

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
6 August 2009 (06.08.2009)

PCT

(10) International Publication Number
WO 2009/097558 A2

(51) International Patent Classification:

G06F 17/27 (2006.01)

(21) International Application Number:

PCT/US2009/032695

(22) International Filing Date: 30 January 2009 (30.01.2009)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:

61/063,047 30 January 2008 (30.01.2008) US

12/341,926 22 December 2008 (22.12.2008) US

(71) Applicant (for all designated States except US): **THOMSON REUTERS GLOBAL RESOURCES** [IE/CH]; Landis + Gyr-Str. 3, CH-6300 Zug (CH).

(71) Applicants and

(72) Inventors: **SCHILDER, Frank** [DE/US]; 1764 Hague Avenue, St. Paul, MN 55104 (US). **DOZIER, Christopher** [US/US]; 4825 Knox Avenue So., Minneapolis, MN 55419 (US). **KONDADADI, Ravi Kumar** [IN/US]; 3575 Lexington Avenue So., Eagan, MN 55123 (US).

(74) Agent: **STEFFEY, Charles E.**, Uspto Re; 121 So. Eighth Street, Suite 1600, Minneapolis, MN 55402 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, 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, IS, JP, KE, KG, KM, 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, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, 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, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (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, SE, SI, SK, TR), OAPI (BF, BJ,

[Continued on next page]

(54) Title: FINANCIAL EVENT AND RELATIONSHIP EXTRACTION

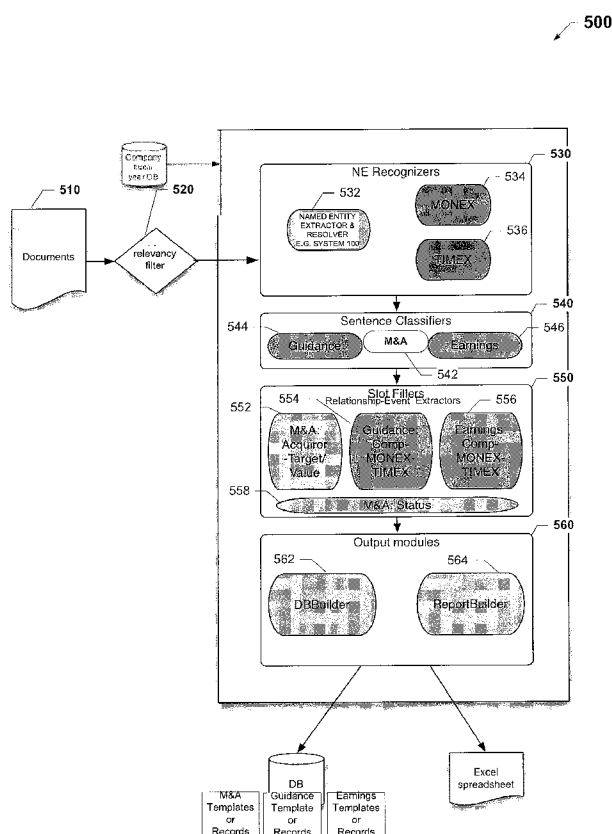


FIGURE 5

(57) Abstract: For automated text processing, the inventors devised, among other things, an exemplary system that automatically extracts financial events from various unstructured text based sources, such as press releases and news articles. Extracted events, such as mergers & acquisitions, earnings guidance reports, and actual earnings announcements, are represented as structured data records which can be linked, searched, and displayed and used as a basis for controlling accessing to the source documents and other related financial documents for named entities.



CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN,
TD, TG).

Published:

- *without international search report and to be republished
upon receipt of that report*

FINANCIAL EVENT AND RELATIONSHIP EXTRACTION

5 Copyright Notice and Permission

A portion of this patent document contains material subject to copyright protection. The copyright owner has no objection to the facsimile reproduction by anyone of the patent document or the patent disclosure, as it appears in the Patent and Trademark Office patent files or records, but otherwise reserves all
10 copyrights whatsoever. The following notice applies to this document:
Copyright © 2007-2008, Thomson Reuters Global Resources.

Related Applications

This application claims priority to U.S. Patent Application 12/341,926,
15 which was filed December 22, 2008, and to U.S. Provisional Application
61/063,047 which was filed January 30, 2008. Both of these applications are
incorporated herein by reference.

Technical Field

20 Various embodiments of the present invention concern extraction of data
and related information from documents, such as identifying and tagging names
and events in text and automatically inferring relationships between tagged
entities, events, and so forth.

25 Background

The present inventors recognized a need to provide information
consumers relational and event information about entities, such as companies,
30 persons, cities, that are mentioned in electronic documents, particularly financial
documents. For example, documents, such as news feeds, SEC (Securities and
Exchange Commission) filings may indicate that Company A merged with or is
rumored to be merging with Company B, or that Company C announced actual
or projected earnings of X dollars per share.

35 However, because of language variations and the unstructured nature of
many of the documents, automatically discerning the relational and event

information about these entities is difficult and time consuming even with state-of-the art computing equipment.

Summary

5 To address this and/or other needs, the present inventors devised, among other things, systems and methods for named-entity tagging and event and relationship extraction from documents, such as financial news articles and press releases.

10 The exemplary system automatically extracts financial events from various unstructured text based sources, such as press releases and news articles. Extracted events, such as mergers & acquisitions, earnings guidance reports, and actual earnings announcements, are represented as structured records.

To achieve this end, the exemplary system includes a set of recognizers, a set of text segment classifiers, and a set of relationship extractors. The set of
15 recognizers receive input text and tag, resolve, and normalize entities, monetary amounts, and temporal indicators in the text segments, such as sentences. Receiving the text segments, the text segment classifiers classify what types of events, such as financial events, the text segments may include, and route to an appropriate one of a set of relationship extractors. For a text segment that
20 includes an event, such as a mergers and acquisition event, the relationship extractor determines the role of named entities in the text segment within the event, associated monetary values, and/or timing or status of the event.

25 Brief Description of the Drawings

Figure 1 is a block and flow diagram of an exemplary system for named-entity tagging, resolving and event extraction, which corresponds to one or more embodiments of the present invention.

