



(51) International Patent Classification:

G06Q 30/02 (2012.01) G06F 17/30 (2006.01)
G06Q 30/06 (2012.01) G06F 17/28 (2006.01)

(21) International Application Number:

PCT/US2014/048089

(22) International Filing Date:

24 July 2014 (24.07.2014)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

13/952,470 26 July 2013 (26.07.2013) US
13/952,459 26 July 2013 (26.07.2013) US

(71) Applicant: **GENESYS TELECOMMUNICATIONS LABORATORIES, INC.** [US/US]; 2001 Junipero Serra Boulevard, Daly City, California 94014 (US).

(72) Inventors: **LEV-TOV, Amir**; c/o GENESYS TELECOMMUNICATIONS LABORATORIES, INC., 2001 Junipero Serra Boulevard, Daly City, California 94014 (US). **FA-IZAKOF, Avraham**; c/o GENESYS TELECOMMUNICATIONS LABORATORIES, INC., 2001 Junipero Serra Boulevard, Daly City, California 94014 (US). **OLLINGER, David**; c/o GENESYS TELECOMMUNICATIONS LABORATORIES, INC., 2001 Junipero Serra Boulevard, Daly City, California 94014 (US). **KONIG,**

Yochai; c/o GENESYS TELECOMMUNICATIONS LABORATORIES, INC., 2001 Junipero Serra Boulevard, Daly City, California 94014 (US).

(74) Agent: **CHANG, Josephine E.**; Christie, Parker & Hale, LLP, P.O. Box 29001, Glendale, California 91209-9001 (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, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, 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,

[Continued on next page]

(54) Title: SYSTEM AND METHOD FOR DISCOVERING AND EXPLORING CONCEPTS




FIG. 1

(57) Abstract: A method for identifying concepts in a plurality of interactions includes: filtering, on a processor, the interactions based on intervals; creating, on the processor, a plurality of sentences from the filtered interactions; computing, on the processor, a saliency of each the sentences; pruning away, on the processor, sentences with low saliency for generating a set of informative sentences; clustering, on the processor, the sentences of the set of informative sentences for generating a plurality of sentence clusters, each of the clusters corresponding to a concept of the concepts; computing, on the processor, a saliency of each of the clusters; and naming, on the processor, each of the clusters.



TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, —
KM, ML, MR, NE, SN, TD, TG).

*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))*

Published:

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

1 **SYSTEM AND METHOD FOR DISCOVERING AND EXPLORING CONCEPTS**

FIELD

5 **[0001]** Aspects of the present invention relate to speech processing, indexing, searching, and analytics. In particular, aspects of the present invention relate to analyzing recorded and live information to categorize conversations and to identify concepts and trends.

BACKGROUND

10 **[0002]** An organization's contact center typically receives a multitude of calls regarding a variety of issues. For example, a sales department of a contact center may receive calls with questions about the feature sets and pricing of various products offered by the organization; a customer support department may receive calls regarding particular problems with using the products or the quality of the services being delivered; an accounts department may receive calls about changes in billing policy, incorrect charges, and other issues.

15 **[0003]** Generally, it is useful for an organization to be able to identify concepts and patterns within the conversations (or "interactions") in order to categorize the calls and identify underlying issues to be addressed (e.g., specific complaints about products or general dissatisfaction with services). However, conventional systems for doing so generally involve the manual survey of data collected by customer support agents and manual analysis of this data. This manual process of analysis can be time consuming and there may be long delays between collecting the data and determining results from the analysis.

20 **[0004]** In some conventional systems, conversations can be tagged or categorized based on their containing predefined keywords or phrases. For example, through the above discussed manual (human) analysis of phrases that are either identified by a human listener or identified by a computer system using phrase recognition, one might infer that conversations with a call center that contain the phrases "I would like to speak to your manager" and "Can I talk to your supervisor?" lead to the escalation of the call to a higher level representative. As such, any call containing these phrases would be categorized as containing an "escalation attempt."

25 **[0005]** As such, an organization can identify trends and infer conditions based on the number of such interactions falling into various categories. For example, a large number of interactions originating from a particular area and categorized as indicating a "service outage" or "poor network performance" could alert an internet service provider to take action to address system problems within that particular area.

30 **[0006]** However, conversations containing phrases that were not previously identified would not be categorized appropriately. For example, if the phrase "Let me talk to your boss" was not previously identified as being associated with escalation attempts, then a conversation containing that phrase would not be categorized as an "escalation attempt."

1 [0007] In addition, some conventional systems use Bayesian networks to identify correlations between events. However, developing these Bayesian networks requires human input to specify various parameters (e.g., the nodes of the Bayesian network).

5 SUMMARY

[0008] Aspects of embodiments of the present invention are directed to addressing two issues related to discovery and exploration of data. One aspect of the present invention is directed to the automatic discovery and extraction of concepts from a set of documents without human assistance. In addition, embodiments of the present invention are directed to understanding why a certain document or phone call is classified into a predefined category or why the document or phone call contains a particular concept (which may be newly identified). In other words, embodiments of the present invention are directed to discovering what is the “root cause” for various “symptoms” reported by customers and inferring a predefined category or newly discovered concept “A” that is the root cause for the call to be categorized with a predefined category or newly discovered concept “B”, so an organization or a user can resolve the underlying issue.

[0009] According to one embodiment of the present invention, a method for identifying concepts in a plurality of interactions includes: filtering, on a processor, the interactions based on intervals; creating, on the processor, a plurality of sentences from the filtered interactions; computing, on the processor, a saliency of each the sentences; pruning away, on the processor, sentences with low saliency for generating a set of informative sentences; clustering, on the processor, the sentences of the set of informative sentences for generating a plurality of sentence clusters, each of the clusters corresponding to a concept of the concepts; computing, on the processor, a saliency of each of the clusters; and naming, on the processor, each of the clusters.

[0010] The interactions may include an output of a voice recognition system.

[0011] The method may further include filtering the output of the voice recognition system based on word confidence.

[0012] The voice recognition system may be a large-vocabulary continuous speech recognition system.

[0013] The intervals may be time intervals.

[0014] The clustering the sentences may include: selecting a plurality of template sentences from the set of informative sentences, each of the template sentences corresponding to one of the clusters; computing a similarity of each sentence in the set of informative sentences to the template sentences; assigning each of the sentences of the set of informative sentences to a cluster of the clusters in accordance with the computed similarities with the template sentences; and removing clusters having fewer than a threshold number of assigned sentences.

1 [0015] The clustering the sentences may further include, iteratively, selecting additional
template sentences from sentences that were not assigned to a cluster and repeating the
selecting the plurality of template sentences, computing the similarity of each sentence to the
template sentences; the assigning the each of the sentences to the clusters; and removing
5 clusters having fewer than the threshold number of assigned sentences until all sentences are
assigned or until an iteration limit is reached.

[0016] The naming each of the clusters may include, for each cluster: computing a term
frequency-inverse document frequency for each sentence in the cluster and naming the
cluster with a sentence of the cluster having the highest term frequency-inverse document
10 frequency.

[0017] According to another embodiment of the present invention, a method for
exploring concepts automatically identified in a plurality of interactions includes: receiving,
on a processor, a query comprising a concept; retrieving, by the processor, a cluster
containing the concept; and displaying the retrieved cluster, wherein the concepts are
15 automatically identified by: filtering, on the processor, the interactions based on intervals;
creating, on the processor, a plurality of sentences from the filtered interactions; computing,
on the processor, a saliency of each the sentences; pruning away, on the processor, sentences
with low saliency for generating a set of informative sentences; and clustering, on the
processor, the sentences of the set of informative sentences for generating a plurality of
20 sentence clusters, each of the clusters corresponding to a concept of the concepts.

[0018] According to one embodiment of the present invention, a computer system
includes: a processor; and a memory, wherein the memory has stored thereon instructions
that, when executed by the processor, causes the processor to identify concepts in a plurality
of interactions by: filtering the interactions based on intervals; creating a plurality of
25 sentences from the filtered interactions; computing a saliency of each the sentences; pruning
away sentences with low saliency for generating a set of informative sentences; clustering the
sentences of the set of informative sentences for generating a plurality of sentence clusters,
each of the clusters corresponding to a concept of the concepts; computing a saliency of each
of the clusters; and naming each of the clusters.

30 [0019] The interactions may include an output of a voice recognition system.

[0020] The instructions may further include instructions for filtering the output of the
voice recognition system based on word confidence.

[0021] The voice recognition system may be a large-vocabulary continuous speech
recognition system.

35 [0022] The intervals may be time intervals.

[0023] The clustering the sentences may include: selecting a plurality of template
sentences from the set of informative sentences, each of the template sentences corresponding
to one of the clusters; computing a similarity of each sentence in the set of informative

1 sentences to the template sentences; assigning each of the sentences of the set of informative sentences to a cluster of the clusters in accordance with the computed similarities with the template sentences; and removing clusters having fewer than a threshold number of assigned sentences.

5 [0024] The clustering the sentences may further include, iteratively, selecting additional template sentences from sentences that were not assigned to a cluster and repeating the selecting the plurality of template sentences, computing the similarity of each sentence to the template sentences; the assigning the each of the sentences to the clusters; and removing clusters having fewer than the threshold number of assigned sentences until all sentences are
10 assigned or until an iteration limit is reached.

[0025] The naming each of the clusters may include, for each cluster: computing a term frequency-inverse document frequency for each sentence in the cluster and naming the cluster with a sentence of the cluster having the highest term frequency-inverse document frequency.

[0026] According to one embodiment of the present invention, a computer system
15 includes: a processor; and a memory, wherein the memory has stored thereon instructions that, when executed by the processor, causes the processor to respond to requests for exploration of concepts by: receiving a query comprising a concept; retrieving a cluster containing the concept; and displaying the retrieved cluster, wherein the concepts are automatically identified by: filtering the interactions based on intervals; creating a plurality of
20 sentences from the filtered interactions; computing a saliency of each the sentences; pruning away sentences with low saliency for generating a set of informative sentences; and clustering the sentences of the set of informative sentences for generating a plurality of sentence clusters, each of the clusters corresponding to a concept of the concepts.

[0027] Aspects of embodiments of the present invention are directed to addressing two
25 issues related to discovery and exploration of data. One aspect of the present invention is directed to the automatic discovery and extraction of concepts from a set of documents without human assistance. In addition, embodiments of the present invention are directed to understanding why a certain document or phone call is classified into a predefined category or why the document or phone call contains a particular concept (which may be newly
30 identified). In other words, embodiments of the present invention are directed to discovering what is the “root cause” for various “symptoms” reported by customers and inferring a predefined category or newly discovered concept “A” that is the root cause for the call to be categorized with a predefined category or newly discovered concept “B”, so an organization or a user can resolve the underlying issue.

