US 20160019470A1

(19) **United States**

(12) **Patent Application Publication**

LIGHTNER et al.

(10) **Pub. No.: US 2016/0019470 A1**

(43) **Pub. Date: Jan. 21, 2016**

(54) **EVENT DETECTION THROUGH TEXT ANALYSIS USING TRAINED EVENT TEMPLATE MODELS**

(71) Applicant: **QBASE, LLC**, Reston, VA (US)

(72) Inventors: **Scott LIGHTNER**, Leesburg, VA (US); **Rakesh DAVE**, Dayton, OH (US); **Sanjay BODDHU**, Dayton, OH (US); **Robert FLAGG**, Portland, ME (US)

(73) Assignee: **QBASE, LLC**, Reston, VA (US)

(21) Appl. No.: **14/867,138**

(22) Filed: **Sep. 28, 2015**

Related U.S. Application Data

(63) Continuation of application No. 14/558,300, filed on Dec. 2, 2014, now Pat. No. 9,177,254.

(60) Provisional application No. 61/910,809, filed on Dec. 2, 2013.

Publication Classification

(51) **Int. Cl.**

G06N 99/00 (2006.01)

G06N 5/04 (2006.01)

(52) **U.S. Cl.**

CPC ***G06N 99/005*** (2013.01); ***G06N 5/04*** (2013.01)

(57) **ABSTRACT**

A system and method for detecting events based on input data from a plurality of sources. The system may receive input from a plurality of sources containing information about possible events. A method for event detection involves pre-processing and normalizing a data input from a plurality of sources, extracting and disambiguating events and entities, associate event and entities, correlate events and entities associated from a data input to results from a different data sources to determine if an event has occurred, and store the detected events in a data storage.

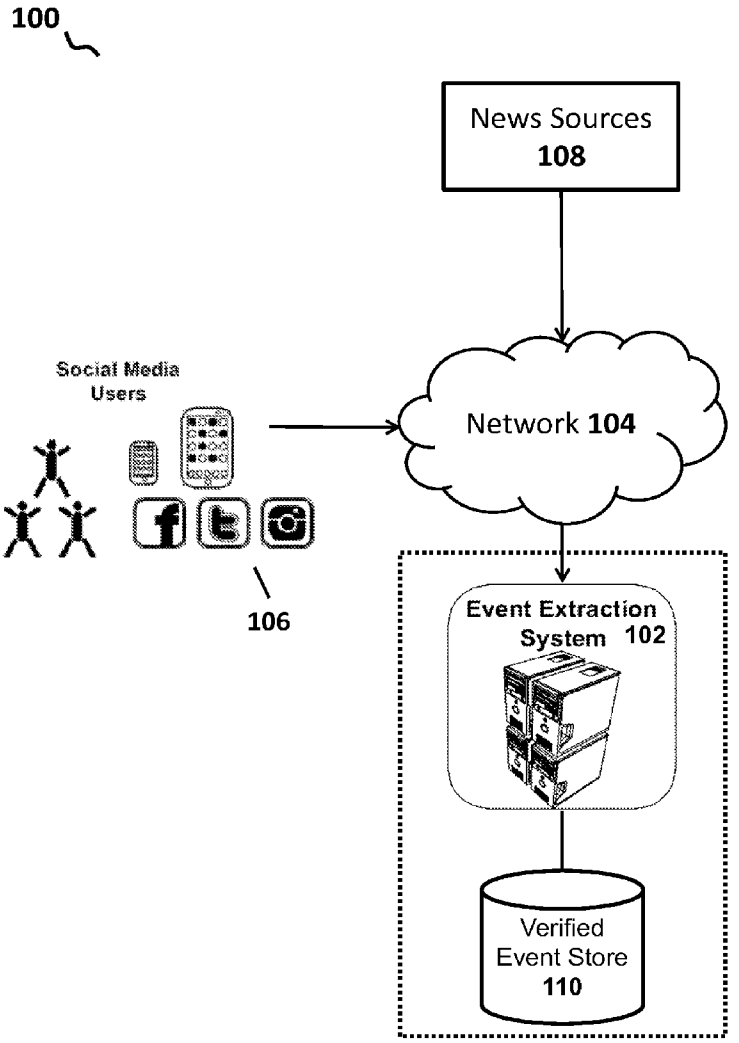

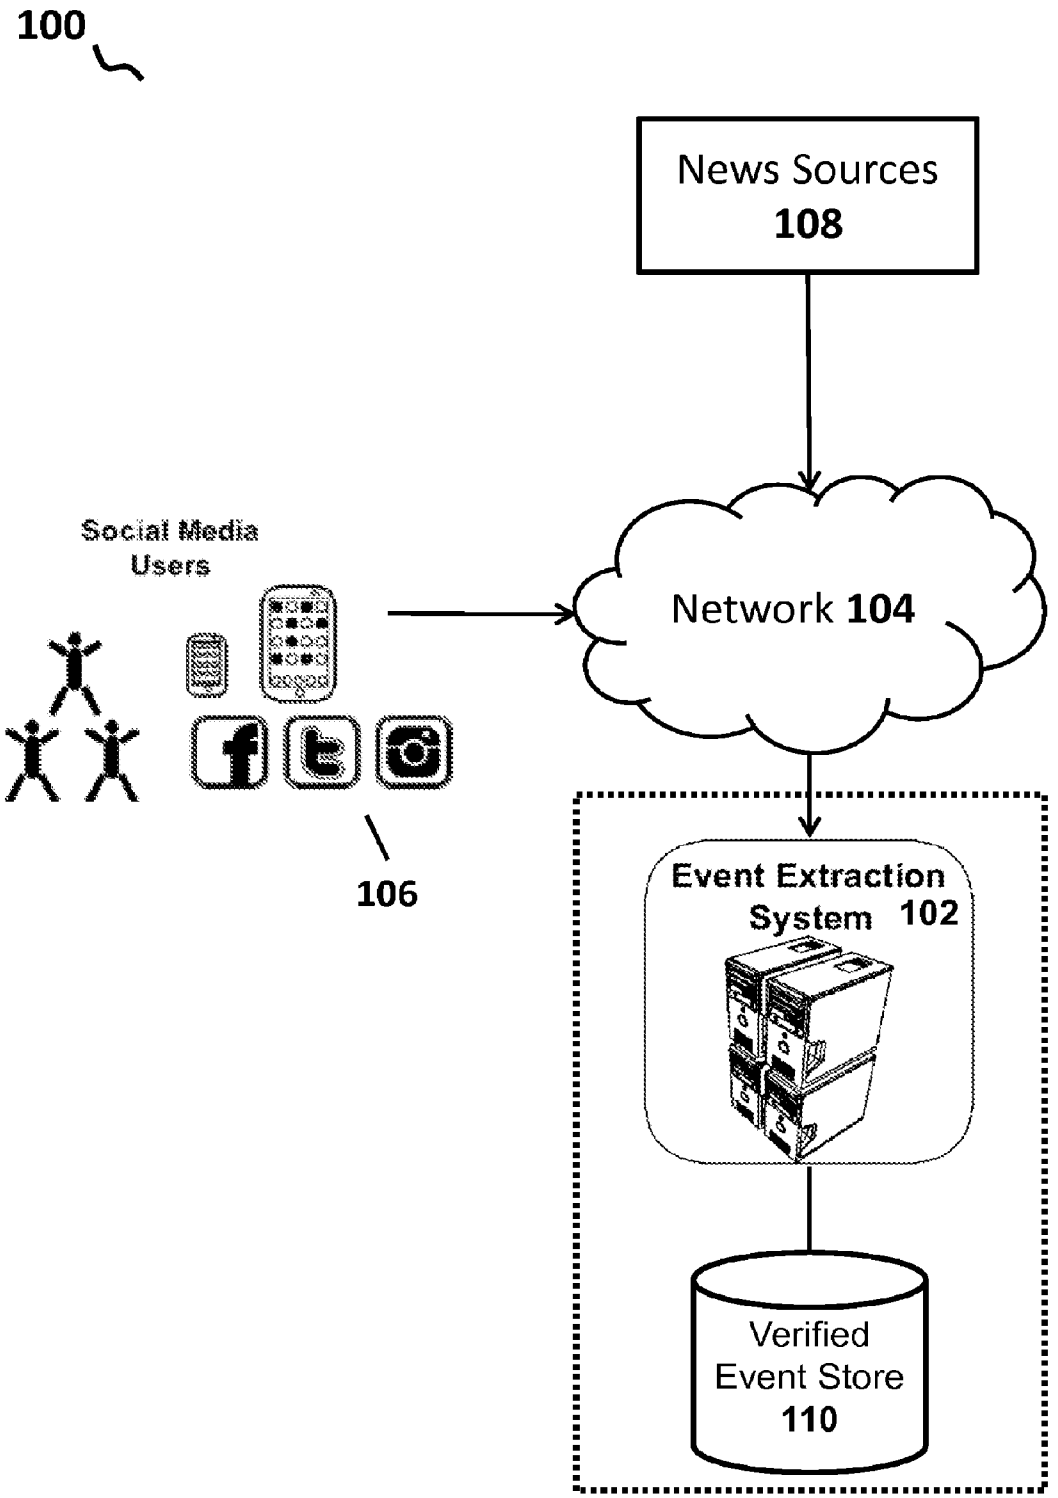
100
News Sources
108
Social Media
Users
Network 104
106
Event Extraction
System 102
Verified
Event Store
110