Figure 2 is a diagram illustrating guided sequence decoding for named-
30 entity tagging which corresponds to one or more embodiments of the present invention.

Figure 3 is a block diagram of an exemplary named-entity tagging, resolution, and event extraction system corresponding to one or more embodiments of the present invention.

Figure 4 is a flow chart of an exemplary method of named-entity tagging and resolution and event extraction corresponding to one or more embodiments of the present invention.

Figure 5 is a block and flow diagram of another exemplary system for
5 named-entity tagging and resolving, and event extraction, which corresponds to one or more embodiments of the present invention.

Detailed Description of the Exemplary Embodiment(s)

This description, which incorporates the Figures and the claims,
10 describes one or more specific embodiments of an invention. These embodiments, offered not to limit but only to exemplify and teach the invention, are shown and described in sufficient detail to enable those skilled in the art to implement or practice the invention. Thus, where appropriate to avoid obscuring the invention, the description may omit certain information known to those of
15 skill in the art.

Exemplary Named-Entity Tagging and Resolution System

Figure 1 shows an exemplary named entity tagging and resolving system 100. In addition to processors 101 and a memory 102, system 100 includes an entity tagger 110, an entity resolver 120, and authority files 130. (Tagger 110,
20 resolver 120, and authority files 130 are implemented using machine-readable data and/or machine-executable instructions stored on memory 102, which may take a variety of consolidated and/or distributed forms.

Entity tagger 110, which receives textual input in the form of documents
25 or other text segments, such as a sentence 109, includes a tokenizer 111, a zoner 112, and a statistical tagger 113.

Tokenizer 111 processes and classifies sections of a string of input characters, such as sentence 109. The process of tokenization is used to split the sentence or other text segment into word tokens. The resulting tokens are output
30 to zoner 112..

Zoner 112 locates parts of the text that need to be processed for tagging, using patterns or rules. For example, the zoner may isolate portions of the

document or text having proper names. After that determination, the parts of the text that need to be processed further are passed to statistical sequence tagger 113.

Statistical sequence tagger 113 (or decoder) uses one or more unambiguous name lists (lookup tables) 114 and rules 115 to tag the text within sentence 109 as company, person, or place or as a non-name. The rules and lists are regarded herein as high-precision classifiers.

Exemplary pattern rules can be implemented using regex+Java, Jape rules within GATE, ANTLR, and so forth. A sample rule for illustration dictates that “if a sequence of words is capitalized and ends with “Inc.” then it is tagged as a company or organization. The rules are developed by a human (for example, a researcher) and encoded in a rule formalism or directly in a procedural programming language. These rules tag an entity in the text when the preconditions of the rule are satisfied.

Exemplary name lists identify companies, such as Microsoft, Google, AT&T, Medtronics, Xerox; places, such as Minneapolis, Fort Dodge, Des Moines, Hong Kong; and drugs, such as Vioxx, Viagra, Aspirin, Penicillin. In the exemplary embodiment, the lists are produced offline and made available during runtime. To produce the list, a large corpus of documents, for example, a set of news stories, is passed through a statistical model and/or various rules (for example, a conditional random field (CRF) model) to determine if the name is considered unambiguous. Exemplary rules for creating the lists include: 1) being listed in a common noun dictionary; and 2) being used as company name more than ninety percent of the time the name is mentioned in a corpus. The lookup tagger also finds systematic variants of the names to add to the unambiguous list. In addition, the lookup tagger guides and forces partial solutions. Using this list assists the statistical model (the sequence tagger) by immediately pinning that exact name without having to make any statistical determinations.

Examples of statistical sequence classifiers include linear chain conditional random field (CRF) classifiers, which provide both accuracy and speed. Integrating such high precision classifiers with the statistical sequence labeling approach entails first modifying the feature set of the original statistical model by including features corresponding to the labels assigned by the high-

precision classifiers, in effect turning "on" the appropriate label features depending on the label assigned by the external classifier. Second, at run time, a Viterbi decoder (or a decoder similar in function) is constrained to respect the partially labeled or tagged sequences assigned by the high- precision classifiers.

5 This form of guided decoding provides several benefits. First, the speed of the decoding is enhanced, because the search space is constrained by the pretagging. Second, results are more consistence, because three sources of knowledge are taken account: the lists, the rules, and trained decoder statistical model. The third benefit is an ease of customization that stems from an
10 elimination of a need to retrain the decoder if new rules and list items are added.

Figure 2 is a conceptual diagram showing how a text segment "Microsoft on Monday announced a" is pretagged and how this pretagging (or pinning) constrains the possible tags or labeling options that a decoder, such as Viterbi
15 decoder, has to process. In the Figure, the term Microsoft is tagged or pinned as a company based on its inclusion in a list of company names; the term Monday is marked as "out" based on its inclusion of a list of terms that should always be marked as "out"; and the term "on" is marked as out based on a rule that it should be marked as "out", if it is followed by an term that is marked as "out" in
20 this case the term "Monday."

In the exemplary embodiment, the statistical sequence tagger calculates the probability of a sequence of tags given the input text. The parameters of the model are estimated from a corpus of training data, that is, text where a human has annotated all entity mentions or occurrences. (Unannotated text may also be
25 used to improve the estimation of the parameters.) The statistical model then assembles training data, develops a feature set and utilizes rules for pinning. Pinning is a specific way to use a statistical model to tag a sequence of characters and to integrate many different types of information and methods into the tagging process.

30 The statistical model locates the character offset positions (that is, beginning and end) in the document for each named entity. The document is a sequence of characters; therefore, the character offset positions are determined. For example, within the sentence "Hank's Hardware, Inc. has a sale going on right now," the piece of text "Hank's Hardware, Inc." has an offset position of

(0, 20). The sequence of characters has a beginning point and an ending point; however the path in between those points varies.