35 [0028] According to one embodiment of the present invention, a method for determining a cause of events detected in a plurality of interactions includes: identifying, on a processor, a plurality of elements in the interactions; detecting, on the processor, a plurality of sequences of elements in the interactions; mining, on the processor, the plurality of sequences for

1 generating a set of supported patterns; computing, on the processor, association rules from the set of supported patterns; and returning the computed association rules.

[0029] The elements may include defined topics and identified concepts.

5 [0030] The identified concepts may be derived by: filtering, on the processor, the interactions based on intervals; creating, on the processor, a plurality of sentences from the filtered interactions; computing, on the processor, a saliency of each the sentences; pruning away, on the processor, sentences with low saliency for generating a set of informative sentences; clustering, on the processor, the sentences of the set of informative sentences for generating a plurality of sentence clusters, each of the clusters corresponding to a concept of the concepts; computing, on the processor, a saliency of each of the clusters; and naming, on the processor, each of the clusters.

10 [0031] The detecting the sequences of elements may include: sorting the identified elements in each interaction by timestamp within the interaction for generating the plurality of sequences; and condensing each of the sequences for removing repeated consecutive elements.

15 [0032] The method may further include condensing, on the processor, the set of supported patterns for removing repeated elements in the sequences.

20 [0033] The computing the association rules from the set of condensed supported patterns may include: computing an association rule for each sequence pattern in the condensed supported patterns, the computing the association rule including, if the sequence pattern includes more than one element: dividing the sequence pattern into a first portion and a second portion, the second portion including the last element in the sequence pattern; generating a proposed association rule, the proposed association rule being a logical implication from the first portion to the second portion; computing a confidence of the proposed association rule; if the computed confidence is greater than a threshold confidence level, storing the proposed association rule and moving the last element in the first portion to the second portion and iteratively repeating generating the proposed association rule and computing the confidence; and if the computed confidence is less than a threshold confidence level, ending analysis of the sequence pattern and returning a previously stored proposed association rule as the association rule or returning no association rule if no such proposed association rules were stored.

25 [0034] The method may further include computing a lift and a saliency of the association rule.

30 [0035] The method may further include computing a lift and a saliency of the proposed association rule.

35 [0036] According to one embodiment of the present invention, a method for determining a root cause of an event detected in a plurality of interactions includes: receiving, on a processor, a query for the root cause of the event; searching, on the processor, a plurality of

1 association rules, each of the association rules including one or more first portion elements
and one or more second portion elements, each of the association rules being a logical
implication from the first portion to the second portion; and returning, from the processor,
one or more association rules matching the query, the second portion elements of each of the
5 matching association rules including the event.

[0037] Each of the association rules may have a corresponding confidence value, the
query may include a confidence threshold, and the returning the one or more association rules
matching the query may further include returning association rules matching the query that
have confidence values exceeding the confidence threshold.

10 [0038] Each of the association rules may have a corresponding lift value and a
corresponding saliency value, the query may include at least one of a lift threshold and a
saliency threshold, and the returning the one or more association rules matching the query
may further include returning association rules matching the query that have lift values or
saliency values exceeding the lift threshold or the saliency threshold.

15 [0039] According to one embodiment of the present invention, a computer system
includes: a processor; and a memory, wherein the memory has stored thereon instructions
that, when executed by the processor, causes the processor to compute association rules
between events detected in a plurality of interactions by: identifying a plurality of elements in
the interactions; detecting a plurality of sequences of elements in the interactions; mining the
20 plurality of sequences for generating a set of supported patterns; computing association rules
from the set of supported patterns; and returning the computed association rules.

[0040] The elements may include defined topics and identified concepts.

[0041] The identified concepts may be derived by: filtering, on the processor, the
interactions based on intervals; creating, on the processor, a plurality of sentences from the
25 filtered interactions;

[0042] computing, on the processor, a saliency of each the sentences; pruning away, on
the processor, sentences with low saliency for generating a set of informative sentences;
clustering, on the processor, the sentences of the set of informative sentences for generating a
plurality of sentence clusters, each of the clusters corresponding to a concept of the concepts;
30 computing, on the processor, a saliency of each of the clusters; and naming, on the processor,
each of the clusters.

[0043] The processor may be further configured to compute association rules between
events detected in a plurality of interactions by: sorting the identified elements in each
interaction by timestamp within the interaction for generating the plurality of sequences; and
35 condensing each of the sequences for removing repeated consecutive elements.

[0044] The processor may be further configured to compute association rules between
events detected in a plurality of interactions by condensing the set of supported patterns for
removing repeated elements in the sequences.

1 [0045] The computing the association rules from the set of condensed supported patterns
may include: computing an association rule for each sequence pattern in the condensed
supported patterns, the computing the association rule including, if the sequence pattern
includes more than one element: dividing the sequence pattern into a first portion and a
5 second portion, the second portion including the last element in the sequence pattern;
generating a proposed association rule, the proposed association rule being a logical
implication from the first portion to the second portion; computing a confidence of the
proposed association rule; if the computed confidence is greater than a threshold confidence
level, storing the proposed association rule and moving the last element in the first portion to
10 the second portion and iteratively repeating generating the proposed association rule and
computing the confidence; and if the computed confidence is less than a threshold confidence
level, ending analysis of the sequence pattern and returning a previously stored proposed
association rule as the association rule or returning no association rule if no such proposed
association rules were stored.

15 [0046] The computing the association rule may further include computing a lift and a
saliency of the association rule.

[0047] The computing the association rule may further include computing a lift and a
saliency of the proposed association rule.

20 [0048] According to one embodiment of the present invention, a computer system
includes: a processor; and a memory, wherein the memory has stored thereon instructions
that, when executed by the processor, causes the processor to respond to a query for the root
cause of an event by: receiving the query for the root cause of the event; searching a plurality
of association rules, each of the association rules including one or more first portion elements
and one or more second portion elements, each of the association rules being a logical
25 implication from the first portion to the second portion; and returning one or more association
rules matching the query, the second portion elements of each of the matching association
rules including the event.

BRIEF DESCRIPTION OF THE DRAWINGS

30 [0049] FIG. 1 is a schematic block diagram of a system supporting a contact center that is
configured to provide access to searchable transcripts to customer service agents according to
one exemplary embodiment of the invention.

[0050] FIG. 2A is a block diagram of a computing device according to an embodiment of
the present invention.

35 [0051] FIG. 2B is a block diagram of a computing device according to an embodiment of
the present invention.

[0052] FIG. 2C is a block diagram of a computing device according to an embodiment of
the present invention.

1 [0053] FIG. 2D is a block diagram of a computing device according to an embodiment of the present invention.

[0054] FIG. 2E is a block diagram of a network environment including several computing devices according to an embodiment of the present invention.

5 [0055] FIG. 3 is a screenshot of a category distribution report according to one embodiment of the present invention.

[0056] FIG. 4 is a screenshot illustrating an interface for customizing and defining predefined categories according to one embodiment of the present invention.

10 [0057] FIG. 5 is a screenshot illustrating an interface for exploring relationships between topics in a plurality of interactions according to one embodiment of the present invention.

[0058] FIG. 6 illustrates a user interface for exploring clustering of key terms according to one embodiment of the present invention.

[0059] FIG. 7 is a flowchart illustrating a method for detecting topics within interactions according to one embodiment of the present invention.

15 [0060] FIG. 8 is a flowchart illustrating a method for clustering sentences according to one embodiment of the present invention.

[0061] FIG. 9 is a screenshot listing deduced association rules between causes and events along with support, confidence, lift, and saliency levels for each of the derived inference rules according to one embodiment of the present invention.

20 [0062] FIG. 10 is a flowchart illustrating a method for determining causes of events according to one embodiment of the present invention.

[0063] FIG. 11 is a flowchart illustrating a method for generating association rules according to one embodiment of the present invention.

25 [0064] FIG. 12 is an illustration an output of the method for determining causes of events according to one embodiment of the present invention.

DETAILED DESCRIPTION

30 [0065] In the following detailed description, only certain exemplary embodiments of the present invention are shown and described, by way of illustration. As those skilled in the art would recognize, the invention may be embodied in many different forms and should not be construed as being limited to the embodiments set forth herein. Like reference numerals designate like elements throughout the specification.

35 [0066] As described herein, various applications and aspects of the present invention may be implemented in software, firmware, hardware, and combinations thereof. When implemented in software, the software may operate on a general purpose computing device such as a server, a desktop computer, a tablet computer, a smartphone, or a personal digital assistant. Such a general purpose computer includes a general purpose processor and memory.

1 [0067] Some embodiments of the present invention will be described in the context of a
contact center. However, embodiments of the present invention are not limited thereto and
may also be used in under other conditions involving searching recorded audio such as in
5 computer based education systems, voice messaging systems, medical transcripts, or any
speech corpora from any source.

[0068] Aspects of embodiments of the present invention are directed to a system and
method for automatically inferring and deducing topics of discussion (or “concepts”) from a
body of recorded or live interactions (or conversations). These interactions may include, for
example, telephone conversations, text-based chat sessions, email conversation threads, and
10 the like. The inferring of these concepts does not require manual categorization by a human
and can be performed by the system (or the “analytics system”) according to embodiments of
the present invention. Therefore, new, previously unidentified topics of conversation can
quickly be identified and brought to the attention of an organization without performing a
manual analysis of conversation logs.

15 [0069] For example, suppose a company released a new version product that added
Bluetooth® connectivity and there were no predefined categories in the interactions analytics
system to match the phrases “Bluetooth connection” or “Bluetooth pairing” to issues with
Bluetooth® connections. In conventional systems, this category might go undetected until
those phrases were manually added to the analytics system.

20 [0070] In contrast, embodiments of the present invention are directed to a system and
method for identifying salient phrases, generating new categories (or “concepts”) based on
these identified phrases, and categorizing interactions based on these automatically identified
categories. As a result, embodiments of the present invention can be used to alert
organizations to newly trending topics within interactions (e.g., conversations with
25 customers), thereby allowing faster responses to changing circumstances. See, e.g., FIG. 3,
which is a screenshot of a portion of a category distribution report 1 showing exemplary
categories “New Customer,” “Emergency,” “Identification,” “Billing,” and “Payment
Inquiry” along the number of interactions categorized into each of these categories and the
percentages of all calls that involve these categories. Note that the percentages add up to
30 more than 100% because any given interaction may be assigned to multiple categories or not
assigned to any category. Viewing this category distribution report, an organization can
assess the most frequently discussed topics.

