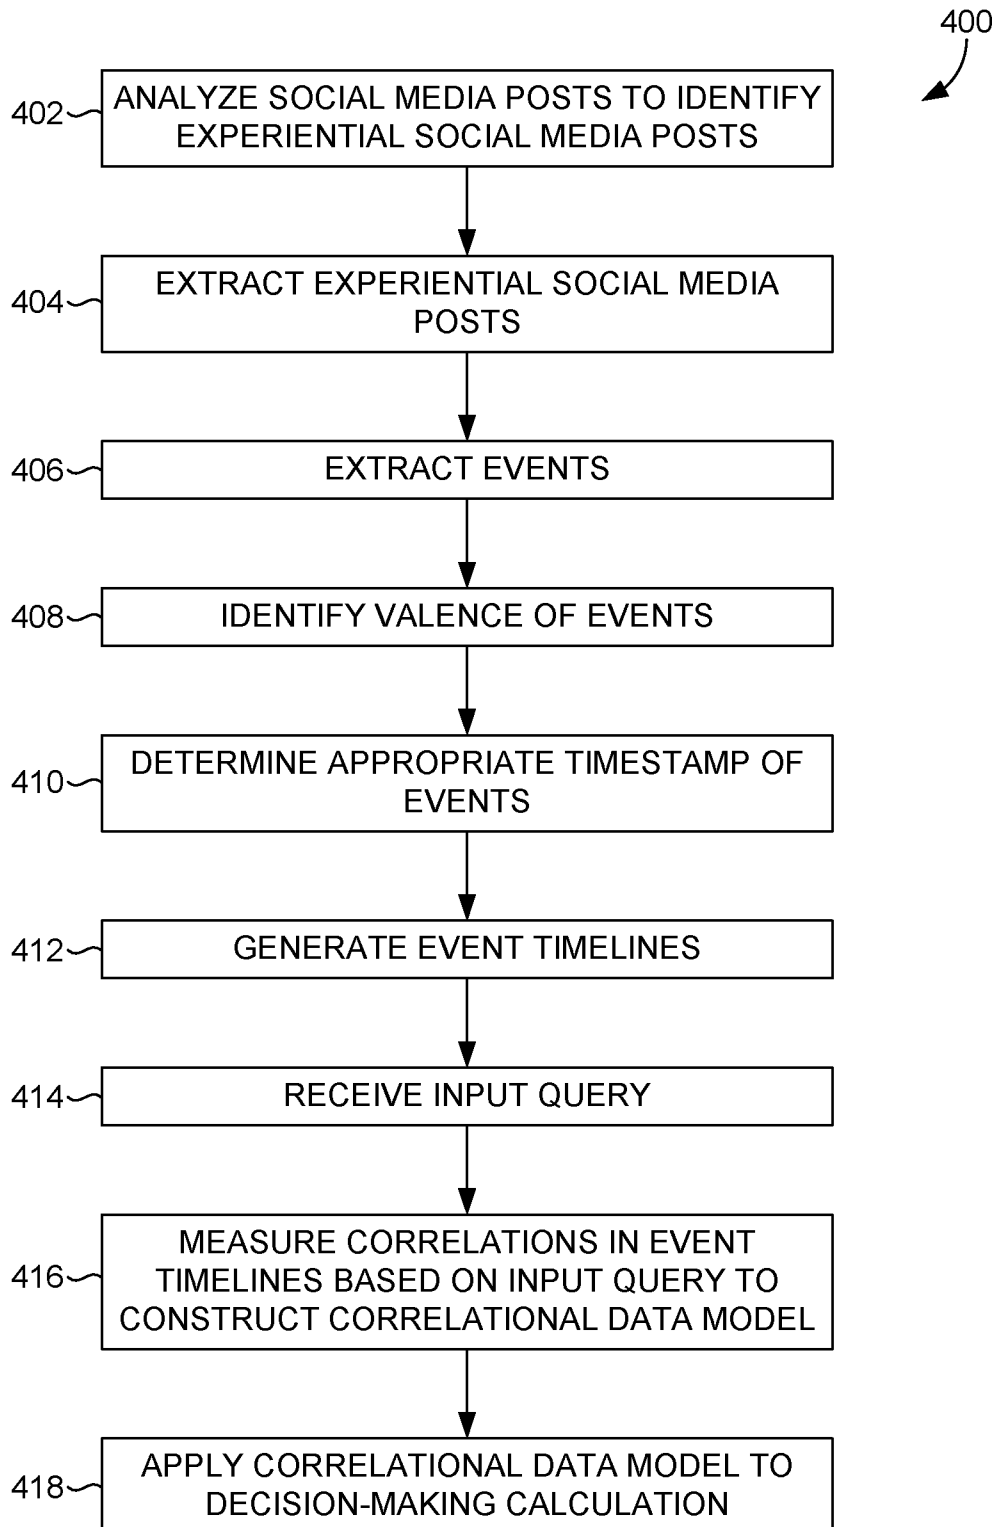


FIG. 4

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2015/064354

A. CLASSIFICATION OF SUBJECT MATTER
 INV. G06Q30/02 G06Q50/00 G06F17/30
 ADD.
 According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED
 Minimum documentation searched (classification system followed by classification symbols)
 G06Q G06F

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

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

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2012/089616 A1 (SA ADON ROEE ROBERT [IL]) 12 April 2012 (2012-04-12) the whole document	1-10
X	US 2013/124437 A1 (PENNACCHIOTTI MARCO [US] ET AL) 16 May 2013 (2013-05-16) the whole document	1-10
X	WO 2010/048172 A1 (LUMER ERIK [IT]; CASCAAD SRL [IT]) 29 April 2010 (2010-04-29) the whole document	1-10
X	US 2010/268830 A1 (MCKEE MARTIN W [US] ET AL) 21 October 2010 (2010-10-21) the whole document	1-10

Further documents are listed in the continuation of Box C.

See patent family annex.

* Special categories of cited documents :

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier application or patent but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

- "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
- "&" document member of the same patent family

Date of the actual completion of the international search 16 February 2016	Date of mailing of the international search report 25/02/2016
--	---

Name and mailing address of the ISA/ European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016	Authorized officer Chauvet, Christophe
--	--

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2015/064354

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2012089616 A1	12-04-2012	US 2012089616 A1	12-04-2012
		US 2014108409 A1	17-04-2014
		US 2014108429 A1	17-04-2014
		US 2014108430 A1	17-04-2014

US 2013124437 A1	16-05-2013	NONE	

WO 2010048172 A1	29-04-2010	US 2012001919 A1	05-01-2012
		WO 2010048172 A1	29-04-2010

US 2010268830 A1	21-10-2010	NONE	