After the character offset positions are located, information about the entity is identified through the use of features. This information ranges from general information (that is, determining text is last name) to specific information (e.g., unique identifier). The exemplary embodiment uses the features discussed below, but other embodiments use other types and numbers amounts of features:

- Regular expressions: contains an uppercase letter, last char is a dot, Acronym format, contains a digit, punctuation
- Single word lists: last names, job titles, loc words, etc.
- Multi-word lists: country names, country capitals, universities, company names, state names, etc.
- Combination features: title@-1 AND (firstname OR last)
- Copy features: copies features from one token to neighboring tokens, for example, the token two to the left of me is capitalized (Cap@-2)
- The word itself features: “was” has the feature was@0
- First-sentence features: copy features from 1st sentence words to others
- Abbreviation feature: copy features of name to mentions of abbr.

The features computation does not calculate features for isolated pinned tokens. The computations combine hashes, combine tries, and combine regular expressions. Features are only computed when necessary (for example punctuation tokens are not in any hashes so do not look them up). Once the model has been trained, the Viterbi algorithm (or an algorithm similar in function) is used to efficiently find the most probable sequence of tags given the input and the trained model. After the algorithm determines the most probable sequence of tags, the text, such as tagged sentence 119, where the entities are located is passed to a resolver, such as entity resolver 120.

Entity resolver 120 provides additional information on an entity by matching an identifier for an external object within authority files 130 to which the entity refers. The resolver in the exemplary embodiment uses rules instead of a statistical model to resolve named entities. In the exemplary embodiment, the external object is a company authority file containing unique identifiers. The exemplary embodiment also resolves person names.

The exemplary resolver uses three types of rules to link names in text to authority file entries: rules for massaging the authority file entries, rules for

normalizing the input text, and rules for using prior links to influence future links. Other embodiments include integrating the statistical model and resolver.

This list along with the original text is the input to an entity resolver module. The entity resolver module takes these tagged entities and decides which element in an authority file the tagged entity refers. In the exemplary embodiment, authority file 130 is a database of information about entities. For example an authority file entry for Swatch might have an address for the company, a standard name such as Swatch Ltd., the name of the current CEO, and a stock exchange ticker symbol. Each authority file entry has a unique identity. In the previous example a unique id could be, ID:345428, "Swatch Ltd.", Nicholas G. Hayek Jr., UHRN.S. The goal of the resolver is to determine which entry in the authority file matches corresponds a name mention in text. For example, it should figure out the Swatch Group refers to entity ID:345428. Of course, resolving names like Swatch is relatively easy in comparison to a name like Acme. However, even for names like Swatch, a number of related but different companies may be possible referents. What follows is a heuristic resolver algorithm used in the exemplary embodiment:

Heuristic Resolver Algorithm for Companies

Iterate through entities tagged by the CRF:
If entity tagged as ORG:
If a "do not resolve" ORG (i.e., stock exchange abbreviations):
set ID attribute to "NOTRESOLVED"
Else:
If entity in the company authority file,
set ID attribute to company ID
Else:
set ID attribute to "NOTRESOLVED"
Iterate through NOTRESOLVED entities:
If E is a left-anchored substring of a resolved company:
set ID attribute to already resolved company substring match ID,
change the tag kind to ORG, if necessary
If E is an acronym of an already-resolved company:
set ID attribute to already resolved non-acronym company ID,
change the tag kind to ORG, if necessary

Note that the exemplary entity tagger and variations thereof is not only useful for named entity tagging. Many important data mining tasks can be framed as sequence labeling. In addition, there are many problems for which

high precision (but low recall) external classifiers are available that may have been trained on a separate training set.

Exemplary Event and Relationship Extraction System

5 Figure 3 shows an exemplary system 300 which builds onto the components of system 100 with a classifier 310 and a template extractor 320, which are shown as part of memory 102, and understood to be implemented using machine-readable and machine-executable instructions.

Classifier 310, which accepts tagged and resolved text such as sentence
10 129 from resolver 120, identifies sentences that contain extractable relationship information pertaining to a specific relationship class. For example, if one is interested in the hiring relationship where the relationship is hire(firm, person), the filter (or classifier) 312 identifies sentence (1.1) as belonging to the class of sentences containing a hiring or job-change event and sentence (1.2) as not
15 belonging to the class.

(1.1) *John Williams has joined the firm of Skadden & Arps as an associate.*

(1.2) *John Williams runs the billing department at Skadden & Arps.*

The exemplary embodiment implements classifier 310 as a binary
20 classifier. In the exemplary embodiment, building this binary classifier for relationship extraction entails:

- 1) Extracting articles from a target database;
- 2) Splitting sentences in all articles and loading to a single file;
- 3) Tagging and resolving types of entities relevant to a relationship type
25 that occur within each sentence;
- 4) Selecting from set of sentences all sentences that have the minimal number of tagged entities needed to form a relationship of interest. This means for example that at least one person name and one law firm name must be specified in a sentence for it to contain a job change event.
30 Sentences containing requisite number of tagged entity types are called candidate sentences; 5) Identifying 500 positive instances from the candidate set and 500 negative instances. A sentence in the candidate set that actually contains a relationship of interest is called a positive instance. A sentence in the candidate set that does not contain a

relationship of interest is called a negative instance. All sentences within the candidate set are either positive or negative instances. These sampled instances should be representative of their respective sets and should be found as efficiently as possible;

5 6) Creating classifier that combines selected features with selected training methods. Exemplary training methods include naive Bayes and Support Vector Machine (SVM.) Exemplary features include co-occurring terms and syntax trees connecting relationship entities; and

7) Testing the classification of randomly selected sentences from
10 candidate pool. After testing the exemplary embodiment evaluates first hundred sentences classified as positive (for example, job change event containing) and first hundred classified as negative, computing precision and recall and saving evaluated sentences as gold data for future testing.

A range of filters that are either document-dependent filters or complex
15 relation detection filters based on machine learning algorithms are developed and tools that easily retarget new document types. The structure of a document type provides very reliable clues on where the sought after information can be found. Ideally, the filter is flexible and automatically detects promising areas in a document. For example, a filter that includes a machine learning tool (for
20 example Weka) that detects promising areas and produces pipelines that can be changed according to the relevant features needed for the task.

Depending on the requirements, different levels of co-reference resolution can be implemented. In some domains, no co-reference resolution is used. Other situations use a relatively simple set of rules for co-reference
25 resolution, based on recent mentions in the text and identifiable attributes (i.e., gender, plurality, etc.) of the interested named entities. For example, in the job change event, almost all co-reference issues are solved by simply referring backward to the most recent mention of the matching entity type (that is, law firm or lawyer name).