[0071] Another aspect of embodiments of the present invention is directed to systems and
methods for automatically determining possible root causes of events and concepts within a
35 conversation. For example, an internet service provider may be alerted to a large number of
requests to cancel service plans. Using embodiments of the present invention, the
organization would be able to view association rules (or inference rules) automatically
deduced from the interactions to identify one or more causes (or “root causes”) of these

1 events based on other categories and concepts appearing in the interaction. See, e.g., FIG. 9,
which is a screenshot listing deduced association rules between causes (labeled “Left Hand
Side”) and events (labeled “Right Hand Side”) along with support, confidence, lift, and
saliency levels for each of the derived inference rules according to one embodiment of the
5 present invention. A user can search for rules relating to a particular topic (e.g., customer
“Dissatisfaction”) by selecting the topic from the “Target Topic” dropdown box, in which
case rules containing “Dissatisfaction” on the “Right Hand Side” will be shown. In the
example shown in FIG. 9, the rules “Transfer + Dispute → Dissatisfaction” and “No Payment
+ On Hold → Dissatisfaction” would be shown among the rules involving the
10 “Dissatisfaction” concept.

[0072] Therefore, embodiments of the present invention are directed to systems and
methods for providing timely summary of trends in topics of discussion in a collection of
interactions and systems and methods for determining root causes of predefined and inferred
topics of discussion based on, for example, correlations with particular other topics within the
15 same interaction.

[0073] In one embodiment, the above-described systems and methods are used in the
context of a contact center and are used to monitor and infer topics of conversation during
interactions between customers and an organization and to analyze and determine root causes
of events for display to members of the organization.

20 [0074] FIG. 1 is a schematic block diagram of a system supporting a contact center that is
configured to provide customer availability information to customer service agents according
to one exemplary embodiment of the invention. The contact center may be an in-house
facility to a business or corporation for serving the enterprise in performing the functions of
sales and service relative to the products and services available through the enterprise. In
25 another aspect, the contact center may be a third-party service provider. The contact center
may be hosted in equipment dedicated to the enterprise or third-party service provider, and/or
hosted in a remote computing environment such as, for example, a private or public cloud
environment with infrastructure for supporting multiple contact centers for multiple
enterprises.

30 [0075] According to one exemplary embodiment, the contact center includes resources
(e.g. personnel, computers, and telecommunication equipment) to enable delivery of services
via telephone or other communication mechanisms. Such services may vary depending on the
type of contact center, and may range from customer service to help desk, emergency
response, telemarketing, order taking, and the like.

35 [0076] Customers, potential customers, or other end users (collectively referred to as
customers) desiring to receive services from the contact center may initiate inbound calls to
the contact center via their end user devices 10a-10c (collectively referenced as 10). Each of
the end user devices 10 may be a communication device conventional in the art, such as, for

1 example, a telephone, wireless phone, smart phone, personal computer, electronic tablet, and/or the like. Users operating the end user devices 10 may initiate, manage, and respond to telephone calls, emails, chats, text messaging, web-browsing sessions, and other multi-media transactions.

5 [0077] Inbound and outbound calls from and to the end users devices 10 may traverse a telephone, cellular, and/or data communication network 14 depending on the type of device that is being used. For example, the communications network 14 may include a private or public switched telephone network (PSTN), local area network (LAN), private wide area network (WAN), and/or public wide area network such as, for example, the Internet. The
10 communications network 14 may also include a wireless carrier network including a code division multiple access (CDMA) network, global system for mobile communications (GSM) network, and/or any 3G or 4G network conventional in the art.

[0078] According to one exemplary embodiment, the contact center includes a switch/media gateway 12 coupled to the communications network 14 for receiving and
15 transmitting calls between end users and the contact center. The switch/media gateway 12 may include a telephony switch configured to function as a central switch for agent level routing within the center. In this regard, the switch 12 may include an automatic call distributor, a private branch exchange (PBX), an IP-based software switch, and/or any other switch configured to receive Internet-sourced calls and/or telephone network-sourced calls.
20 According to one exemplary embodiment of the invention, the switch is coupled to a call server 18 which may, for example, serve as an adapter or interface between the switch and the remainder of the routing, monitoring, and other call-handling systems of the contact center.

[0079] The contact center may also include a multimedia/social media server for
25 engaging in media interactions other than voice interactions with the end user devices 10 and/or web servers 32. The media interactions may be related, for example, to email, vmail (voice mail through email), chat, video, text-messaging, web, social media, screen-sharing, and the like. The web servers 32 may include, for example, social interaction site hosts for a variety of known social interaction sites to which an end user may subscribe, such as, for
30 example, Facebook, Twitter, and the like. The web servers may also provide web pages for the enterprise that is being supported by the contact center. End users may browse the web pages and get information about the enterprise's products and services. The web pages may also provide a mechanism for contacting the contact center, via, for example, web chat, voice call, email, web real time communication (WebRTC), or the like.

35 [0080] According to one exemplary embodiment of the invention, the switch is coupled to an interactive voice response (IVR) server 34. The IVR server 34 is configured, for example, with an IVR script for querying customers on their needs. For example, a contact center for a bank may tell callers, via the IVR script, to "press 1" if they wish to get an

1 account balance. If this is the case, through continued interaction with the IVR, customers may complete service without needing to speak with an agent.

[0081] If the call is to be routed to an agent, the call is forwarded to the call server 18 which interacts with a routing server 20 for finding an appropriate agent for processing the call. The call server 18 may be configured to process PSTN calls, VoIP calls, and the like. For example, the call server 18 may include a session initiation protocol (SIP) server for processing SIP calls.

[0082] In one example, while an agent is being located and until such agent becomes available, the call server may place the call in, for example, a call queue. The call queue may be implemented via any data structure conventional in the art, such as, for example, a linked list, array, and/or the like. The data structure may be maintained, for example, in buffer memory provided by the call server 18.

[0083] Once an appropriate agent is available to handle a call, the call is removed from the call queue and transferred to a corresponding agent device 38a-38c (collectively referenced as 38). Collected information about the caller and/or the caller's historical information may also be provided to the agent device for aiding the agent in better servicing the call. In this regard, each agent device 38 may include a telephone adapted for regular telephone calls, VoIP calls, and the like. The agent device 38 may also include a computer for communicating with one or more servers of the contact center and performing data processing associated with contact center operations, and for interfacing with customers via a variety of communication mechanisms such as chat, instant messaging, voice calls, and the like.

[0084] The selection of an appropriate agent for routing an inbound call may be based, for example, on a routing strategy employed by the routing server 20, and further based on information about agent availability, skills, and other routing parameters provided, for example, by a statistics server 22. According to one exemplary embodiment of the invention, the statistics server 22 includes a customer availability aggregation (CAA) module 36 for monitoring availability of end users on different communication channels and providing such information to, for example, the routing server 20, agent devices 38a-38c, and/or other contact center applications and devices. The CAA module may also be deployed in a separate application server. The aggregation module 36 may be a software module implemented via computer program instructions which are stored in memory of the statistics server 22 (or some other server), and which program instructions are executed by a processor. A person of skill in the art should recognize that the aggregation module 36 may also be implemented via firmware (e.g. an application-specific integrated circuit), hardware, or a combination of software, firmware, and hardware.

[0085] According to one exemplary embodiment, the aggregation module 36 is configured to receive customer availability information from other devices in the contact

1 center, such as, for example, the multimedia/social media server 24. For example, the
multimedia/social media server 24 may be configured to detect user presence on different
websites including social media sites, and provide such information to the aggregation
module 36. The multimedia/social media server 24 may also be configured to monitor and
5 track interactions on those websites.

[0086] The multimedia/social media server 24 may also be configured to provide, to an
end user, a mobile application 40 for downloading onto the end user device 10. The mobile
application 40 may provide user configurable settings that indicate, for example, whether the
user is available, not available, or availability is unknown, for purposes of being contacted by
10 a contact center agent. The multimedia/social media server 24 may monitor the status settings
and send updates to the aggregation module each time the status information changes.

[0087] The contact center may also include a reporting server 28 configured to generate
reports from data aggregated by the statistics server 22. Such reports may include near real-
time reports or historical reports concerning the state of resources, such as, for example,
15 average waiting time, abandonment rate, agent occupancy, and the like. The reports may be
generated automatically or in response to specific requests from a requestor (e.g.
agent/administrator, contact center application, and/or the like).

[0088] According to one exemplary embodiment of the invention, the routing server 20 is
enhanced with functionality for managing back-office/offline activities that are assigned to
20 the agents. Such activities may include, for example, responding to emails, responding to
letters, attending training seminars, or any other activity that does not entail real time
communication with a customer. Once assigned to an agent, an activity an activity may be
pushed to the agent, or may appear in the agent's workbin 26a-26c (collectively referenced as
26) as a task to be completed by the agent. The agent's workbin may be implemented via any
25 data structure conventional in the art, such as, for example, a linked list, array, and/or the like.
The workbin may be maintained, for example, in buffer memory of each agent device 38.

[0089] According to one exemplary embodiment of the invention, the contact center also
includes one or more mass storage devices 30 for storing different databases relating to agent
data (e.g. agent profiles, schedules, etc.), customer data (e.g. customer profiles), interaction
30 data (e.g. details of each interaction with a customer, including reason for the interaction,
disposition data, time on hold, handle time, etc.), and the like. According to one embodiment,
some of the data (e.g. customer profile data) may be provided by a third party database such
as, for example, a third party customer relations management (CRM) database. The mass
storage device may take form of a hard disk or disk array as is conventional in the art.

35 [0090] According to one embodiment of the present invention, the contact center 102 also
includes a call recording server 40 for recording the audio of calls conducted through the
contact center 102, a call recording storage server 42 for storing the recorded audio, a speech

1 analytics server 44 configured to process and analyze audio collected in the from the contact center 102, and a speech index database 46 for providing an index of the analyzed audio.

[0091] The speech analytics server 44 may be coupled to (or may include) an analytics server 45 including a topic detecting module 45a, a root cause mining module 45b, and a user interface module 45c. The analytics server 45 may be configured to provide the automatic detection of topics from interactions recorded by the call recording server 40 and stored on the call recording storage server 42. The analytics server 45 may also access data stored on, for example, the multimedia/social media server 24 in order to process interactions from various chat, social media, email, and other non-voice interactions.