FIG. 1

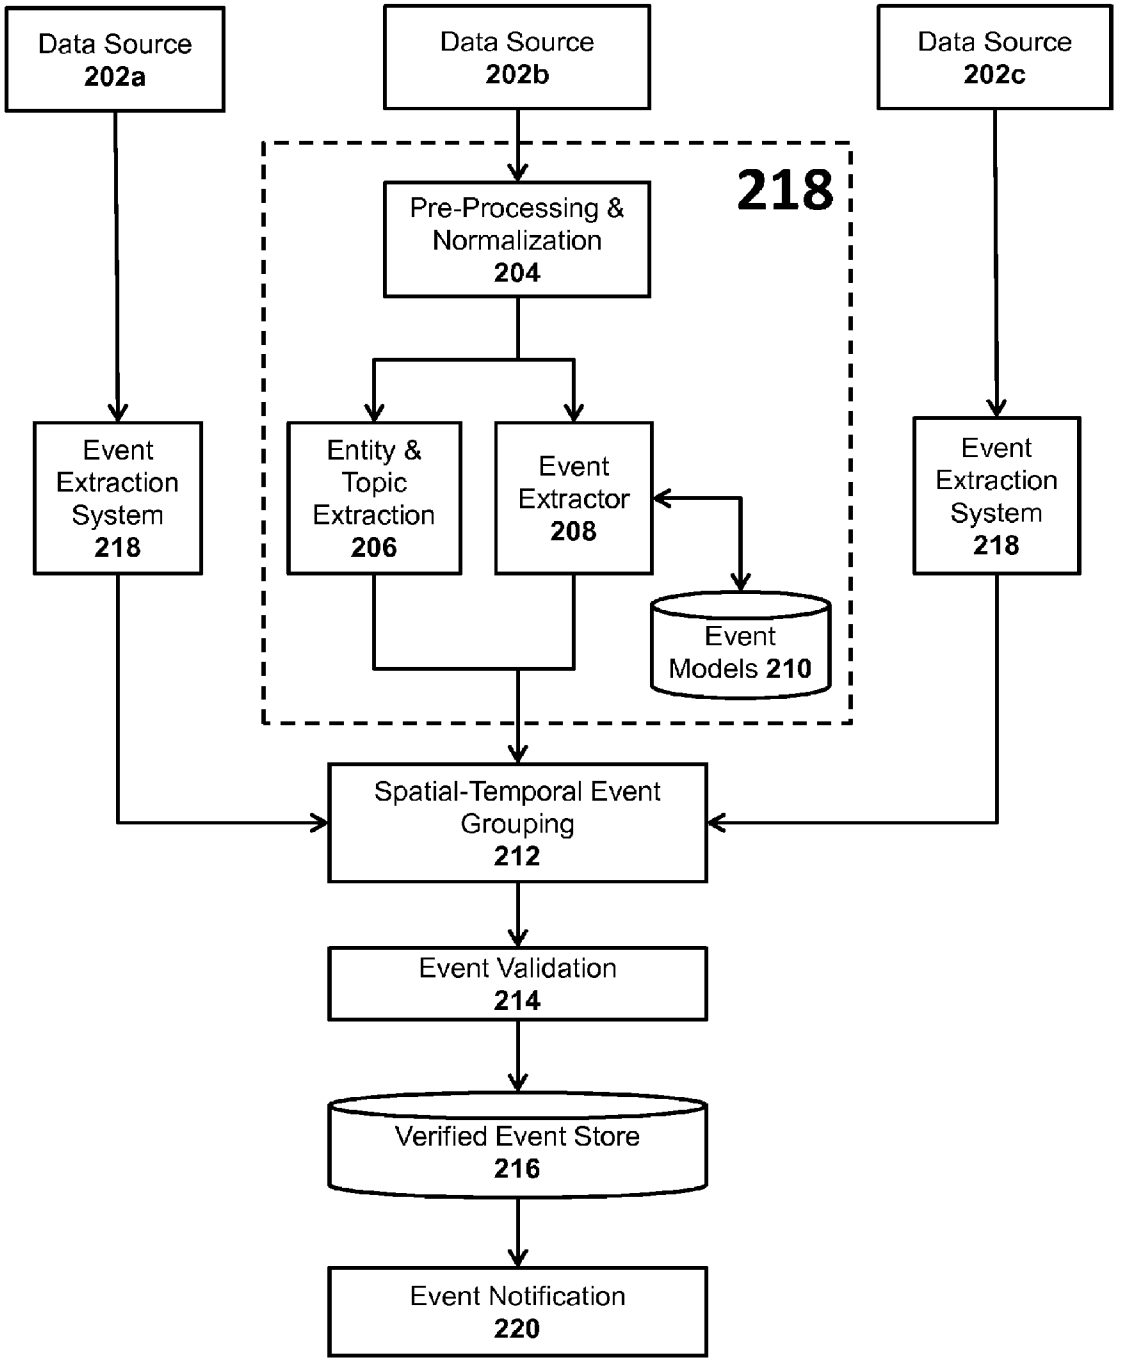
Data Source
202a
Data Source
202b
Data Source
202c
218
Pre-Processing &
Normalization
204
Event
Extraction
System
218
Entity &
Topic
Extraction
206
Event
Extractor
208
Event
Models 210
Event
Extraction
System
218
Spatial-Temporal Event
Grouping
212
Event Validation
214
Verified Event Store
216
Event Notification
220