30 Template extractor 320 extracts event templates from positively classified sentences, such as sentence 319, from classifier 310. In the exemplary embodiment, extracting templates from sentences involves identifying the name entities participating in the relationship and linking them together so that their

respective roles in the relationship are identified. A parser is utilized to identify noun phrase chunks and to supply a full syntactic parse of the sentence.

In the exemplary embodiment, implementing extractor 320 entails:

- 5 1) Creating gold data by taking positive example sentences from classification phase and manually generating appropriate template records. The user is automatically presented with all possible templates which could be generated from the sentence and asking the user to select the one that is correct;
- 10 2) Taking 400 sentences from gold data set for training data and develop extraction programs based on one or more of the following technologies: association rules, chunk kernel based on chunks, CRF, and tree kernel based on syntactic structure;
- 3) Testing solutions on 100 held out test samples;
- 15 4) Combining classifier with extractor to test precision using unseen data. For instance, a sentence containing a job change event is one that describes an attorney joining a law firm or other organization in a professional capacity. The target corpora from which job change events are extracted are legal newspaper databases. The minimal number of tagged entities which qualify a sentence for inclusion in the candidate set
- 20 is one lawyer name and one legal organization name. One way to efficiently collect positive and negative training instances is to stratify samplings. This can be done by sorting the sentences according to the head word of the verb phrase that connects a person with a law firm in the sentence. Then collect all head verbs that occur at least five times
- 25 under a single bucket. After collection, select five example sentences from each bucket randomly and mark them as either positive or negative examples. For each bucket that yields only positive examples, add all remaining instances to the positive example pool. And for each bucket that yields only negative examples, add all examples to the negative
- 30 examples group. If there are less than 500 positive examples or less than 500 negative examples, manually score randomly selected sentences until 500 examples of each time are identified. The job change event extractor moves identified entities from a positively classified job change event

sentence into a structured template record. The template record identifies the roles the named entities and tagged phrases play in the event.

The template below (which also represents a data structure) is in reference to sentence 1.1 above.

Role	Value	Entity ID
Attorney	John Williams	A23456
Firm	Skadden & Arps	F56748
Position	Associate	P234

5

Various assumptions are incorporated in the exemplary embodiment.

One main assumption is that the identity of the entities is usually independent of the way of talking about an event or relationship. Another assumption is that the extraction of sentences deemed paraphrases based upon the equality of constituent entities and time window is relatively error-free. The precision of this latter filtering step is improved by having other checks such as on the cosine similarity between the documents in which the two sentences are found, similarity of titles of the documents etc. This approach entails:

- 1) Providing a large corpus of documents preferably having the property that several documents talking about the same event or relationship from different authors are easy to find. One example is a time-stamped news corpus from different news sources, where the same event is likely to be covered by different sources;
- 2) Using a named entity recognizer to tag the entities in the corpus with reasonable accuracy. Clearly the set of entities that need to be covered by the NER (named-entity resolver) depends upon the extraction problem;
- 3) Providing an indexer for efficient search and retrieval from the corpus;
- 4) Providing a human generated list of high-precision sentences with the entities replaced by wild-cards. For example, for MA, a human might provide a rule "ORG1 acquired ORG2" means this is an MA sentence with ORG1 being the buyer and ORG2 being the target.

Exemplary Methods of Operating a Named-Entity Tagging, Resolution and
Event and Relationship Extraction System

Figure 4 shows a flow chart 400 of an exemplary method of operating a named entity tagging, resolution, and event extraction system, such as system 300 in Figure 3. Flow chart 300 includes blocks 410- 460, which are arranged and described serially. However, other embodiments also provide different functional partitions or blocks to achieve analogous results.

Block 410 entails breaking the extracted text into tokens. Execution proceeds at block 220.

10 Block 420 entails locating parts of the extracted text that need to be processed. In the exemplary embodiment, this entails use of zoner 112 to locate candidate sentences for processing. Execution then advances to block 230.

 Block 430 entails finding the named entities within the processed parts of extracted text. Then the entities of interest in the candidate sentences are tagged.
15 Candidate sentences are sentences from target corpus that might contain a relationship of interest. For example, one embodiment identifies text segments that indicate job-change events; another identifies segments that indicate merger and acquisition activity; a yet another identifies segments that may indicate corporate income announcements. Execution continues at block 440.

20 Block 440 entails resolving the named entities. Each entity is attached to a unique ID that maps the entity to a unique real world object, such as an entry in an authority file. Execution then advances to block 250.

 Block 250 classifies the candidate sentences. The candidate sentences are classified into two sets: those that contain the relationship of interest and
25 those that do not. For example, one embodiment identifies text segments that indicate job-change events; another identifies segments that indicate merger and acquisition activity; a yet another identifies segments that may indicate corporate income announcements. When the text is classified, executes advances to block 260.

30 Block 260 entails extracting the relationship of interest using a template. More specifically, this entails extracting entities from text containing the relationship and place the entities in a relationship template that properly defines the relationship between the entities. When the template is completed, the

extracted data may be stored in a database but it may also involve more complex operations such as representing the data according a time line or mapping it to an index.

Some embodiments of the present invention are implemented using a number of pipelines that add annotations to text documents, each component receiving the output of one or more prior components. These implementations use the Unstructured Information Management Architecture (UIMA) framework and ingest plain text and decomposes the text into components. Each component implements interfaces defined by the framework and provide self-describing metadata via XML descriptor files. The framework manages these components and the data flow between them. Components are written in Java or C++; the data that flows between components is designed for efficient mapping between these languages. UIMA additionally provides a subsystem that manages the exchange between different modules in the processing pipeline. The Common Analysis System (CAS) holds the representation of the structured information Text Analysis Engines (TAEs) add to the unstructured data. The TAEs receive results from other UIMA components and produce new results that are added to the CAS. At the end of the processing pipeline, all results stored in the CAS can be extracted from there by the invoking application (for example, database population) via a CAS consumer. Primitive TAEs (for example, tokenizer, sentence splitter) can be bundled into an aggregate TAE. Other embodiments use alternatives to the UIMA.framework.