10 [0092] The various servers of FIG. 1 may each include one or more processors executing computer program instructions and interacting with other system components for performing the various functionalities described herein. The computer program instructions are stored in a memory implemented using a standard memory device, such as, for example, a random access memory (RAM). The computer program instructions may also be stored in other non-transitory computer readable media such as, for example, a CD-ROM, flash drive, or the like. Also, although the functionality of each of the servers is described as being provided by the particular server, a person of skill in the art should recognize that the functionality of various servers may be combined or integrated into a single server, or the functionality of a particular server may be distributed across one or more other servers without departing from the scope of the embodiments of the present invention.

20 [0093] Each of the various servers in the contact center may be a process or thread, running on one or more processors, in one or more computing devices 500 (e.g., FIG. 2A, FIG. 2B), executing computer program instructions and interacting with other system components for performing the various functionalities described herein. The computer program instructions are stored in a memory which may be implemented in a computing device using a standard memory device, such as, for example, a random access memory (RAM). The computer program instructions may also be stored in other non-transitory computer readable media such as, for example, a CD-ROM, flash drive, or the like. Also, a person of skill in the art should recognize that a computing device may be implemented via
30 firmware (e.g. an application-specific integrated circuit), hardware, or a combination of software, firmware, and hardware. A person of skill in the art should also recognize that the functionality of various computing devices may be combined or integrated into a single computing device, or the functionality of a particular computing device may be distributed across one or more other computing devices without departing from the scope of the
35 exemplary embodiments of the present invention. A server may be a software module, which may also simply be referred to as a module. The set of modules in the contact center may include servers and other modules.

1 [0094] FIG. 2A and FIG. 2B depict block diagrams of a computing device 500 as may be employed in exemplary embodiments of the present invention. As shown in FIG. 2A and FIG. 2B, each computing device 500 includes a central processing unit 521, and a main memory unit 522. As shown in FIG. 2A, a computing device 500 may include a storage device 528, a removable media interface 516, a network interface 518, an input/output (I/O) controller 523, one or more display devices 530c, a keyboard 530a and a pointing device 530b, such as a mouse. The storage device 528 may include, without limitation, storage for an operating system and software. As shown in FIG. 2B, each computing device 500 may also include additional optional elements, such as a memory port 503, a bridge 570, one or more additional input/output devices 530d, 530e and a cache memory 540 in communication with the central processing unit 521. Input/output devices, e.g., 530a, 530b, 530d, and 530e, may be referred to herein using reference numeral 530.

15 [0095] The central processing unit 521 is any logic circuitry that responds to and processes instructions fetched from the main memory unit 522. It may be implemented, for example, in an integrated circuit, in the form of a microprocessor, microcontroller, or graphics processing unit (GPU), or in a field-programmable gate array (FPGA) or application-specific integrated circuit (ASIC). Main memory unit 522 may be one or more memory chips capable of storing data and allowing any storage location to be directly accessed by the central processing unit 521. In the embodiment shown in FIG. 2A, the central processing unit 521 communicates with main memory 522 via a system bus 550. FIG. 2B depicts an embodiment of a computing device 500 in which the central processing unit 521 communicates directly with main memory 522 via a memory port 503.

25 [0096] FIG. 2B depicts an embodiment in which the central processing unit 521 communicates directly with cache memory 540 via a secondary bus, sometimes referred to as a backside bus. In other embodiments, the central processing unit 521 communicates with cache memory 540 using the system bus 550. Cache memory 540 typically has a faster response time than main memory 522. In the embodiment shown in FIG. 2A, the central processing unit 521 communicates with various I/O devices 530 via a local system bus 550. Various buses may be used as a local system bus 550, including a Video Electronics Standards Association (VESA) Local bus (VLB), an Industry Standard Architecture (ISA) bus, an Extended Industry Standard Architecture (EISA) bus, a MicroChannel Architecture (MCA) bus, a Peripheral Component Interconnect (PCI) bus, a PCI Extended (PCI-X) bus, a PCI-Express bus, or a NuBus. For embodiments in which an I/O device is a display device 530c, the central processing unit 521 may communicate with the display device 530c through an Advanced Graphics Port (AGP). FIG. 2B depicts an embodiment of a computer 500 in which the central processing unit 521 communicates directly with I/O device 530e. FIG. 2B also depicts an embodiment in which local busses and direct communication are mixed: the

1 central processing unit 521 communicates with I/O device 530d using a local system bus 550 while communicating with I/O device 530e directly.

[0097] A wide variety of I/O devices 530 may be present in the computing device 500. Input devices include one or more keyboards 530a, mice, trackpads, trackballs, microphones, and drawing tablets. Output devices include video display devices 530c, speakers, and printers. An I/O controller 523, as shown in FIG. 2A, may control the I/O devices. The I/O controller may control one or more I/O devices such as a keyboard 530a and a pointing device 530b, e.g., a mouse or optical pen.

[0098] Referring again to FIG. 2A, the computing device 500 may support one or more removable media interfaces 516, such as a floppy disk drive, a CD-ROM drive, a DVD-ROM drive, tape drives of various formats, a USB port, a Secure Digital or COMPACT FLASHTM memory card port, or any other device suitable for reading data from read-only media, or for reading data from, or writing data to, read-write media. An I/O device 530 may be a bridge between the system bus 550 and a removable media interface 516.

[0099] The removable media interface 516 may for example be used for installing software and programs. The computing device 500 may further comprise a storage device 528, such as one or more hard disk drives or hard disk drive arrays, for storing an operating system and other related software, and for storing application software programs. Optionally, a removable media interface 516 may also be used as the storage device. For example, the operating system and the software may be run from a bootable medium, for example, a bootable CD.

[00100] In some embodiments, the computing device 500 may comprise or be connected to multiple display devices 530c, which each may be of the same or different type and/or form. As such, any of the I/O devices 530 and/or the I/O controller 523 may comprise any type and/or form of suitable hardware, software, or combination of hardware and software to support, enable or provide for the connection to, and use of, multiple display devices 530c by the computing device 500. For example, the computing device 500 may include any type and/or form of video adapter, video card, driver, and/or library to interface, communicate, connect or otherwise use the display devices 530c. In one embodiment, a video adapter may comprise multiple connectors to interface to multiple display devices 530c. In other embodiments, the computing device 500 may include multiple video adapters, with each video adapter connected to one or more of the display devices 530c. In some embodiments, any portion of the operating system of the computing device 500 may be configured for using multiple display devices 530c. In other embodiments, one or more of the display devices 530c may be provided by one or more other computing devices, connected, for example, to the computing device 500 via a network. These embodiments may include any type of software designed and constructed to use the display device of another computing device as a second display device 530c for the computing device 500. One of ordinary skill in the art will

1 recognize and appreciate the various ways and embodiments that a computing device 500 may be configured to have multiple display devices 530c.

[00101] A computing device 500 of the sort depicted in FIG. 2A and FIG. 2B may operate under the control of an operating system, which controls scheduling of tasks and access to system resources. The computing device 500 may be running any operating system, any embedded operating system, any real-time operating system, any open source operating system, any proprietary operating system, any operating systems for mobile computing devices, or any other operating system capable of running on the computing device and performing the operations described herein.

[00102] The computing device 500 may be any workstation, desktop computer, laptop or notebook computer, server machine, handheld computer, mobile telephone or other portable telecommunication device, media playing device, gaming system, mobile computing device, or any other type and/or form of computing, telecommunications or media device that is capable of communication and that has sufficient processor power and memory capacity to perform the operations described herein. In some embodiments, the computing device 500 may be a virtualized computing device and the virtualized computing device may be running in a networked or cloud based environment. In some embodiments, the computing device 500 may have different processors, operating systems, and input devices consistent with the device.

[00103] In other embodiments the computing device 500 is a mobile device, such as a Java-enabled cellular telephone or personal digital assistant (PDA), a smart phone, a digital audio player, or a portable media player. In some embodiments, the computing device 500 comprises a combination of devices, such as a mobile phone combined with a digital audio player or portable media player.

[00104] As shown in FIG. 2C, the central processing unit 521 may comprise multiple processors P1, P2, P3, P4, and may provide functionality for simultaneous execution of instructions or for simultaneous execution of one instruction on more than one piece of data. In some embodiments, the computing device 500 may comprise a parallel processor with one or more cores. In one of these embodiments, the computing device 500 is a shared memory parallel device, with multiple processors and/or multiple processor cores, accessing all available memory as a single global address space. In another of these embodiments, the computing device 500 is a distributed memory parallel device with multiple processors each accessing local memory only. In still another of these embodiments, the computing device 500 has both some memory which is shared and some memory which may only be accessed by particular processors or subsets of processors. In still even another of these embodiments, the central processing unit 521 comprises a multicore microprocessor, which combines two or more independent processors into a single package, e.g., into a single integrated circuit

1 (IC). In one exemplary embodiment, depicted in FIG. 2D, the computing device 500 includes at least one central processing unit 521 and at least one graphics processing unit 521'.

[00105] In some embodiments, a central processing unit 521 provides single instruction, multiple data (SIMD) functionality, e.g., execution of a single instruction simultaneously on
5 multiple pieces of data. In other embodiments, several processors in the central processing unit 521 may provide functionality for execution of multiple instructions simultaneously on multiple pieces of data (MIMD). In still other embodiments, the central processing unit 521 may use any combination of SIMD and MIMD cores in a single device.

[00106] A computing device may be one of a plurality of machines connected by a
10 network, or it may comprise a plurality of machines so connected. FIG. 2E shows an exemplary network environment. The network environment comprises one or more local machines 502a, 502b (also generally referred to as local machine(s) 502, client(s) 502, client node(s) 502, client machine(s) 502, client computer(s) 502, client device(s) 502, endpoint(s) 502, or endpoint node(s) 502) in communication with one or more remote machines 506a,
15 506b, 506c (also generally referred to as server machine(s) 506 or remote machine(s) 506) via one or more networks 504. In some embodiments, a local machine 502 has the capacity to function as both a client node seeking access to resources provided by a server machine and as a server machine providing access to hosted resources for other clients 502a, 502b. Although only two clients 502 and three server machines 506 are illustrated in FIG. 2E, there
20 may, in general, be an arbitrary number of each. The network 504 may be a local-area network (LAN), e.g., a private network such as a company Intranet, a metropolitan area network (MAN), or a wide area network (WAN), such as the Internet, or another public network, or a combination thereof.