FIG. 2

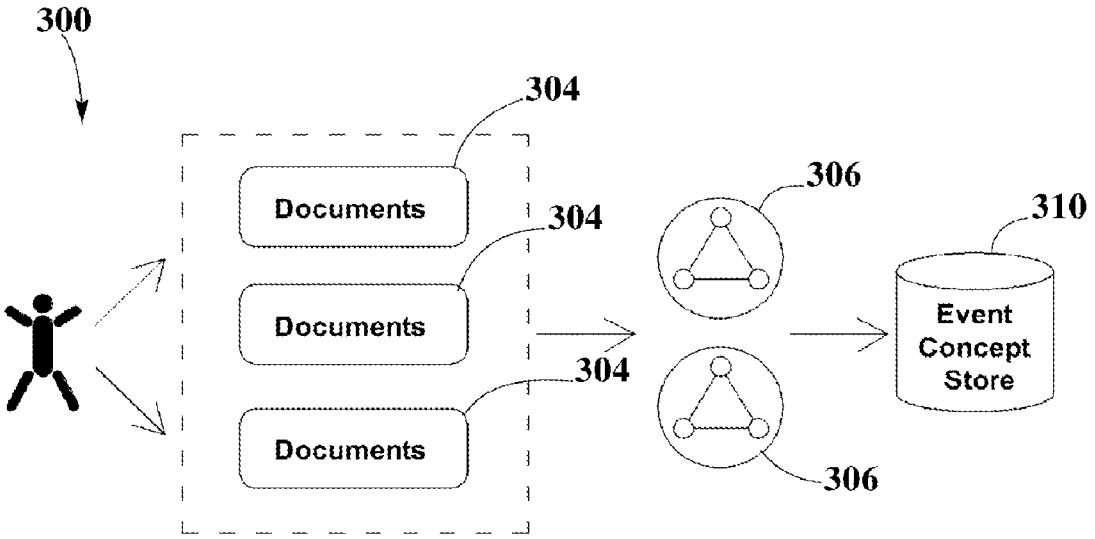
300
304
Documents
304
Documents
304
Documents
306
306
310
Event
Concept
Store

FIG. 3

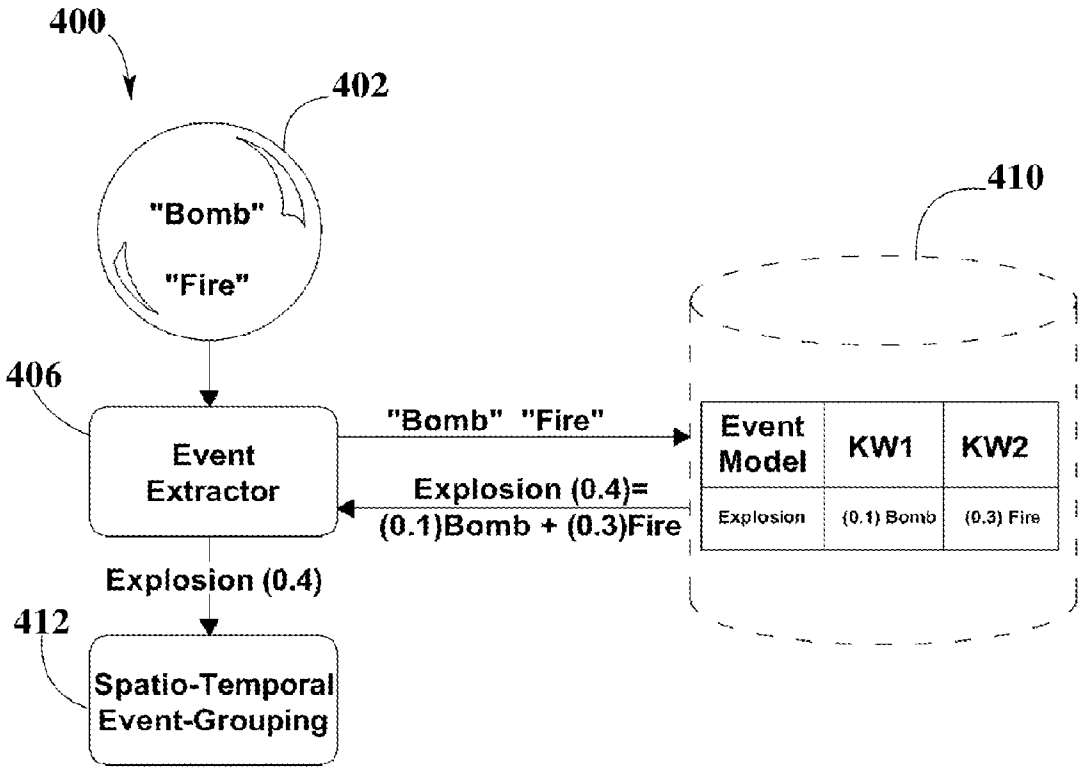
400
402
"Bomb"
"Fire"
410
406
Event Extractor
"Bomb" "Fire"
Explosion (0.4)= (0.1)Bomb + (0.3)Fire
Event Model
KW1
KW2
Explosion
(0.1) Bomb
(0.3) Fire
Explosion (0.4)
412
Spatio-Temporal Event-Grouping