Exemplary Financial Event Extraction and Resolution System and Method

Figure 5 shows an extension or enhancement of system 300 in the form of a system 500 that automatically extracts and resolves financial events from text documents. Although not explicitly shown in this drawing, system 100 is implemented using one or more processors and memory devices, which store data and machine-readable and executable instructions sets. The processors and memory devices may be organized or arranged in any desirable centralized or distributed computing architecture. Some embodiments implement system 500 as a Java pipeline which can easily be integrated into an editorial workflow. The system can be configured to work in batch mode or as a web service.

Additionally, the system can be configured to operate in batch mode or as a web service.

In particular, system 500 includes a set of electronic documents 510, a relevancy filter 520, recognizers 530, text segment classifiers 540, template or slot fillers 550, and output module 560.

Documents 510 includes a set of unstructured and/or structured textual documents. For example, in the exemplary embodiment documents 510 includes press releases, news wire stories, SEC (Securities and Exchange Commission) documents. Documents 510 are input in batch or serial fashion to relevancy filter 520.

Relevancy filter 520 includes one or more financial event classifiers. In the exemplary embodiment, filter 520 determines, using one or more machine-learning-based classifiers, whether the documents are likely to include text that is representative of a financial event that can be extracted by the system. Exemplary financial events include merger & acquisitions, earnings announcements, or earnings guidance reports. Determinations can be based, for example, on whether two companies are mentioned in a single sentence or within some other defined text segment, such as a paragraph or within a certain distance of each other, or whether a monetary amount is mentioned in proximity to a company name or proximate terms correlated with occurrence of a financial event. Determinations may also be based on inclusion of terms such as merger, acquisition, earnings, and related roots, stems, synonyms, and so forth. Documents that are determined unlikely to include a financial event are excluded from further processing, whereas those that are deemed likely to include such events are input to recognizers 530.

Recognizers 530 extracts and resolves companies, percentage and money amounts in the same general manner as described for system 100. In particular recognizers 530 includes a named entity extractor and resolver 532, a monetary extractor 534, and a temporal extractor 536. Named entity extractor and resolver 532 in the exemplary embodiment is identical to system 100 shown in Figure 1. Monetary extractor 534 identifies and tags percent expressions, monetary expressions, including monetary ranges, the color of the money (actual earnings, or projected earnings, etc.), and possibly a trend (for example, up or down). In the exemplary embodiment, this entails normalizing the percent and money

amount to U.S. currencies, for example. Temporal extractor 536 identifies and tags temporal terms and/or windows. In the exemplary embodiment, the temporal extractor (for example ANTLR lexer, which is also used for parsing monetary expressions) also grounds time expressions (e.g., Q2 means second
 5 quarter of the current year) and converts to an ISO time value. The exemplary embodiment implements this extractor programmatically using the following:

- TIMEX :{timex.initialize();}
 (VAGUE|PERIOD|SPECIFIC|INDEXICAL)
 {timex.computeValue();}
- 10 – Class TimexGroundedInfo records the temporal meaning of the expression and computes the grounded time.
- Indexicals: Today, tomorrow, Wednesday
- Specific: 2008-05-06T02:30:30
- Periods: 3 months
- 15 – Vague: Late Monday
- Anaphoric expressions: This period

To achieve this grounding functionality, the exemplary system utilizes a database containing fiscal year information for various companies. Some embodiment restrict tagging of time expressions to those greater than one month
 20 and those that are current relative to the publication date of the document. Also, if there are multiple valid time expressions the one closest to any monetary expression is tagged, and the other omitted unless there is a corresponding monetary expression. If there is a valid time expression, it is extracted Output of recognizers 530, which takes the form of tagged sentences or other text
 25 segments, is feed to sentence classifiers 540.

Sentence classifiers 540 (more generally text segment classifiers) include a set of classifiers for directing processing of the sentences or text segments to one or more of record or template filling modules within slot fillers 550. Specifically, sentence classifiers 540 includes an M&A (mergers & acquisitions)
 30 event classifier 542, a guidance event classifier 544, and an earnings event classifier 546.

M&A classifier 542 determines whether tagged and resolved sentences (or more generally text segments) from recognizers 530 include an M&A event. Within the exemplary embodiment, an M&A event is defined as a relation

between two companies and a money amount (or a percentage of stake). The two companies in an M&A event are the acquirer and the target. An M&A event also has a status (i.e., rumor, intended, announced, pending, completed, withdrawn). An example text containing an M&A event is shown below along with the corresponding structured event record (data structure) produced by M&A slot filler (relationship extractor) 552 and status classifier 558.

Sample Merger & Acquisition Text

Under the deal announced Thursday, Glu Mobile (GLUU) will pay about \$14.7 million in AGGREGATE VALUE to acquire Beijing Zhangzhong MIG Information Technology Co. Ltd.

Extracted Merger & Acquisition Template (Record)

Acquirer	GLU MOBILE
Target	BEIJING ZHANGZHONG MIG INFORMATION TECHNOLOGY CO.
Value	1.47E+07
Value Type	USD
Value Measure	AGGREGATE VALUE
Value Text	\$14.7 million
Status	ANNOUNCED

In the exemplary embodiment, creating a structured template given an input document involves identifying whether the document contains an M&A event and filling the template(s) with the correct entity information, such as company name, company IDs, or normalized money amount.

M&A classifier 542 is implemented using a semi-supervised machine-learning approach to determine which sentences have acquirer-target pairs of companies. Rules-based approach is then used to associate one or more merger valuation figures or values with the acquirer-target pair. M&A status classifier 558 determines a status for the M&A event. The exemplary embodiment implements classifier 558 using a semi-supervised machine learning approach.

The success of any supervised machine learning approach relies on having high quality training data. But training data requires the manual tagging

of hundreds of examples, and can therefore be expensive and time consuming to generate. To alleviate this bottleneck, the exemplary embodiment employs a framework for generating large amounts of training data semi-automatically from an unlabeled, time-stamped news corpus. Such methods are called ‘semi-supervised’, because they require less human intervention in the training process. Sometimes, multiple algorithms can be used to train each other (co-training) or high recall features can be used to train other features (surrogate learning). Based on a small set of 15 seed patterns (e.g., "acquisition of ORG"), we derived the training data from a large unlabeled news corpus. The training data was then used to learn models that identify the different pieces of information required to extract a structured record for each M&A event from the input document.