[00107] The computing device 500 may include a network interface 518 to interface to the
25 network 504 through a variety of connections including, but not limited to, standard telephone lines, local-area network (LAN), or wide area network (WAN) links, broadband connections, wireless connections, or a combination of any or all of the above. Connections may be established using a variety of communication protocols. In one embodiment, the computing device 500 communicates with other computing devices 500 via any type and/or
30 form of gateway or tunneling protocol such as Secure Socket Layer (SSL) or Transport Layer Security (TLS). The network interface 518 may comprise a built-in network adapter, such as a network interface card, suitable for interfacing the computing device 500 to any type of network capable of communication and performing the operations described herein. An I/O device 530 may be a bridge between the system bus 550 and an external communication bus.

[00108] According to various embodiments of the present invention, exploration and
35 discovery technologies are directed toward discovering interesting phenomena without the user input—in other words, identifying information that is relevant to the user without the user explicitly looking for this information. Categorization technologies are focused on

1 classifying documents (e.g., text, audio, and video) into predefined categories such as “all the calls in which a customer has asked to speak to a supervisor.”

[00109] FIG. 3 is a screenshot of a category distribution report according to one embodiment of the present invention. In this report, the voice calls, customer-agent phone
5 conversations (or interactions) that have occurred in the last 7 days have been classified into categories (e.g., predefined categories) that represent the set of known reasons for calls. In other embodiments, conversations are aggregated over different time periods (e.g., over the past day, over the past hour, over the past month, since a particular date, or between two arbitrary dates) may be aggregated. In addition, in other embodiments, the interactions may
10 be limited to particular communication channels, such as one or more of telephone, email, chat, and social media, limited to interactions from particular contact centers, or limited to interactions from particular departments (e.g., sales or customer support).

[00110] FIG. 4 is a screenshot illustrating an interface for customizing and defining predefined categories according to one embodiment of the present invention. Each predefined
15 category can be defined as some Boolean expression of topics where each topic may be defined as a union of phrases or words, thereby producing a set of categorizing rules used to classify the interactions. For example, FIG. 4 illustrates the definition of the “Repeat Call or Contact” category, which is defined by interactions having ‘Found topic “Repeat Calls” at least once with Very-Low strictness OR Found topic “Repeat Contacts” at least once with
20 Very-Low strictness’. The “Repeat Calls” and “Repeat Contact” topics may be triggered, for example, by detecting particular triggering events such as a record of multiple calls from a particular phone number or by identifying particular phrases in the interaction such as “thanks for calling again”.

[00111] When one of these phrases of the Boolean expression is spoken in a conversation,
25 various speech recognition technologies can recognize it in the audio. (One such technology is phrase recognition as described in U.S. Pat. 7,487,094 “System and method of call classification with context modeling based on composite words” the content of which is incorporated herein by reference) In other embodiments, the interactions are conducted over other media (for example, text chat) and other appropriate methods of detecting phrases are
30 used. Upon detecting one of these phrases, it triggers the detection of topics to which these phrases belong. The detection of topics feed the appropriate categorizing rules matching the given category are triggered and the interaction is labeled in accordance with the matching categories.

[00112] Therefore, according to one embodiment, the analytics server 45 can generate the
35 category distribution report by counting the number of interactions within a given time period that fall within each category.

[00113] According to one embodiment of the present invention, an analytics server 45 provides a user with the ability to view or “explore” related words, as illustrated, for example,

1 in FIG. 5. A user can start from a single word and explore the co-occurrence of the starting word with other words in various conversations. For instance, FIG. 5 depicts the relationship or co-occurrence of the word “credit” with other words in the set of relevant calls.

5 [00114] According to one embodiment of the invention, the analytics server 45 and the topic detecting module 45a provide a user interface through user interface module 45c for a user to select a category from a set of categories (e.g., a set of predefined categories) and a group of calls that were classified into this category in a certain interval (e.g., last 7 days) and request information regarding why the group of calls were classified into the given category. The analytics server 45 and the user interface module 45c can present to the user, words and
10 phrases that often occur and that are informative around occurrences of this category (e.g., words that appear with higher frequency than their occurrence in general within this category), thus characterizing the cause, or what is caused by, this category. Systems and methods for categorizing interactions will be described below with respect to FIG. 7.

15 [00115] FIG. 6 illustrates a user interface 6 for exploring clustering of key terms (or “buzz around categories”) according to one embodiment of the present invention. Referring to FIG. 6, a time range of documents may be specified in field 61. The language of the interactions (or documents) and general field (e.g., business category) may also be specified in field 62, and the types of interactions (e.g., audio from telephone calls, chats, email, and social media interactions) may be selected in checkboxes 63. As shown in FIG. 6, various defined
20 categories 67 of calls, such as “Account Balance,” “Billing Issue,” “Escalation,” “Dissatisfaction,” and “Positive Feedback” are listed, with each row showing the percentage of interactions being classified into the given categories, the total number of interactions matching that category, and percentage change from a prior period (e.g., if the percentages currently shown are for the past week, the percentage change may be shown in comparison
25 to, for example, two weeks ago or, as another example, the same week last year).

[00116] When a row corresponding to a category is selected (e.g., in FIG. 6, the “Dissatisfaction” row is selected), additional information regarding the selected category may be shown in pane 65 of the user interface. As seen in FIG. 6, phrases from the interactions that are common to interactions within the selected category are shown in pane 65. In some
30 embodiments, the size of the words may correspond to the frequency with which the phrases appear in the interaction. In addition, in some embodiments the user interface may also show a graph showing the number of interactions in this category over time.

[00117] FIG. 7 is a flowchart of a process executed by the analytics server 45 and the topic detecting module 45a for extracting concepts from interactions (e.g., text and text
35 transcriptions of audio) according to one embodiment of the present invention. According to one embodiment, sentences that are semantically related are grouped together (or clustered) as conveying the same idea. Clustering is a machine learning technique that can be used to take sentences as input and to cluster the sentences together when the important portions of

the sentences are appear to be the similar or the same. Each one of these clusters is a concept as mentioned above. Thus, to find buzz (or activity) around (or related to) particular categories, the text that appears around the location of a phrase associated with a category (for example, 30 seconds before and after the phrase), and perform the clustering on this text, as explained in more detail below.

[00118] Referring to FIG. 7, according to one embodiment of the present invention, concepts or phrases are extracted from interactions by supplying the entire body of interactions (or the entire body of text) to the system (e.g., the analytics server 45 as shown in FIG. 1, which may be a computer system 500 as shown in FIG. 2A, including the topic detecting module 45a as shown in FIG. 1) configured to perform the categorization.

[00119] When the interactions being processed are the output of large-vocabulary continuous speech recognition (LVCSR), then low confidence words are filtered out in operation 100 and only words with high confidence remain in the text to be processed. If the exploration is done on email, chat, or other text, the entire text is used.

[00120] For exploration around topics, according to one embodiment, only text from a within a particular time boundary before and after the topic recognition (e.g., 30 seconds before and after the topic) are used and other data is filtered out in operation 101.

[00121] In operation 102, “sentences” are created out of the filtered text. As used and defined herein, “sentence” refers to a block of consecutive words in the text and this block of words does not necessarily correspond to a particular grammatical or orthographical unit (e.g., a complete grammatical sentence or a sequence of words beginning with a capital letter and ending in a period). For example, in one embodiment, n -grams with overlaps (overlapping n -grams) are created, where n may be 4, which means that every consecutive 4-words form a “sentence.” However, in other embodiments of the present invention, n can be any other whole number greater than 1. In addition, n -grams are merely one way to create sentences from words, and in other embodiments, other methods of forming “sentences” or blocks of consecutive words can be applied.

[00122] In operation 103, the saliency of a sentence is computed from the saliency of the words in the sentence. According to one embodiment, the inverse document frequency (IDF) of a word is used to measure the saliency of word w , and the saliency of the sentence is given by the square of the sum of the IDF of each of the words in the sentence:

$$IDF(w) = \log\left(\frac{N}{DF(w)}\right)$$

$$Saliency(sentence) = \left(\sum_{w \in sentence} IDF(w) \right)^2$$

where N is the total number of documents in the collection and $DF(w)$ is the number of documents in which the word w appears.

[00123] To make the clustering faster, in one embodiment, only the most salient sentences are used. As such, in operation 104, the sentences are pruned by sorting the sentences by saliency and discarding the sentences with low saliency relative to the top ones. For example, in one embodiment, sentences with less than 5% of the saliency of the top ones are removed from consideration. In operation 105, the sentences are clustered to group together similar sentences that differ from one another only by less-salient words. The similarity of sentences can be measured based on various text mining measures, and is described in more detail below.

[00124] In operation 106, the saliency of each cluster is computed based on text mining measures. According to one embodiment, the saliency of a cluster is constructed from a weighted sum of the saliencies of the sentences of the cluster:

$$S(\text{Cluster}) = \sum_{n\text{-gram} \in \text{cluster}} \frac{\text{sentence weight}}{\text{cluster size}} \text{Saliency}(\text{sentence})$$

[00125] In a manner similar to that described for sentence pruning, in one embodiment, only the top clusters will be presented to the user and clusters with lower saliencies may be pruned away.

[00126] In operation 107, the clusters are named with a sequence of words that is both informative and frequent in the cluster is selected. For all the sentences (or n -grams) in the clusters of orders $N_{\text{clustering}} \geq N \geq 2$, where N is the n -gram level, the length of the sentence, we compute:

$$TFIDF_{NG} = IDF_{NG} \sqrt{TF_{NG}}$$

where the IDF of an n -gram is defined above, and the TF of the n -gram is the number of times this n -gram appears in the cluster.

[00127] Note that this formulation balances between long and short sequences of words. Long sequences will have larger IDF , since the IDF is constructed from the sum of word $IDFs$, but a lower TF . Short sequences will have a larger count (or TF) but lower IDF .

[00128] As such, the cluster can be named by the n -gram of the cluster having the largest $TFIDF_{NG}$.

[00129] According to one embodiment, a measure of similarity between n -grams is formulated in which two n -grams are more similar if they contain the same important words, and less similar if the important words are not shared between them. The following similarity measure Sim exhibits these characteristics:

$$\begin{aligned} \text{Both}(\text{sent}_1, \text{sent}_2) &= \sum_{w \in \text{both sentences}} IDF(w) \\ \text{SIOnly}(\text{sent}_1, \text{sent}_2) &= \sum_{w \in \text{sent}_1 \text{ but not } \text{sent}_2} IDF(w) \end{aligned}$$

$$S2Only(sent_1, sent_2) = \sum_{w \in sent_2 \text{ but not } sent_1} IDF(w)$$

$$Sim(sent_1, sent_2) = \frac{Both(sent_1, sent_2)}{Both(sent_1, sent_2) + 0.5Max(S1Only(sent_1, sent_2), S2Only(sent_1, sent_2))}$$

where $sent_1$ and $sent_2$ are sentences to be compared in similarity.