FIG. 4

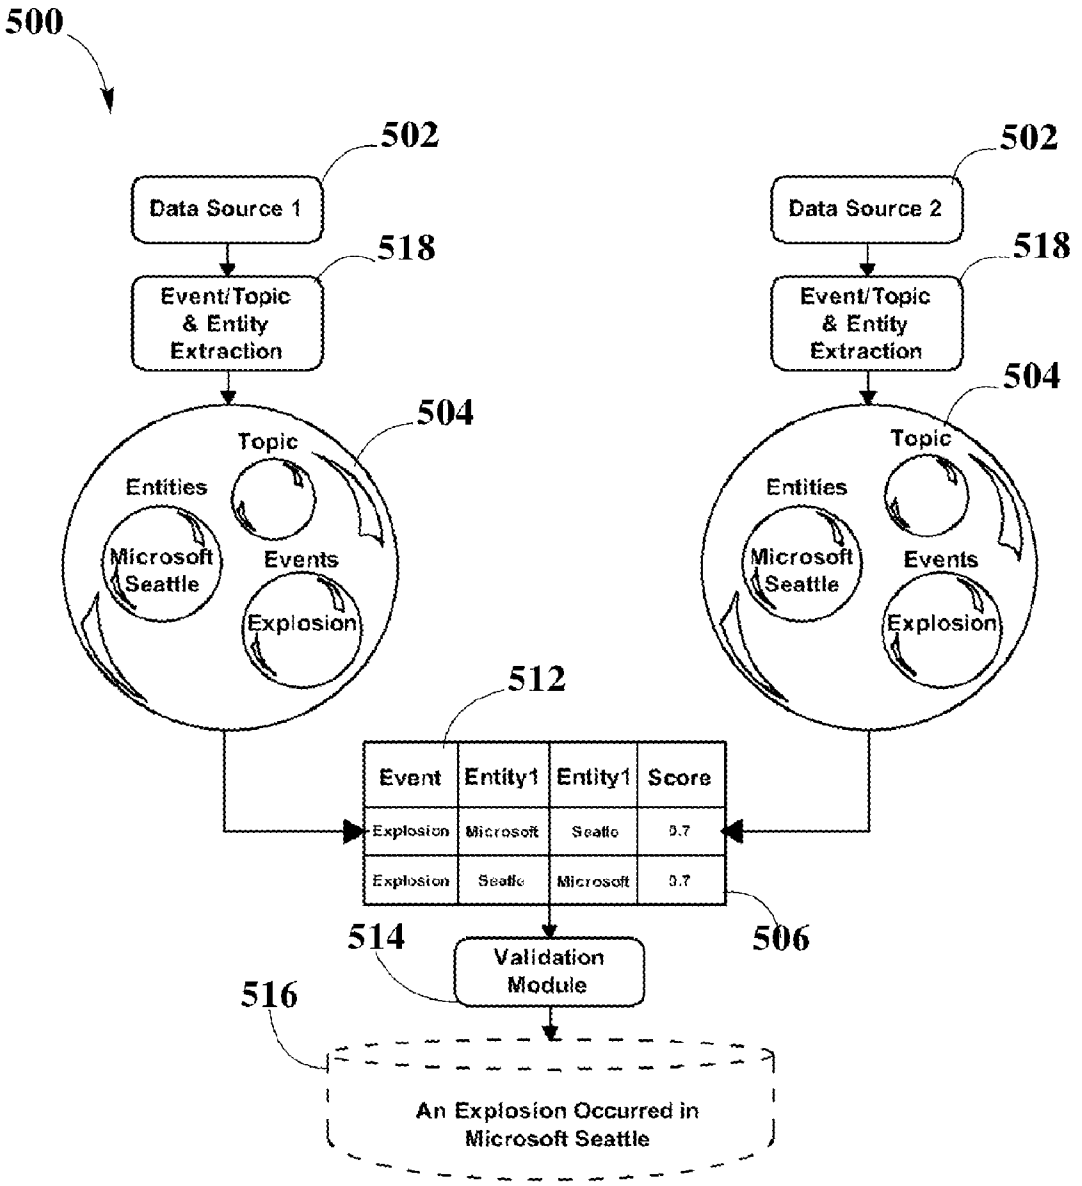
500
502
Data Source 1
518
Event/Topic & Entity Extraction
504
Topic
Entities
Microsoft Seattle
Events
Explosion
502
Data Source 2
518
Event/Topic & Entity Extraction
504
Topic
Entities
Microsoft Seattle
Events
Explosion
512
Event
Entity1
Entity1
Score
Explosion
Microsoft
Seattle
0.7
Explosion
Seattle
Microsoft
0.7
506
514
Validation Module
516
An Explosion Occurred in Microsoft Seattle

FIG. 5

EVENT DETECTION THROUGH TEXT ANALYSIS USING TRAINED EVENT TEMPLATE MODELS

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application is a continuation of U.S. patent application Ser. No. 14/558,330, filed on Dec. 2, 2014, which is a non-provisional patent application that claims the benefit of U.S. Provisional Application Ser. No. 61/910,809, entitled "Event Detection Through Text Analysis Using Trained Event Template Models," filed Dec. 2, 2013, all of which are hereby incorporated by reference in their entirety.

[0002] This application is related U.S. patent application Ser. No. 14/558,342, entitled "Event Detection Through Text Analysis Using Dynamic Self Evolving/Learning Module," filed Dec. 2, 2014, and U.S. patent application Ser. No. 14/558,254, entitled "Design And Implementation Of Clustered In-Memory Database," filed Dec. 2, 2014, each of which are hereby incorporated by reference in their entirety.

TECHNICAL FIELD

[0003] The present disclosure generally relates to information data mining from various media sources, and more specifically to event detection, extraction and validation from media sources.

BACKGROUND

[0004] The internet provides several sources of information which may be exploited. Internet news feeds and websites that allow users to interact with one another have exploded in popularity in the last few years. News feed channels such as CNN®, social networking websites sites such as Facebook® or LinkedIn®, and microblogging websites such as Twitter® enjoy widespread use. Millions of users post messages, images and videos on such websites on a daily, even hourly basis. Often, information gathered from these sources may refer to events taking place in real time. Such publicly accessible media may serve as a rich mine of information that may be used in different applications. For example, consider a scenario where a wide area emergency such as an earthquake or a flood has occurred and conventional emergency service lines are stressed beyond capacity; in this case users may turn to social media in order to request assistance. Another example of an event taking place in real time may be news feed reporting on civilians trapped under a building.

[0005] The high proliferation of information generated by media sources makes proper identification of events troublesome. Media data may contain ambiguous features which may hinder the ability of associating events with specific names, places or organizations. For example, a news feed may refer to a "Paris kidnapping"; however, in general, Paris may refer to a city in France, a city in Texas, or it may even refer to a person.

[0006] Thus a need exists for a method of detecting, extracting and validating events from media sources and effectively associate them with independent entities.

SUMMARY

[0007] A system and method for detecting events based on input data from a plurality of sources such as, social media, news feeds, and/or a corpus of documents. The system may receive input from a plurality of sources in the way of natural language unstructured text containing information about real time events. The system may use natural language processing techniques in order to separate individual entities and keywords. The process may then proceed with an entity disambiguation step and identify specific entities the source may be referring. The system may then identify independent events and associate them with the specific entities identified in the same source. The process may then validate events based on overlapping and co-occurrence of events and entities from other data sources. The system allows for the detection of events happening, and their proper association to disambiguated entities through text analysis of different sources.