The minimal number of tagged entities which qualifies a sentence for inclusion in the candidate set is two company names. To help collect training data, the exemplary embodiment uses structured records from merger and acquisitions database on Westlaw® information-retrieval system (or other suitable information-retrieval system) to identify merger and acquisition events that have taken place in the recent past.

To efficiently identify positive training instances from the candidate set, the exemplary embodiment finds sentences that contain the names of entities that match these records and were published during the time frame over which the merging event took place. To identify negative instances, the exemplary embodiment selects sentences that contain companies known to not have been involved in a merger or acquisition. Once the system determines that a text segment includes an M&A event, the segment is passed to M&A event extractor 552 which copies or places identified entities and tagged expressions from a positively classified M & A change event sentence (text segment) into a structured template record that identifies the roles of the named entities and tagged expressions in the event.

Guidance event classifier 544 determines whether tagged and resolved sentences (or more generally text segments) from recognizers 530 include a guidance event. Within the exemplary embodiment, a guidance event is defined as a relation between a company, a complex money amount and a future time period. The complex money amount is called MONEX for our purposes and can contain a money amount (or range), the color of the money (e.g., earnings) and

possibly a trend (for example, up or down). An example of a guidance statement and the corresponding event template produced by guidance event extractor 554 is shown below.

5

Sample Guidance Text

CA boosted its **full-year 2008** forecast, now expecting **earnings of 87 cents to 91 cents a share** and revenue in the range of \$4.15 billion to \$4.2 billion.

(tagged terms or phrases are highlighted in bold.)

10

Extracted Guidance Template

Company	CA Inc.
Company Id	C000001193
Period	2008P1Y
Measure	EPS
Low Value	0.87
High Value	0.91
Value Type	Denominated
Currency	USD
Trend	Up

Because the language used in guidance events is somewhat formulaic, the exemplary guidance event classifier uses a rule-based approach to determine if a text segment includes a guidance event. One aspect of this determination is determining whether a time period tagged in the text segment is a future time period relative to a current time period or publication date associated with the document that contains the text segment. In addition, the color of the MONEX is determined. Earnings of \$10-\$12 a share describes a MONEX containing the following slots: [MinValue: 10, MaxValue: 12, Currency: USD, Measure: EPS]. Then, it identifies the respective company and the time period.

Earnings event classifier 546 determines whether tagged and resolved sentences (or more generally text segments) from recognizers 530 include an earnings event. The exemplary embodiment defines an earnings event as a relation between a company, a complex money amount and a past time period. The complex money amount is called MONEX for our purposes and can contain

a money amount (or range), the color of the money (e.g., earnings) and possibly a trend (e.g. up). An example of an earnings event and its corresponding structured record produced by earnings event extractor 556 are shown below.

5

Sample Earnings Text

Genpact Ltd., (G) the Gurgaon, India, manager of business processes for companies, reported **third-quarter** earnings rose 27% on 32% higher revenue. **Earnings** reached **\$16.3 million** from \$12.8 million in the year-earlier period.

10

Extracted Earnings Template

Company	Genpact
Company Id	C902357116
Period	2007F3Q
Measure	CSH
Value	1.63E+07
Value Type	Denominated
Currency	USD
Trend	Up

Similar to the processing of guidance events, the exemplary embodiment uses a rule-based approach to classify earnings events because the underlying language is generally formulaic. In some embodiments, the minimal number of tagged entities which qualifies a sentence for inclusion in the candidate set (that is, as potentially including an earnings event) is one company name and the phrase "net income" or the word "profit". To efficiently find positive instances, the exemplary embodiment extracts net income information from SEC documents for particular companies and finds positive candidates when the named company in the sentence and the dollar amount or percentage increase in profit for a time period line up with information from an SEC document. Negative instances are found when the data for a particular company does not line up with SEC filings. The earnings event extractor 556 (net income announcement event extractor) moves identified entities from a positively classified net income announcement

(earnings) event sentence into a structured template record. The template record identifies the roles the named entities and tagged phrases play in the event.

For a text segment to include a guidance or earnings event, some embodiments impose the rule that it must include at least one resolved company name which is not an analyst company (e.g., Thomson First Call, or
5 MarketWatch) and one monetary expression.

In addition to text segment classifiers 540 and relationship extractors (slot fillers 550, system 500 includes output modules 560.

10 Output modules 560 includes a database builder module 562 and a report builder module 564. Database builder module 562 builds a database from the event templates or records that are filled by relationship extractors 550, enabling one to, for example, to readily access the event data using conventional search .
Report builder

15

Exemplary Extraction of Information From Tables Found In Text

System 500 makes use of SEC filing data for example to determine timing, discern earnings trends, etc. To facilitate use of this data, the exemplary embodiment employs a novel system and methodology for extracting
20 information from tables found in the text of these documents. One component of the table-data extractions system is an SVM classifier (or another classifier similar in function) that distinguishes tables from non-tables. Tables that are only used for formatting reasons are identified as non-tables. In addition, tables are classified as tables of interest, such as background, compensation, etc. The
25 feature set comprises text before and after the tables as well as n-grams of the text in the table. The tables of interest are then processed according to the following:

1) label/value detection. The table has to be partitioned in the labels and the values. For the exemplary table below, the system determines that the
30 money amounts are values and the rest are labels;

2) label grouping. Some labels are grouped together. For example, Eric Schmidt and his current position are one label. On the other hand, a table that contains a year and a list of term names (i.e. Winter, Spring, Fall) are not grouped together;

3) abstract table derivation. A derived Cartesian coordinate system leads to the notation that defines every value accordingly. [Name and Principal Position.Eric Schmidt Chairman of the Executive Committee and Chief Executive Officer.Year.2005, Annual Compensation.Salary(\$)] = 1;