[00130] Using this measure, saliency that is contained in both sentences will increase the similarity and saliency that is contained in only one of the sentences will decrease the similarity.

[00131] FIG. 8 is a flowchart illustrating a method executed by the analytics server 45 and the topic detecting module 45a for clustering sentences according to one embodiment of the present invention.

[00132] Referring to FIG. 8, in operation 200, sentences are randomly selected to serve as centers (templates) for the clusters. This is a standard procedure in clustering algorithms, however, in embodiments of the present invention, the centers are not numeric vectors (as would be typical), but instead are sequences of words, and the clustering process is based on words that occur in both sentences, in a way that somewhat resembles. See, e.g., H. Ye and S. Young, "A Clustering Approach to Semantic Decoding" ICSLP 2006, Pittsburgh, PA (2006), the entire disclosure of which is incorporated herein by reference.

[00133] In operation 201, each of the sentences involved in the clustering process is compared to each of the cluster centers (templates) using a similarity formula such as the sentence similarity function Sim described above.

[00134] In operation 202, each sentence is assigned to the cluster that has highest similarity to it, provided that this similarity is also greater than a threshold similarity (e.g., a predefined threshold).

[00135] In operation 203, the process removes templates to which very few sentences or sentences are assigned.

[00136] In operation 204, if all of the sentences have been tried as templates or a certain number of iterations have been reached, then the ending conditions have been satisfied, and the clustering process ends. If the ending conditions are not satisfied, then the process returns to operation 200 and selects additional sentences to serve as templates, where the additional sentences are randomly selected from the set of sentences that have not yet been tried as templates.

[00137] Therefore, embodiments of the present invention can automatically identify and assign names to new topics of conversation based on detecting related phrases and label (or tag or classify) interactions as involving these topics without manual entry of trigger phrases by a user.

[00138] According to another embodiment of the invention, a root cause mining procedure can be employed on the concepts and/or categories assigned in an earlier processing stage to infer association rules (for example, logical implication) between them. For example, an

1 association rule between categories and/or concepts A, B, and C such as $A, B \rightarrow C$, can indicate that the root cause of C is having A and B occurring in the same call or document prior to C. Here, A, B, and C can be predefined categories or newly discovered concepts, or a mixture thereof. For illustration this is described in FIG. 9. FIG. 9 is a screenshot illustrating
5 a user interface for querying and viewing inferred association rules. The association rules are shown along with support and confidence levels of each of them, along with lift and saliency measures. Systems and methods for performing root cause analysis to generate the inferred association rules shown in FIG. 9 will be described below with respect to FIGS. 10, 11, and 12.

10 **[00139]** In another embodiment of the invention, the two approaches described above are combined: First, the user can select a category (e.g., category “C”) and then choose to “search” for relationships to other categories. The root cause mining module 45b of the analytics server 45 then performs mining of association rules of categories that cause category C, *i.e.*, rules having C on their right hand side (RHS) (for example $B \rightarrow C$). For
15 deeper investigation, the user can then choose to see the “buzz around relation.” The root cause mining module 45b can then extract concepts only from the set of interactions having this relation, in a similar way to “Buzz around categories” as shown in FIG. 5, which is a screenshot illustrating an interface for exploring relationships between topics in a plurality of interactions according to one embodiment of the present invention. In addition, the root cause
20 mining module 45b can be configured and constrained to look only on a part of an interaction that starts just before category B and ends just after category C.

[00140] In general, as used herein, a “term” (or “phrase”) is a part of a “topic” and a “topic” a part of “category.” The systems and methods described herein can be applied to any of these levels (term/phrase, topic, or category). For instance, the term “I want to speak to a
25 supervisor” can be part of “Escalation” topic which can be part of a category “the customer asks to escalate more than once at the same call”. Embodiments of the present invention can be applied to any and all of these levels and, for the sake of convenience, are described herein with respect to topics. However, embodiments of the present invention are not limited to use with topics and may be used at other levels, such as terms, phrases, and categories.

30 **[00141]** FIG. 9 is a screenshot listing deduced association rules between causes and events along with support, confidence, lift, and saliency levels for each of the derived inference rules according to one embodiment of the present invention. Support, confidence, lift, and saliency are computed metrics that may be viewed by a user to evaluate the quality of each of the derived inference rules and are described in more detail below. In this figure we present
35 association rules between topics but association rules between terms/phrases or categories can be generated and used in substantially the same way. A user can search for rules relating to a particular topic (e.g., customer “Dissatisfaction”) by selecting the topic from the “Target Topic” dropdown box, in which case rules containing “Dissatisfaction” on the “Right Hand

Side” will be shown. A user can also restrict the results to rules matching particular criteria (e.g., minimum lift and/or minimum confidence). In the example shown in FIG. 9, only rules satisfying a minimum confidence of 70 were returned and, if the “Dissatisfaction” topic were selected under “Target Topic,” only the rules “Transfer + Dispute → Dissatisfaction” and “No Payment + On Hold → Dissatisfaction” would be shown among the rules involving the “Dissatisfaction” concept.

[00142] As described above, embodiments of the present invention can be used to mine correlations and causal relationships between predefined topics or categories, between discovered concepts, or both. For the sake of convenience, the mined objects, whether they are instances of predefined topics or instances of discovered concepts, will be referred to as “elements.”

[00143] FIG. 10 is a flowchart illustrating a process executed by the analytics server 45 and the root cause mining module 45b for determining causes of events according to one embodiment of the present invention.

[00144] The set of documents to be analyzed can be the entire set of interactions or any subset of documents defined by the application or selected by the user (e.g., only calls, only chats, all interactions other than sales calls, etc.). To identify frequent sequences within the set of documents, for each document in a given set of documents, elements below a certain confidence value are filtered out in operation 300 and a sequence of element IDs, sorted by ascending element start time, is created in operations 301 and 302. In the case of topics or categories, the start time of the element is the recognition start time of the term contained in the topic. In the case of a discovered concept, for example a cluster of sentences, the element start time is the recognition start time of the first word in the sentence contained in the cluster. The set of sequences can then be condensed by eliminating occurrences of consecutive similar topics in operation 303.

[00145] The resulting set is then mined in operation 304 using an algorithm for mining frequent sequences such as PrefixSpan, (see Pei J, Han J, Mortazavi-Asl B, Wang J, Pinto H, Chen Q, Dayal U, Hsu M-C, *Mining Sequential Patterns by Pattern-Growth: The PrefixSpan Approach*, IEEE Trans. Knowl. Data Eng. 16:1424–1440 (2004)), resulting in a set F of patterns with support greater than a minimum support threshold s .

[00146] The set of patterns F is condensed in operation 305 by deleting sequences with repetitions (even non-consecutive ones). For example, the pattern (a, d, a, b) is not much more interesting than the pattern (a, d, b). In order to reduce the size of the output presented to the end user, the pattern (a, d, a, b) would be deleted. In addition, deleting non-repeating patterns can ensure that association rules derived from the pattern do not contain the same element in both the conditions (LHS) and the consequences (RHS). The pattern (a, d, a, b) can be removed from the set of patterns F without losing information regarding the

relationship between elements {a, d, b} because both patterns (a, d, b) and (d, a, b) remain in the set F after condensing the set F .

[00147] More formally, given a set of frequent patterns F and a pattern p with a repetition of item e , the pattern p' created by removing the repetition e , exists in F :

$$p \in F \Rightarrow \sup(p) > s$$

$$p' \text{ is subsequence of } p \Rightarrow \sup(p') \geq \sup(p) > s \Rightarrow p' \in F$$

where $\sup(p)$ is the support function, which is defined as the number of times pattern p appears in the across all documents.

[00148] Moreover, this procedure can be performed without updating the supports of (a, d, b) and (d, a, b) because every instance of (a, d, a, b) is already counted for the patterns (a, d, b) and (d, a, b).

[00149] After condensing the set of patterns in F in operation 305, association rules for the patterns in F are computed in operation 306.

[00150] FIG. 11 is a flowchart illustrating, in more detail, the operation of generating association rules 306 of the process shown in FIG. 10 according to one embodiment of the present invention. Referring to FIG. 11, according to another aspect of embodiments of the present invention, the association rules can be computed by the analytics server 45 and the topic detection module 45a according to the following method:

[00151] Patterns F that were mined in operation 304 (see FIG. 10) are processed in generating the association rules. For each sequence pattern p in patterns F (400), check the length of pattern p in operation 401. If p only has one element, then the pattern p is discarded and not considered and the process continues with the next pattern p in F (see operation 408). If p has more than one element, then p is split into a first portion p_1 (or “left-hand side” or LHS) and a second portion p_2 (or “right hand side” or RHS) to initialize the first and second portions p_1 and p_2 in operation 402, where p_1 is initialized to include all but the last element of p and p_2 is initialized to include only the last element of p .

[00152] An inference rule from the first portion p_1 to the second portion p_2 ($p_1 \rightarrow p_2$ or LHS $\rightarrow$ RHS) is generated and stored. A confidence of the inference rule is then computed in operation 403, where confidence is computed as:

$$Conf(LHS \Rightarrow RHS) = \frac{Sup(LHS, RHS)}{Sup(LHS)} = \frac{|(LHS, RHS)|}{|(LHS)|}$$

The confidence of a rule represents the probability of having its RHS in the documents, given that they contain its LHS. In other words, given the conditions (LHS), what is the probability to see the consequences (RHS) afterwards?

[00153] The computed confidence is compared with a threshold in operation 404 and, if the confidence is above a given threshold, then the rule is added to the results set in operation 407. To further generate more potential rules from the same pattern, the last element of p_1 is

1 moved to the start of p_2 . If this has not caused the p_1 to be empty (as checked in operation 406), a new inference rule is generated with the modified first and second portions p_1 and p_2 and the confidence of the new inference rule is computed in operation 403.

[00154] If the computed confidence is less than the threshold or if p_1 no longer has any
 5 elements, then generation of rules from the current pattern is stopped. . If there are more patterns in F , as determined in operation 408, then the process continues with the next pattern in operation 401. Otherwise, the rules in the stored results set above the threshold are returned in operation 409.