[0008] In one embodiment, a system for detecting and extracting events comprises an event concept store storing one or more event models, wherein an event model corresponds to an event candidate and comprises a threshold event score and a set of one or more features having a corresponding weight; an entity and topic extraction processor configured to extract a set of entities and a set of topics from a data stream and then disambiguate each topic and each entity; an event extraction processor configured to identify each of the features of each event model that occur in the data stream, calculate an event score for one or more event candidates having an identified feature using the corresponding event model, and then extract the event candidate when an event score satisfies the threshold event score of the event model; and a spatial-temporal event processor configured to associate each topic and entity extracted from each of the data streams with each of the event candidates extracted from each of the data streams to form a spatial-temporal event grouping comprising one or more records, wherein a record is the event candidate and the associated topic or entity of a data stream.

[0009] In another embodiment, a computer-implemented method of event extraction and detection comprises normalizing, by a computer, one or more source items received over a network from a plurality of sources into one or more data streams adequate for computer-automated processing, wherein each data stream is drawn from a corresponding source item; extracting, by the computer, one or more entities from a data stream when entities are identified by the computer; extracting, by the computer, one or more topics from the data stream when topics are identified by the computer; disambiguating, by the computer, each of the entities and each of the topics; identifying, by the computer, one or more features in the data stream matching a feature in an event model, wherein an event model corresponds to an event candidate, and comprises a threshold event score and a set of one or more features having a corresponding weight; calculating, by the computer, an event score for each event model based on the weights assigned to the identified features; extracting, by the computer, each event candidate corresponding to an event model having an event score threshold satisfied by the event score; associating, by the computer, each event candidate with each topic or entity into a spatial-temporal grouping, wherein each event candidate and the associated topic or entity extracted from a source item are a record in the spatial-temporal grouping.

[0010] In one embodiment, a computer-implemented method comprises receiving, by a computer, a data stream comprising data representing text strings from a server of a data source; identifying, by the computer, in the data stream one or more features matching a model feature in an event model stored in non-transitory machine-readable storage of an event concept store, wherein at least one feature in the one

or more features is an event candidate corresponding to the event model, and wherein at least one feature in the one or more features is an entity; assigning, by the computer, a weight to each respective entity according to the event model corresponding to the event candidate; calculating, by the computer, an event probability score based on one or more weights respectively assigned to the one or more entities; associating, by the computer, the event candidate with each respective entity into a first record of a spatial-temporal grouping, in response to the computing device determining that the event probability score satisfies an event likelihood threshold score of the event model; and storing, by the computer, the first record containing the event candidate, each respective entity associated with the event candidate, and the event probability score into a non-transitory machine-readable spatial-temporal grouping memory.

[0011] In another embodiments, a system for detecting and extracting events, the system comprising an event concept store comprising non-transitory machine-readable memory storing one or more event models, wherein an event model corresponds to an event candidate and comprises a threshold event score and a set of one or more features having a corresponding weight; an entity and topic extraction processor configured to extract a set of entities and a set of topics from a data stream and then disambiguate each topic and each entity; an event extraction processor configured to identify each of the features of each event model that occur in the data stream, calculate an event score for one or more event candidates having an identified feature using the corresponding event model, and then extract the event candidate when an event score satisfies the threshold event score of the event model; and a spatial-temporal event processor configured to associate each topic and entity extracted from each of the data streams with each of the event candidates extracted from each of the data streams to form a spatial-temporal event grouping comprising one or more records, wherein a record is the event candidate and the associated topic or entity of a data stream.

[0012] Additional features and advantages of an embodiment will be set forth in the description which follows, and in part will be apparent from the description. The objectives and other advantages of the invention will be realized and attained by the structure particularly pointed out in the exemplary embodiments in the written description and claims hereof as well as the appended drawings.

[0013] It is to be understood that both the foregoing general description and the following detailed description are exemplary and explanatory and are intended to provide further explanation of the invention as claimed.

BRIEF DESCRIPTION OF THE DRAWINGS

[0014] The present disclosure can be better understood by referring to the following figures. The components in the figures are not necessarily to scale, emphasis instead being placed upon illustrating the principles of the disclosure. In the figures, reference numerals designate corresponding parts throughout the different views.

[0015] FIG. **1** shows components of an event extraction system, according to an embodiment.

[0016] FIG. **2** is a flow diagram illustrating a process by which events and entities from different sources are extracted, validated and stored, according to an embodiment.

[0017] FIG. **3** shows a process of how an event concept store may be developed, according to an exemplary process embodiment.

[0018] FIG. **4** is an of a detailed event extraction process using an event concept store, according to an embodiment.

[0019] FIG. **5** is an of an event and entity extraction and validation using different data sources.

DEFINITIONS

[0020] As used herein, the following terms have the following definitions:

[0021] "Entity Extraction" refers to information processing methods for extracting information such as names, places, and organizations.

[0022] "Corpus" refers to a collection of one or more documents

[0023] "Features" is any information which is at least partially derived from a document.

[0024] "Event Concept Store" refers to a database of Event template models.

[0025] "Event" refers to one or more features characterized by at least the features' occurrence in real-time.

[0026] "Event Model" refers to a collection of data that may be used to compare against and identify a specific type of event.

[0027] "Module" refers to a computer or software components suitable for carrying out at least one or more tasks.

DETAILED DESCRIPTION

[0028] Reference will now be made in detail to the preferred embodiments, examples of which are illustrated in the accompanying drawings. The embodiments described above are intended to be exemplary. One skilled in the art recognizes that numerous alternative components and embodiments may be substituted for the particular examples described herein and still fall within the scope of the invention. Other embodiments may be used and/or other changes may be made without departing from the spirit or scope of the present disclosure. The illustrative embodiments described in the detailed description are not meant to be limiting of the subject matter presented here.

[0029] Reference will now be made to the exemplary embodiments illustrated in the drawings, and specific language will be used here to describe the same. It will nevertheless be understood that no limitation of the scope of the invention is thereby intended. Alterations and further modifications of the inventive features illustrated here, and additional applications of the principles of the inventions as illustrated here, which would occur to one skilled in the relevant art and having possession of this disclosure, are to be considered within the scope of the invention.

[0030] The present disclosure describes a system and method for detecting, extracting and validating events from a plurality of sources. Sources may include news sources, social media websites and/or any sources that may include data pertaining to events.

[0031] Various embodiments of the systems and methods disclosed here collect data from different sources in order to identify independent events.

[0032] FIG. **1** shows components of a system **100** comprising external data sources **106**, **108** communicatively coupled over a network **104** to an event extraction system **102**. Event extraction system **102** may receive data from a plurality of data sources **106**, **108** through a network **104**. Non-limiting examples of data sources **106**, **108** may include social media **106**, subscription or news sources **108**, though other data

sources **106**, **108** that store and/or publish information may be implemented such as, for example, a document corpus of historical events. Such data sources **106**, **108** may store and/or publish machine-readable data representing unstructured texts such as, for example, Tweets® (i.e., text strings), a news article, or a Facebook® status message.