- 5 4) relation extraction. Given the abstract table representation, the desired relations are derived. The compensation relation, for example, is filled with: NAME: Eric Schmidt; COMPENSATION TYPE: salary; AMOUNT: 1; CURRENCY: \$. Finally, an interpreter for the tables of interest is created. The input to the interpreter is a table and the output is a list of relations represented by the table.
- 10

Name and Principal Position	Year	Annual Compensation		
		Salary(\$)	Bonus(\$)	other Annual Compensation(\$)
Eric Schmidt Chairman of the Executive Committee and Chief Executive Officer	2005	1	1,630	24,741
	2004	81,432	1,556	0

- For the exemplary embodiment, we downloaded hundreds of documents from Edgar database (EDGAR) and annotated 150 of them for training and evaluation. We converted the documents into XHTML using Tidy (Raggett) before annotating them.
- 15

Name and Principal Position	Fiscal Year	Annual Compensation			Long-Term Compensation Awards Options (#)	All Other Compensation (\$)(2)
		Salary(\$)	Bonus(\$)(1)	Other Annual Compensation (\$)		
John T. Chambers	2005	350,000	1,300,000	0	1,500,000	8,977
President, Chief Executive Officer and Director	2004	1	1,900,000	0	0	0
	2003	1	0	0	4,000,000	0
Mario Mazzola	2005	447,120	557,737	0	600,000	7,424
Former Senior Vice President, Chief Development Officer (3)	2004	464,317	666,850	0	600,000	5,726
	2003	447,120	764,897	0	500,000	2,905
Charles H. Giancarlo						
...						

Table 3: A compensation table

Our information extraction system for genuine tables involve the following processes:

1. table classification
2. label row and column classification
3. table structure recognition
4. table understanding

Process 1, which enhances efficiency, entails identifying tables that have a reasonable chance of containing the desired relation before other more computationally expensive processes are applied. The tables containing the desired information are quickly identified using relation-specific classifiers based on supervised machine learning.

Process 2 entails distinguish between label column and label rows from values inside those tables. This time, the same supervised machine learning approach is used, but the training data is different from that in Step 1.

In process 3, after those label rows and label column are identified, an elaborate procedure is applied to these complex tables to ensure that semantically coherent labels are not separated into multiple cells, or multiple distinct labels are not squashed into a cell. The goal here is to associate each value with their labels in the same column and the same row. The result of the Step 3 is a list of attribute-value pairs.

In process 4, a rule-based inference module goes through each of the attribute-value pairs and identifies the desirable ones to populate the officers and directors database.

The exemplary embodiment makes use of an annotation in performing the supervised learning employed in both process 1 and process 2. To make the exemplary system more robust against lexical variations and table variations, supervised machine learning is used in processes 1 and 2. In supervised learning, one of the most challenging and time-consuming tasks is to obtain the labeled examples. To facilitate reuse across different domains, the exemplary embodiment uses a scheme that reduces or minimizes the human annotation effort needed.

For the tables containing the desired information, the exemplary embodiment uses the following annotations:

1. isGenuine: a flag indicates that this is a genuine table or a non-genuine table.
2. relations: the relations that a table contain, such as “name+title”, “name+age”, name+year+salary” or “name+year+bonus”, or a combination of them.
3. isContinuous: a flag indicates that if this table is a continuation of the previous genuine table.
4. lastLabelRow: the row number of the last label row.
5. lastLabelColumn: the column number of the last label column associated with each relation.
6. valueColumn: the number of the column that contains the desired values for each relation.

The specified relations are used as training instances to build models for process 1. The information lastLabelRow and lastLabelColumn are used to build models to classify rows and column as labels rows or columns in process 2. In our guideline to annotators, we specifically ask them to annotate the column number of the last label column for each relation.

The need for such fine-grained annotation is best illustrated using an example. In Table 3, for relation “name+title”, the last label column is 1, the column “name and principal position”. But for relation “name+year+bonus”, the last label column is 3, “fiscal year”. For extracting multiple relations in a table, these relations might share the same last label column, but this is not always the case. As a result, there is a need to annotate the associated label column for each relation separately. The flag isContinuous indicates if the current table is a continuation of the previous table. If it is, the current table can “borrow” the boxhead from previous table since such information is missing. The exemplary embodiment eliminates tables marked with “isContinuous” flag during training, but kept those tables during evaluation. The annotation valueColumn can be used for automatic evaluation in the future.

There are a few rare instances where the default arrangement of boxhead and stub, as shown in Table 3, are swapped in the corpus. Currently in our annotation, we simply don’t supply “valueColumn” for the relations since they don’t apply. For table classification and table understanding tasks, this is not of

much an issue, but the above annotation scheme would need to be further modified to capture such difference.

Table classification: The exemplary embodiment classifies or screens tables based on whether they are likely to include desired relational information before attempting detailed extraction processes. To identify tables that contain desired relations, we employed LIBSVM (Chang & Lin 2001), a well-known implementation of support vector machine. Based on the annotated tables, a separate model is trained for each desired relation. In the SEC domain, a table might contain multiple relations.

Exemplary features for use in the SVM include:

- top 1000 words inside tables in the corpus, and top 200 words in text preceding the tables. These thresholds are based on experiments using LIBSVM 5-fold cross validation. A stop word list was used.
- number of words in tables that are label words
- number of cells containing single word
- number of cells containing numbers
- maximum cell string size
- number of names
- number of label words in the first row

The exemplary embodiment then uses a model for each desired relation. Because “name+year+salary” and “name+year+bonus” cooccur 100% of the time in the annotated corpus, the same classifier was for both relations. In this domain, the number of negative instances is significantly larger than positive instances, perhaps because having both signature tables and tables containing background information in sentences format create significant overlap between positive and negative instances. To address this, the exemplary embodiment only uses a subset of negative instances for training (75% of our training instance are negative instances). We also trained a separate module to distinguish between a genuine and non-genuine tables based on annotated data. This second model is relation independent. The feature set is similar to the feature set outlined above.

To identify which words are likely to be names, we downloaded the list of names from (U.S. Census Bureau). The list of names is further filtered by removing the common words, such as “white”, “cook”, or “president”, based on a English word list (Atkinson August 2004). Although it is feasible to use a list

of common title words, the exemplary embodiment does not use such information so that it may operate more readily across other domains. However, in embodiment that do use such a domain-specific list, this information would probably significantly improve the precision and recall for extracting relation
 5 “name+title”.