[00155] In addition to computing the confidence in each of the rules, lift and saliency can
 10 also be computed for the rules, either in operation 403 or after the result set is returned in operation 409. The properties lift and saliency can be computed as follows:

$$Lift(LHS \Rightarrow RHS) = \frac{Conf(LHS \Rightarrow RHS)}{P(RHS)}$$

$$Saliency(pattern) = \sum_{e \in pattern} IDF(e)$$

Where

$$IDF(e) = \log\left(\frac{N}{DF(e)}\right)$$

20 Where N is the total number of patterns in the collection and $DF(e)$ is the number of patterns in which the element e appears. This formulation is very similar to the sentence saliency computation used for the clustering algorithm described above.

[00156] Lift is a measure of probabilistic dependence of the RHS on the LHS. Lower lift indicates that the events (that is, the LHS and RHS) are more independent; for example,
 25 having lift ≤ 1 . Higher lift indicates that they are more dependent on each other; for example, having lift > 1 .

[00157] Saliency is a measure of the amount of information contained in the pattern, given the whole set of patterns. It can be viewed as the relative importance or prominence of the pattern among all of the patterns that appear in the set. Higher saliency indicates that the
 30 pattern is more interesting.

[00158] Rules in the result set can also be sorted by confidence, lift, and saliency.

[00159] FIG. 12 is an illustration of an output of a root cause mining process as shown, for example, in FIGS. 10 and 11, according to one embodiment of the present invention. As shown in FIG. 12, each of the ovals represents a particular detected element or mined object,
 35 where the elements include:

Table 1:

TRT	TRT Issue Buyer
-----	-----------------

1	Pur	Purchase Process
	Pos	Postponed Cancelled Rainout Contingent
	A	Applying Coupon and Fan Codes
	T	Ticket Details
5	Z	Zones and Restricted Events
	S	Sourcing

However, in other embodiments and when processing various other interactions in different contexts (e.g., customer support, sales of different products, etc.), different elements would appear in the patterns. Each path of arrows from left to right represents a different sequence of elements and darker or more strongly bolded arrows represent patterns that are better in terms of higher lift. As such, paths with more strongly bolded arrows indicate likely pathways between events that lead toward events on the right.

[00160] For example, the events “TRT, Pur” have strong connections with the path $Z \rightarrow \text{Pos} \rightarrow A$, suggesting that the sequence of “Zones and Restricted Events”, then “Postponed Cancelled Rainout Contingent”, and “Applying Coupon and Fan Codes” lead to the “Purchase Process” and “TRT Issue Buyer” events. When a user would like to explore ways to increase the number of interactions that result in the “Purchase Process” event, the user can supply a query to the analytics server 45 through the user interface module 45c (see, e.g., FIG. 9) to search for the “purchase process event” in the derived rules, the analytics server 45 can search the data structure illustrated in FIG. 12 to identify a set of events that generally lead to the desired event. As seen in FIG. 12, the above discussed path of $Z \rightarrow \text{Pos} \rightarrow A$ and the path $Z \rightarrow \text{Pos}$ both leads strongly to the event TRT Pur and, if these paths satisfied the threshold conditions set by the user (e.g., minimum confidence levels), then these paths would be returned and displayed to the user, as seen in, for example, FIG. 9.

[00161] Embodiments of the invention can be practiced as methods or systems. Computer devices or systems including, for example, a microprocessor, memory, a network communications device, and a mass storage device can be used to execute the processes described above in an automated or semi-automated fashion. In other words, the above processes can be coded as computer executable code and processed by the computer device or system.

[00162] It should also be appreciated from the above that various structures and functions described herein may be incorporated into a variety of apparatus. In some embodiments, hardware components such as processors, controllers, and/or logic may be used to implement the described components or circuits. In some embodiments, code such as software or firmware executing on one or more processing devices may be used to implement one or more of the described operations or components.

[00163] As would be understood by one of ordinary skill in the art, the processes described herein and as illustrated in the flowcharts in the figures may be implemented by instructions

1 stored in computer memory to control a computer processor to perform the described
functions. In addition, steps and operations shown in the flowchart do not need to be executed
in the order shown and person of ordinary skill in the art at the time the invention was made
the order of the steps and operations performed may vary without deviating from or
5 substantially altering the underlying technique.

[00164] While the present invention has been described in connection with certain
exemplary embodiments, it is to be understood that the invention is not limited to the
disclosed embodiments, but, on the contrary, is intended to cover various modifications and
equivalent arrangements included within the spirit and scope of the appended claims, and
10 equivalents thereof.

1 WHAT IS CLAIMED IS:

1. A method for identifying concepts in a plurality of interactions, the method comprising:

filtering, on a processor, the interactions based on intervals;

5 creating, on the processor, a plurality of sentences from the filtered interactions;

computing, on the processor, a saliency of each the sentences;

pruning away, on the processor, sentences with low saliency for generating a set of informative sentences;

10 clustering, on the processor, the sentences of the set of informative sentences for generating a plurality of sentence clusters, each of the clusters corresponding to a concept of the concepts;

computing, on the processor, a saliency of each of the clusters; and

naming, on the processor, each of the clusters.

15 2. The method of claim 1, wherein the interactions comprise an output of a voice recognition system.

3. The method of claim 2, further comprising filtering the output of the voice recognition system based on word confidence.

20 4. The method of claim 2, wherein the voice recognition system is a large-vocabulary continuous speech recognition system.

5. The method of claim 1, wherein the intervals are time intervals.

25 6. The method of claim 1, wherein the clustering the sentences comprises: selecting a plurality of template sentences from the set of informative sentences, each of the template sentences corresponding to one of the clusters;

30 computing a similarity of each sentence in the set of informative sentences to the template sentences;

assigning each of the sentences of the set of informative sentences to a cluster of the clusters in accordance with the computed similarities with the template sentences; and

removing clusters having fewer than a threshold number of assigned sentences.

35 7. The method of claim 6, wherein the clustering the sentences further comprises, iteratively, selecting additional template sentences from sentences that were not assigned to a cluster and repeating the selecting the plurality of template sentences, computing the similarity of each sentence to the template sentences; the assigning the each of the sentences

1 to the clusters; and removing clusters having fewer than the threshold number of assigned sentences until all sentences are assigned or until an iteration limit is reached.

5 8. The method of claim 1, wherein the naming each of the clusters comprises, for each cluster:

computing a term frequency-inverse document frequency for each sentence in the cluster and

naming the cluster with a sentence of the cluster having the highest term frequency-inverse document frequency.

10 9. A method for exploring concepts automatically identified in a plurality of interactions, the method comprising:

receiving, on a processor, a query comprising a concept;

retrieving, by the processor, a cluster containing the concept; and

15 displaying the retrieved cluster,

wherein the concepts are automatically identified by:

filtering, on the processor, the interactions based on intervals;

creating, on the processor, a plurality of sentences from the filtered interactions;

20 computing, on the processor, a saliency of each the sentences;

pruning away, on the processor, sentences with low saliency for generating a set of informative sentences; and

25 clustering, on the processor, the sentences of the set of informative sentences for generating a plurality of sentence clusters, each of the clusters corresponding to a concept of the concepts.

10. A computer system comprising:

a processor; and

30 a memory, wherein the memory has stored thereon instructions that, when executed by the processor, causes the processor to identify concepts in a plurality of interactions by:

filtering the interactions based on intervals;

creating a plurality of sentences from the filtered interactions;

computing a saliency of each the sentences;

35 pruning away sentences with low saliency for generating a set of informative sentences;

clustering the sentences of the set of informative sentences for generating a plurality of sentence clusters, each of the clusters corresponding to a concept of the concepts;

computing a saliency of each of the clusters; and

1 naming each of the clusters.

11. The system of claim 10, wherein the interactions comprise an output of a voice
recognition system.

5 12. The system of claim 11, wherein the instructions further comprise instructions
for filtering the output of the voice recognition system based on word confidence.

10 13. The system of claim 11, wherein the voice recognition system is a large-
vocabulary continuous speech recognition system.

14. The system of claim 10, wherein the intervals are time intervals.

15 15. The system of claim 10, wherein the clustering the sentences comprises:
selecting a plurality of template sentences from the set of informative sentences, each
of the template sentences corresponding to one of the clusters;
computing a similarity of each sentence in the set of informative sentences to the
template sentences;
assigning each of the sentences of the set of informative sentences to a cluster of the
20 clusters in accordance with the computed similarities with the template sentences; and
removing clusters having fewer than a threshold number of assigned sentences.

25 16. The system of claim 15, wherein the clustering the sentences further
comprises, iteratively, selecting additional template sentences from sentences that were not
assigned to a cluster and repeating the selecting the plurality of template sentences,
computing the similarity of each sentence to the template sentences; the assigning the each of
the sentences to the clusters; and removing clusters having fewer than the threshold number
of assigned sentences until all sentences are assigned or until an iteration limit is reached.

30 17. The system of claim 10, wherein the naming each of the clusters comprises,
for each cluster:

computing a term frequency-inverse document frequency for each sentence in the
cluster and

35 naming the cluster with a sentence of the cluster having the highest term frequency-
inverse document frequency.

18. A computer system comprising:
a processor; and

1 a memory, wherein the memory has stored thereon instructions that, when executed
by the processor, causes the processor to respond to requests for exploration of concepts by:
receiving a query comprising a concept;
retrieving a cluster containing the concept; and
5 displaying the retrieved cluster,
wherein the concepts are automatically identified by:
filtering the interactions based on intervals;
creating a plurality of sentences from the filtered interactions;
computing a saliency of each the sentences;
10 pruning away sentences with low saliency for generating a set of informative
sentences; and
clustering the sentences of the set of informative sentences for generating a
plurality of sentence clusters, each of the clusters corresponding to a concept of the concepts.

15 19. A method for determining a cause of events detected in a plurality of
interactions, the method comprising:
identifying, on a processor, a plurality of elements in the interactions;
detecting, on the processor, a plurality of sequences of elements in the interactions;
mining, on the processor, the plurality of sequences for generating a set of supported
20 patterns;
computing, on the processor, association rules from the set of supported patterns; and
returning the computed association rules.

25 20. The method of claim 19, wherein the elements comprise defined topics and
identified concepts.

21. The method of claim 20, wherein the identified concepts are derived by:
filtering, on the processor, the interactions based on intervals;
creating, on the processor, a plurality of sentences from the filtered interactions;
30 computing, on the processor, a saliency of each the sentences;
pruning away, on the processor, sentences with low saliency for generating a set of
informative sentences;
clustering, on the processor, the sentences of the set of informative sentences for
generating a plurality of sentence clusters, each of the clusters corresponding to a concept of
35 the concepts;
computing, on the processor, a saliency of each of the clusters; and
naming, on the processor, each of the clusters.