[0033] A network **104** may be a connection between the different sources and event extraction system **102** through the Internet or an intranet. The network **104** may comprise any suitable collection of hardware and software components (e.g., network interface cards, routers, switches, firewalls, antennas, towers, hubs, trunks) capable of supporting networked communications between computing devices through any suitable protocol (e.g., TCP/IP, 3G, 4G, Bluetooth).

[0034] Event extraction system **102** may include a plurality of components (not illustrated in FIG. **1**) capturing and processing data received from a plurality of data sources **106**, **108**. Event extraction system **102** may comprise software with programmatic logic that may process inputs from the data sources **106**, **108**, and then identity and extract independent events and entities. Event extraction system **102** may be implemented in a single server computer or in a distributed architecture across a plurality of server computers.

[0035] Event extraction system **102** may store extracted events in event store **110**. Event store **110** may be a database implemented in known in the art database management systems (DBMS) such as, for example, MySQL®, PostgreSQL, SQLite, Microsoft SQL Server®, Microsoft Access®, Oracle®, SAP®, dBASE, FoxPro®, IBM DB2®, LibreOffice Base®, FileMaker Pro®, and/or any other type of database that may organize collections of data. Event store **110** may also be a No-SQL database such as, for example, MongoDB®, Couchbase®, H-Base®, Cassandra®, Accumulo®, and/or any other type of database that may organize collections of data.

[0036] Data sources **106**, **108** may be any form of computing service that stores, publishes, transmits, or otherwise makes available over a network **104** data representing information about events and entities. Data sources **106**, **108** may comprise one or more computing devices, servers, and other computing hardware capable of storing data, such as a database, and publishing data over a network **104**, such as a webserver. Non-limiting examples of data sources **106**, **108** may include social media networks **106**, online news sources **108**, blogs, educational portals (e.g., Blackboard®, online university libraries), online journals and magazines, among others. Social media **106** may be any computing service hosting on one or more servers information exchanges between users. Social media **106** users may publish webpages containing text, hyperlinks, and/or other forms of media that is then viewable by other users. New sources **108** may be any computing service hosting on one or more servers a web-based new outlet that publishes webpages containing text, hyperlinks, and/or other forms of media. Data sources **106**, **108** may publish data containing information that may be received and analyzed by an event extraction system **102** via webpages (e.g., HTML, PHP), RSS, e-mail, SMS, or other suitable protocol for publishing information across a computing network **104**.

[0037] FIG. **2** is a flow diagram of an event extraction method **200** according to an embodiment. Event extraction method **200** may begin when data is received from one or more data sources **202**. Data sources **202** may include social media computing services **202***a*, web-based news sources **202***b*, and/or any other data sources **202***c* that store and/or publish data containing information related to events.

[0038] In a first step **218**, after event extraction system receives data from one or more data sources **202**, event extraction system may perform event, topic and entity extraction, which may include the sub-steps **204** (pre-processing and data normalization), **206** (entity and topic extraction, and disambiguation), and **208** (event extraction).

[0039] In a first sub-step **204**, pre-processing and data normalization may be performed by a software module implemented on a computer as part of an event extraction system performing event, topic and entity extraction **218**. A module performing pre-processing and data normalization, first sub-step **204**, may contain programmatic logic, which may involve the use of natural language processing techniques (NLP) for identifying key features in data received from a data source **202**. Non-limiting examples of NLP techniques may include removing stop words, tokenization, stemming and part-of speech tagging among others know in the art.

[0040] In a next sub-step **206**, after a pre-processing and data normalization sub-step **204**, normalized pre-processed data may go through an entity/topic extraction and disambiguation, in which a software module of the event extraction system may identify and extract entities from the data and disambiguate independent entities from one another. Non-limiting examples of entities may include people, organizations, geographic locations, medical conditions, weapons, dates, time or any other entities. Entity and topic identification, extraction, and disambiguation of sub-step **206**, may be performed by one or more software module implemented in a computer as part of event extraction system.

[0041] In a simultaneous, subsequent, or previous sub-step **208**, an event extractor software module may identify possible event model candidates in the text received from the data source **202**. Different types of events may include an accident (e.g., car accident, a train accident, etc.), a natural disaster (e.g., an earthquake, a flood, a weather event, etc.), a man-made disaster (e.g., a bridge collapse, a discharge of a hazardous material, an explosion, etc.), a security event (e.g., a terrorist attack, an act of war, etc.), a major sporting event or concert, election day coordination, traffic incident, and/or any other event. Latent Dirichlet Allocation (LDA), or other methods of detecting and extracting events may be used to extract events. The event extractor module performing sub-step **208** may be executed in conjunction with an event concept store **210**. Event concept store **210** may be a database residing on any suitable computing device comprising non-transitory machine-readable storage media that stores event models. Event models may be compared against event model candidates identified in data. That is, in sub-step **208**, the event extractor module may identify types of features, which in this example are keywords, in the normalized pre-processed data received from the data source **202**, and compare the features against event models stored in the event concept store **210**. The event extractor module may then compute a likelihood score representing the likelihood a set of features (e.g., keywords) pertains to a certain event model, based on comparing the features against each of the event models stored in the event concept store **210**. In some implementations, a comparison between features of a event model candidate and a event model yielding a score between determined thresholds may indicate that the event model being compared is actually referenced in the data source.

[0042] In a next step **212**, after event, topic, and entity extraction of step **218**, the process may perform a spatial-temporal event grouping of extracted events and entities. That is, entities extracted from a data received from a data source **202***b* as a result of executing step **206** (entity extraction and disambiguation), and event model candidates identified in data step **202***b* during execution of step **208** (event extraction) may be associated together, as a spatial-temporal grouping, and then stored in non-transitory machine-readable storage memory. In cases having a plurality of data sources **202***a-c*, event model candidates identified in other data sources **202***a*, **202***c* and entities extracted from other data sources **202***a*, **202***c* may also be associated with one another, and then included to the spatial-temporal event grouping.

[0043] In a next step **214**, after generating spatial-temporal event groupings based on entities and event model candidates extracted from data sources **202**, software modules may perform event validation on the event model candidates in the spatial temporal event groupings. Event validation modules may compare spatial-temporal groupings (i.e., event model candidates and associated entities) extracted from different data sources **202***a-c* in order to determine whether a particular event model candidate extracted from a particular data source **202***b* resembles a real-time event being referenced in the different data sources **202***a*, **202***c*. Spatial-temporal groupings of different data sources **202***a*, **202***c* resembling a co-occurrence of event model candidates and entities of the particular data source **202***b* being validated may serve as validation that the event model candidate of the data source **202***b* resembles the event occurring in real-time.

[0044] Once validated in step **214**, the event model candidate and the associated entitles extracted from the data source **202***b* may be stored into a verified event store **216** database. For example, a server publishing text strings of a Twitter® feed may contain information describing a car accident in Washington D.C, while a news feed channel (e.g., text-based RSS) may contain text strings describing a car accident and high traffic jam in an area nearby the location referenced in the Twitter® feed. In this example of step **214**, an event validation software module may calculate a probability score that both text-based streams of data are describing the same real-world event. When the probability score reaches an established threshold, the event may be considered verified and thus stored into the verified event store **216**.

[0045] In some embodiments, a verified event store **216** may be used by different applications in order to query for different events depending on the purpose of the application. For example, an emergency service application may query for events related to vehicle accidents, fires and the like in order to provide first responders assistance. Another example may be a sports application which may query the database in order to determine the latest information in the NFL® Super Bowl®.

[0046] Event notification **220** may be used to push notifications or alerts to subscribers who wish to be notified immediately when events are verified. Once an event is verified, any subscribers who wish to receive notifications for that event type will be notified of the verified event.

[0047] FIG. **3** is an example embodiment of a training process **300** to build an event concept store **310**. A plurality of documents **304** may be manually tagged by a person or an automated process in order to identify features pertaining to specific events and assign weights to those features. Features can be keywords, entities, topics or any other feature derived from the document. For example, an event model **306** for modeling the event, "Explosion"; a person can manually identify in a document **304** relating to an explosion the co-occurrence of keywords such as "Bomb" and/or "Fire". The user may then assign a weight to each word depending on the repetition or the co-occurrence of these keywords with others in a plurality of documents **304** related to explosions, and associate those with an event model **306** for "Explosion" stored in event concept store **310**.

[0048] FIG. **4** is an example embodiment of an event detection process **400**. An event extractor **406** may identify features, which in this example are keywords **402**, from a data input. In the illustrated example, the keywords "Bomb" and "Fire" are identified. The keywords may then be submitted for comparison against event models in the event concept store **410**. In this example, event concept store **410** may assign weights of 0.1 to "Bomb" and 0.3 to "Fire" for the event model of "Explosion". Event extractor **406** may then add up the weighted scores and determine if the resulting score exceeds a determined threshold. In this example a 0.4 score is generated for the probability of the event being an "Explosion"; however other methods of calculating weighted scores may be used and are included within the scope of this disclosure. Event extractor **406** may then transfer an event possibility of 0.4 of explosion to spatial-temporal event grouping **412**.

[0049] After pre-processing and normalization, entity extraction and disambiguation, and event extraction, each of the identified event model candidates and associated entities/topics from each of the different sources may be grouped together in a spatial-temporal event grouping **412**, which may be stored as a record of the spatial-temporal grouping **412**.

[0050] FIG. **5** is an example embodiment of an event grouping process **500**. The process may begin by taking an input from different data sources **502**. Each input from data sources **502** may go through an event, topic and entity extraction **518** process. Entities, topics and event model candidates **504** are extracted from the different data sources **502**. Entities, topics and event model candidates **504** from different data sources **502** may then be grouped together in spatial-temporal grouping **512** and an initial confidence score **506** may be assigned to each entity, topic and event model candidate **504** association.

[0051] Validation module **514** may compare the different records stored in the spatial-temporal grouping, and identify an overlap between entities and event model candidates **504** from each of the different data sources **502**. A score may be calculated using the initial confidence score **506** from the different entities and event model candidates **504** that overlap and/or repeat themselves in different data sources **502**. A score greater than a predetermined threshold may serve as an indication that the event model candidate actually occurred. A verified event may then be stored in verified event store **516**. In the exemplary embodiment illustrated in FIG. **5** an overlap of entities "Microsoft" and "Seattle" are extracted along with the event model candidate "Explosion" from different sources this may serve as an indication that an explosion has occurred at Microsoft®, in Seattle.

[0052] In Example #1 a tweet is extracted from Twitter and ingested into the event extraction system **102**. The Tweet® contains the message "Bill Gates the chairman of Microsoft was Kidnapped in Syria". The process may go through pre-processing and data normalization **204** step where stop words are removed. The process may then continue and extract

entities "Bill", "Gates", "chairman", "Microsoft", "Syria" in entity/topic extraction and disambiguation **206** step and extract the event "Kidnapped" using event extraction method **200**. The entity extraction process may then identify Bill Gates as Chairman of Microsoft®, and associate the entity with the event kidnapped in spatial-temporal event grouping **212**. Event validation **214** may then compare the kidnapping event of Bill Gates to other events from other sources also in spatial-temporal event grouping **212**. Event validation **214** may identify if other events also refer to the kidnapping of the Chairman of Microsoft®, Bill Gates, in Syria, and thus validate if the event is real. If the event is real it may be transferred to verified event store **216**

[0053] While various aspects and embodiments have been disclosed, other aspects and embodiments are contemplated. The various aspects and embodiments disclosed are for purposes of illustration and are not intended to be limiting, with the true scope and spirit being indicated by the following claims.

[0054] The foregoing method descriptions and the process flow diagrams are provided merely as illustrative examples and are not intended to require or imply that the steps of the various embodiments must be performed in the order presented. As will be appreciated by one of skill in the art the steps in the foregoing embodiments may be performed in any order. Words such as "then," "next," etc. are not intended to limit the order of the steps; these words are simply used to guide the reader through the description of the methods. Although process flow diagrams may describe the operations as a sequential process, many of the operations can be performed in parallel or concurrently. In addition, the order of the operations may be re-arranged. A process may correspond to a method, a function, a procedure, a subroutine, a subprogram, etc. When a process corresponds to a function, its termination may correspond to a return of the function to the calling function or the main function.

[0055] The various illustrative logical blocks, modules, circuits, and algorithm steps described in connection with the embodiments disclosed herein may be implemented as electronic hardware, computer software, or combinations of both. To clearly illustrate this interchangeability of hardware and software, various illustrative components, blocks, modules, circuits, and steps have been described above generally in terms of their functionality. Whether such functionality is implemented as hardware or software depends upon the particular application and design constraints imposed on the overall system. Skilled artisans may implement the described functionality in varying ways for each particular application, but such implementation decisions should not be interpreted as causing a departure from the scope of the present invention.

[0056] Embodiments implemented in computer software may be implemented in software, firmware, middleware, microcode, hardware description languages, or any combination thereof. A code segment or machine-executable instructions may represent a procedure, a function, a subprogram, a program, a routine, a subroutine, a module, a software package, a class, or any combination of instructions, data structures, or program statements. A code segment may be coupled to another code segment or a hardware circuit by passing and/or receiving information, data, arguments, parameters, or memory contents. Information, arguments, parameters, data, etc. may be passed, forwarded, or transmitted via any suitable means including memory sharing, message passing, token passing, network transmission, etc.

[0057] The actual software code or specialized control hardware used to implement these systems and methods is not limiting of the invention. Thus, the operation and behavior of the systems and methods were described without reference to the specific software code being understood that software and control hardware can be designed to implement the systems and methods based on the description herein.