1 22. The method of claim 19, wherein the detecting the sequences of elements comprises:

 sorting the identified elements in each interaction by timestamp within the interaction for generating the plurality of sequences; and

5 condensing each of the sequences for removing repeated consecutive elements.

 23. The method of claim 19, further comprising condensing, on the processor, the set of supported patterns for removing repeated elements in the sequences.

10 24. The method of claim 19, wherein the computing the association rules from the set of condensed supported patterns comprises:

 computing an association rule for each sequence pattern in the condensed supported patterns, the computing the association rule comprising, if the sequence pattern comprises more than one element:

15 dividing the sequence pattern into a first portion and a second portion, the second portion comprising the last element in the sequence pattern;

 generating a proposed association rule, the proposed association rule being a logical implication from the first portion to the second portion;

 computing a confidence of the proposed association rule;

20 if the computed confidence is greater than a threshold confidence level, storing the proposed association rule and moving the last element in the first portion to the second portion and iteratively repeating generating the proposed association rule and computing the confidence; and

25 if the computed confidence is less than a threshold confidence level, ending analysis of the sequence pattern and returning a previously stored proposed association rule as the association rule or returning no association rule if no such proposed association rules were stored.

30 25. The method of claim 24, further comprising computing a lift and a saliency of the association rule.

 25. The method of claim 24, further comprising computing a lift and a saliency of the proposed association rule.

35 26. A method for determining a root cause of an event detected in a plurality of interactions, the method comprising:

 receiving, on a processor, a query for the root cause of the event;

1 searching, on the processor, a plurality of association rules, each of the association rules comprising one or more first portion elements and one or more second portion elements, each of the association rules being a logical implication from the first portion to the second portion; and

5 returning, from the processor, one or more association rules matching the query, the second portion elements of each of the matching association rules comprising the event.

27. The method of claim 26, wherein each of the association rules has a corresponding confidence value,

10 wherein the query comprises a confidence threshold, and
 wherein the returning the one or more association rules matching the query further comprises returning association rules matching the query that have confidence values exceeding the confidence threshold.

15 28. The method of claim 26, wherein each of the association rules has a corresponding lift value and a corresponding saliency value,
 wherein the query comprises at least one of a lift threshold and a saliency threshold, and

20 wherein the returning the one or more association rules matching the query further comprises returning association rules matching the query that have lift values or saliency values exceeding the lift threshold or the saliency threshold.

29. A computer system comprising:
 a processor; and
25 a memory, wherein the memory has stored thereon instructions that, when executed by the processor, causes the processor to compute association rules between events detected in a plurality of interactions by:

 identifying a plurality of elements in the interactions;
 detecting a plurality of sequences of elements in the interactions;
30 mining the plurality of sequences for generating a set of supported patterns;
 computing association rules from the set of supported patterns; and
 returning the computed association rules.

30. The computer system of claim 29, wherein the elements comprise defined
35 topics and identified concepts.

31. The computer system of claim 30, wherein the identified concepts are derived by:

1 filtering, on the processor, the interactions based on intervals;
 creating, on the processor, a plurality of sentences from the filtered
interactions;
 computing, on the processor, a saliency of each the sentences;
5 pruning away, on the processor, sentences with low saliency for generating a
set of informative sentences;
 clustering, on the processor, the sentences of the set of informative sentences
for generating a plurality of sentence clusters, each of the clusters corresponding to a concept
of the concepts;
10 computing, on the processor, a saliency of each of the clusters; and
 naming, on the processor, each of the clusters.

32. The computer system of claim 29, wherein the processor is further configured
to compute association rules between events detected in a plurality of interactions by:
15 sorting the identified elements in each interaction by timestamp within the interaction
for generating the plurality of sequences; and
 condensing each of the sequences for removing repeated consecutive elements.

33. The computer system of claim 29, wherein the processor is further configured
20 to compute association rules between events detected in a plurality of interactions by
condensing the set of supported patterns for removing repeated elements in the sequences.

34. The computer system of claim 29, wherein the computing the association rules
from the set of condensed supported patterns comprises:

25 computing an association rule for each sequence pattern in the condensed supported
patterns, the computing the association rule comprising, if the sequence pattern comprises
more than one element:

 dividing the sequence pattern into a first portion and a second portion, the
second portion comprising the last element in the sequence pattern;

30 generating a proposed association rule, the proposed association rule being a
logical implication from the first portion to the second portion;

 computing a confidence of the proposed association rule;

 if the computed confidence is greater than a threshold confidence level, storing
the proposed association rule and moving the last element in the first portion to the second
35 portion and iteratively repeating generating the proposed association rule and computing the
confidence; and

 if the computed confidence is less than a threshold confidence level, ending
analysis of the sequence pattern and returning a previously stored proposed association rule

1 as the association rule or returning no association rule if no such proposed association rules
were stored.

5 35. The computer system of claim 34, wherein the computing the association rule
further comprises computing a lift and a saliency of the association rule.

36. The computer system of claim 34, wherein the computing the association rule
further comprises computing a lift and a saliency of the proposed association rule.

10 37. A computer system comprising:
a processor; and
a memory, wherein the memory has stored thereon instructions that, when executed
by the processor, causes the processor to respond to a query for the root cause of an event by:
receiving the query for the root cause of the event;
15 searching a plurality of association rules, each of the association rules
comprising one or more first portion elements and one or more second portion elements, each
of the association rules being a logical implication from the first portion to the second
portion; and
returning one or more association rules matching the query, the second portion
20 elements of each of the matching association rules comprising the event.

38. A system substantially as hereinbefore described with reference to the
accompanying drawings.

25 39. A method substantially as hereinbefore described with reference to the
accompanying drawings.




FIG. 1




FIG. 2A




FIG. 2B




FIG. 2C




FIG. 2D




FIG. 2E

1 ↗

New Customer		42.0%		84
Emergency		29.5%		59
Identification		19.5%		39
Billing		6.0%		12
Payment Inquiry		4.0%		8

FIG. 3




FIG. 4




FIG. 5

9/15




FIG. 6

10/15




FIG. 7

11/15




FIG. 8

5 Dataset Filters <<
(Standard Search Parameters)

Items on Report

Target Topic (Optional)

None Selected ▼

Minimum Lift

Minimum Confidence

70

[Run Report](#)

Left Hand Side	Right Hand Side	Support %	Confidence ▼	Lift %	Saliency %
Transfer + Dispute	→ Dissatisfaction	70.0	90.0	15	15.0
No Payment + On Hold	→ Dissatisfaction	70.0	85.5	2.0	15.5
Collection Closed	→ Urgency + Courtesy	65.0	85.0	2.0	10.0
Positive Feedback	→ Buying + Sympathy	65.0	77.5	1.0	8.0
Infringement + Claim	→ Fraud	60.0	75.0	1.0	8.0

FIG. 9

13/15




FIG. 10




FIG. 11




FIG. 12

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2014/048089**A. CLASSIFICATION OF SUBJECT MATTER****G06Q 30/02(2012.01)i, G06Q 30/06(2012.01)i, G06F 17/30(2006.01)i, G06F 17/28(2006.01)i**

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 30/02; G10L 15/04; G06F 11/07; G06F 17/30; G06F 17/27; G10L 11/02; G06F 15/18; G06Q 30/06; G06F 17/28

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

Korean utility models and applications for utility models

Japanese utility models and applications for utility models

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

eKOMPASS(KIPO internal) & Keywords: identifying concepts, interaction, saliency, clustering, naming

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 6185527 B1 (DRAGUTIN PETKOVIC et al.) 06 February 2001 See columns 6-9, claims 1-13 and figures 1-13.	1-25(1), 29-36
Y		26-28, 37
Y	US 2012-0096316 A1 (JOHN M. SUIT) 19 April 2012 See abstract, claim 12 and figures 1-10.	26-28, 37
A	US 2010-0161612 A1 (CHIEN-CHIN CHEN et al.) 24 June 2010 See abstract, claim 1 and figures 1-5.	1-37
A	US 2012-0166180 A1 (LAWRENCE AU) 28 June 2012 See abstract, claims 1-8 and figures 1-15.	1-37
A	US 2011-0295903 A1 (CHAOMEI CHEN) 01 December 2011 See abstract, claims 1-10 and figures 1-20.	1-37



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

24 November 2014 (24.11.2014)

Date of mailing of the international search report

25 November 2014 (25.11.2014)

Name and mailing address of the ISA/KR


International Application Division
Korean Intellectual Property Office
189 Cheongsu-ro, Seo-gu, Daejeon Metropolitan City, 302-701,
Republic of Korea

Facsimile No. +82-42-472-7140

Authorized officer

KIM, Sung Gon

Telephone No. +82-42-481-8746



INTERNATIONAL SEARCH REPORTInternational application No.
PCT/US2014/048089**Box No. II Observations where certain claims were found unsearchable (Continuation of item 2 of first sheet)**

This international search report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☐ Claims Nos.:
because they relate to subject matter not required to be searched by this Authority, namely:
2. ☒ Claims Nos.: 38-39
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:
Claims 38-39 are contrary to PCT Rule 6.2(a) because they rely on reference to the drawings instead of definite structure expressly stated in the claims.
3. ☐ Claims Nos.:
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box No. III Observations where unity of invention is lacking (Continuation of item 3 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

1. ☐ As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.
2. ☐ As all searchable claims could be searched without effort justifying an additional fees, this Authority did not invite payment of any additional fees.
3. ☐ As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:
4. ☐ No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:

Remark on Protest

- ☐ The additional search fees were accompanied by the applicant's protest and, where applicable, the payment of a protest fee.
- ☐ The additional search fees were accompanied by the applicant's protest but the applicable protest fee was not paid within the time limit specified in the invitation.
- ☐ No protest accompanied the payment of additional search fees.

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2014/048089

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 6185527 B1	06/02/2001	CN 1261181 A0 JP 2000-259168 A JP 3531729 B2 KR 10-0380947 B1 TW 469422 B	26/07/2000 22/09/2000 31/05/2004 21/04/2003 21/12/2001
US 2012-0096316 A1	19/04/2012	US 2012-0221898 A1 US 8156378 B1 US 8656219 B2	30/08/2012 10/04/2012 18/02/2014
US 2010-0161612 A1	24/06/2010	TW 201025035 A US 8140543 B2	01/07/2010 20/03/2012
US 2012-0166180 A1	28/06/2012	None	
US 2011-0295903 A1	01/12/2011	US 2014-0195539 A1 US 8566360 B2	10/07/2014 22/10/2013