[0058] When implemented in software, the functions may be stored as one or more instructions or code on a non-transitory computer-readable or processor-readable storage medium. The steps of a method or algorithm disclosed herein may be embodied in a processor-executable software module which may reside on a computer-readable or processor-readable storage medium. A non-transitory computer-readable or processor-readable media includes both computer storage media and tangible storage media that facilitate transfer of a computer program from one place to another. A non-transitory processor-readable storage media may be any available media that may be accessed by a computer. By way of example, and not limitation, such non-transitory processor-readable media may comprise RAM, ROM, EEPROM, CD-ROM or other optical disk storage, magnetic disk storage or other magnetic storage devices, or any other tangible storage medium that may be used to store desired program code in the form of instructions or data structures and that may be accessed by a computer or processor. Disk and disc, as used herein, include compact disc (CD), laser disc, optical disc, digital versatile disc (DVD), floppy disk, and blu-ray disc where disks usually reproduce data magnetically, while discs reproduce data optically with lasers. Combinations of the above should also be included within the scope of computer-readable media. Additionally, the operations of a method or algorithm may reside as one or any combination or set of codes and/or instructions on a non-transitory processor-readable medium and/or computer-readable medium, which may be incorporated into a computer program product.

[0059] It is to be appreciated that the various components of the technology can be located at distant portions of a distributed network and/or the Internet, or within a dedicated secure, unsecured and/or encrypted system. Thus, it should be appreciated that the components of the system can be combined into one or more devices or co-located on a particular node of a distributed network, such as a telecommunications network. As will be appreciated from the description, and for reasons of computational efficiency, the components of the system can be arranged at any location within a distributed network without affecting the operation of the system. Moreover, the components could be embedded in a dedicated machine.

[0060] Furthermore, it should be appreciated that the various links connecting the elements can be wired or wireless links, or any combination thereof, or any other known or later developed element(s) that is capable of supplying and/or communicating data to and from the connected elements. The term module as used herein can refer to any known or later developed hardware, software, firmware, or combination thereof that is capable of performing the functionality associated with that element. The terms determine, calculate and compute, and variations thereof, as used herein are used interchangeably and include any type of methodology, process, mathematical operation or technique.

[0061] The preceding description of the disclosed embodiments is provided to enable any person skilled in the art to make or use the present invention. Various modifications to these embodiments will be readily apparent to those skilled in

the art, and the generic principles defined herein may be applied to other embodiments without departing from the spirit or scope of the invention. Thus, the present invention is not intended to be limited to the embodiments shown herein but is to be accorded the widest scope consistent with the following claims and the principles and novel features disclosed herein.

[0062] The embodiments described above are intended to be exemplary. One skilled in the art recognizes that numerous alternative components and embodiments that may be substituted for the particular examples described herein and still fall within the scope of the invention.

What is claimed is:

1. A method comprising:

extracting, by a server, a first entity from a first data of a first data source and a second entity from a second data of a second data source;

identifying, by the server, a first event model candidate from the first data and a second event model candidate from the second data;

comparing, by the server, the first event model candidate against a set of event models stored in a first database and the second event model candidate against the set of event models;

determining, by the server, a first score indicative of a first probability of a first match for the first event model based on the comparing and a second score indicative of a second probability of a second match for the second event model based on the comparing;

associating, by the server, the first score, the first entity, and the first event model candidate as a first spatial-temporal grouping and the second score, the second entity, and the second event model candidate as a second spatial-temporal grouping;

validating, by the server, an event based on an overlap of at least one of the first entity or the first event model candidate from the first spatial-temporal grouping and at least one of the second entity or the second event model candidate from the second spatial-temporal grouping and based on the first score from the first spatial-temporal grouping and the second score from the first spatial-temporal grouping exceeding a threshold representative of the event actually occurring; and

storing, by the server, a record associated with the event in a second database based on the validating.

2. The method of claim **1**, wherein the first database and the second database are one database.

3. The method of claim **1**, wherein the first database and the second database are different databases.

4. The method of claim **1**, wherein at least one of the first database or the second database is an in-memory database.

5. The method of claim **1**, further comprising:

pushing, by the server, a message to a client responsive to the storing, wherein the message is informative of the record.

6. The method of claim **5**, wherein the pushing is based on a rule set by the client.

7. The method of claim **1**, wherein at least one of the first data source or the second data source comprises at least one of social media network, a newsfeed, a blog hosting server, an education portal, or a document.

8. The method of claim **1**, further comprising:

pre-processing, by the server, at least one of the first data or the second data at least one of before the extracting the first entity and the second entity or before the extracting the first event model candidate and the second event model candidate.

9. The method of claim **1**, wherein at least one model from the set of event models is based on an information manually provided to the at least one model by a user.

10. The method of claim **1**, wherein the extracting the first entity and the second entity comprises disambiguating at least one of the first entity or the second entity.

11. A system comprising:

a server;

a first database in communication with the server;

a second database in communication with the server,

wherein the server is configured to:

extract a first entity from a first data of a first data source and a second entity from a second data of a second data source,

identify a first event model candidate from the first data and a second event model candidate from the second data,

compare the first event model candidate against a set of event models stored in the first database and the second event model candidate against the set of event models,

determine a first score indicative of a first probability of a first match for the first event model based on the comparison and a second score indicative of a second probability of a second match for the second event model based on the comparison,

associate the first score, the first entity, and the first event model candidate as a first spatial-temporal grouping and the second score, the second entity, and the second event model candidate as a second spatial-temporal grouping,

validate an event based on an overlap of at least one of the first entity or the first event model candidate from the first spatial-temporal grouping and at least one of the second entity or the second event model candidate from the second spatial-temporal grouping and based on the first score from the first spatial-temporal grouping and the second score from the first spatial-temporal grouping exceeding a threshold representative of the event actually occurring, and

store a record associated with the event in the second database based on the validating.

12. The system of claim **11**, wherein the first database and the second database are one database.

13. The system of claim **11**, wherein the first database and the second database are different databases.

14. The system of claim **11**, wherein at least one of the first database or the second database is an in-memory database.

15. The system of claim **11**, wherein the server is configured to push a message to a client responsive to the storage, wherein the message is informative of the record.

16. The system of claim **15**, wherein the pushing is based on a rule set by the client.

17. The system of claim **11**, wherein at least one of the first data source or the second data source comprises at least one of social media network, a newsfeed, a blog hosting server, an education portal, or a document.

18. The system of claim **11**, wherein the server is configured to pre-process at least one of the first data or the second data at least one of before the extraction of the first entity and

the second entity or before the extraction of the first event model candidate and the second event model candidate.

19. The system of claim **11**, wherein at least one model from the set of event models is based on an information manually provided to the at least one model by a user.

20. The system of claim **11**, wherein the extraction of the first entity and the second entity comprises the server disambiguating at least one of the first entity or the second entity.

* * * * *