Label row and column classification: Based on the annotated data, LIBSVM is again used to classify which rows belong to boxhead and which columns belong to stub. The training data for the models are words in the desired tables that were manually identified as box-head and stubs by using
 10 lastLabelRow and lastLabelColumn features. Other features used include the frequency of label words, the frequency of name words, and frequency of numbers.

For each relation, the exemplary embodiment uses a different label column classifier, since the lastColumnLabel might differ between different
 15 relations, as explained in the Annotation Section.

Table structure recognition: Because tables in the SEC filings are somewhat complex and formatted for visual purpose, a significant amount of effort is needed to normalize the table to facilitate later operations. Once label rows and columns are identified, several normalization operations are carried
 20 out:

1. create duplicate cells based on rowspan and colspan
2. merge cells into coherent label cells
3. identify subheadings
4. split specific column based on conjoin marker, such as “and” or
 25 parenthesis (before last label column)
5. split cells containing multiple labels, such as years “2005, 2006, 2007”

Step 1 specifically addresses the issue with the use of colspan and rowspan in HTML table, as have been done in (Chen, Tsai, & Tsai 2000). In Table 3, without copying the original labels into spanning cells, the label “annual
 30 compensation” would not be attached to the value “1,300,000” using just the HTML specification. By doing this step, we only need to associate all the labels in the box-head in that particular column to the value and ignore other columns.

In Step 2, we use certain layout information, such as underline, empty line, or background color, to determine when a label is really complete. In SEC

filings, there are many instances where a label is broken up into multiple cells in the boxhead or stub. In those cases, we want to recreate the semantically meaningful labels to facilitate later relation extraction – a process that is heavily dependent on the quality of the labels attached to the values. For example, in

5 Table 3, based on the separate in row 5, cells “John T. Chambers”, “President, Chief Executive”, and “Officer and Director” are merged into one cell, with line break marker (#) inserted into the original position. The new cell is “John T. Chambers#President, Chief Executive#Officer and Director”, and it is stored in cell on row 2, and copied to cells on row 3 and 4.

10 In Step 4, heuristic rules were applied to identify subheader. For example, if there is no value in the whole row except for the first label cell, then that label cell is classified as subheader. The subheader label is assigned as part of the label to every cell below it until a new subheader label cell is encountered.

Step 5 splits certain columns into multiple columns to ensure that a value

15 cell does not contain multiple values. For example, in Table 3, the first cell in first column is “name and principal position”. The system detects the word “and” and split the column into two columns, “name” and “principal position”, and do similar operations to all the cells in the original column. Remember in Step 3, cell on row 2 is the result of merge 3 cells, with line break markers between the

20 string in the original cells. By default, we use the first line break marker to break the merged cell into two cells. After this transformation, we have “John T. Chambers” and “President, Chief...” that corresponding to “name” and “principal position”. This type of operation is not only limited to “and”, but also to certain parenthesis, “Nondirector Executive Officer (Age as of February 28, 2006)”.

25 Such cells are broken into two, and so are the other cells in the same column.

Step 6 deals with repeated sequences in last label column. In Table 3, we are fortunate that all the cells under “fiscal year” contains only 1 value. There are instances in our corpus that such information is represented inside the same cell with line break between each value. In such cases, there are no lines between

30 these values, and the resulting table looks cleaner and thus visually more pleasing. It is certainly incorrect to assign all 3 years “2005, 2004, 2003” to the cell containing bonus information “1,300,000”. To address this, our system performs repeated sequence detection on all last label columns. If a sequence pattern, which doesn’t always have to be exactly the same, is detected, the

repeated sequence are broken into multiple cells so that each cell can be assigned to the associated value correctly.

Transforming a normalized table to Wang's representation (Wang 1996) is a trivial process. Given a value cell at (r,c), all the label cells in column (c) and row (r) are its associated labels. In addition, the labels in stub might also have additional associated labels in the boxhead, and those should be associated with the value cell also. For example, the value "1,300,000" will have following 4 associated labels: [annual compensation|bonus\$(1)], [fiscal year|2005], [principal position|president, chief executive officer and director], [name|John T. Chambers]. The characters "|" inside those associate labels indicate hierarchical relation between the labels. For tables with subheading, the subheading labels have already been inserted into all the associated labels in the stubs earlier.

Table understanding: Similar to (Gatterbauer et al. 2007), we consider IE from Wang's model requires further intelligent processing. To populate database based on Wang's representation, a rule-based system is used. We specifically look for certain patterns, such as "name", "title" or "position" in the associated labels in order to populate the "name-title" relation. For different relations, a different set of patterns is used. It's important to perform error analysis at this stage to detect ineffective patterns. For example, several tables with "name-title" information used the phrase "nondirector executive officer" instead of the label for "name". Clearly, we can apply supervised machine learning to make the process more robust. In our annotation, we have asked the annotators to identify the columns that contains the information we want in valueColumn. Such information might be used to train our table understanding module in the future.

The following procedures can be used to tailor our approach to a new application or domain:

- Collect a corpus and annotate the tables with the desired information as described in the Annotation section.
- Modify features to take advantage of knowledge in the new domain.
- Train all the classifiers. Depending on the size of the corpus, different thresholds can be specified to minimize the size of the vocabulary, which is used as features. This training process can be automated.

- Modify table normalization to take advantage of domain knowledge.

For example, in SEC domain, separating the label cell “name and title” is applied in order to simplify later relation extraction operations.

- Modify relation extraction rules. Different relations are signaled by different words in the labels. Currently, we manually specify these rules.
- This process is designed to maximize precision and recall while minimizing the annotation effort. Each component can be modified to take advantage of the domain specific information to improve its performance.

10 Exemplary Generation of Sentence Paraphrases

An additional embodiment of the present invention includes a tool that generates sentence paraphrases starting from the seed templates provided by a user. The tool takes sentences that indicate an event with high precision with the actual entities replaced by their generic types, for example:

- 15 **<ORG> bought <ORG>**
 <ORG>'s merger with <ORG>

- The sentence is searched for in a corpus and actual entity identities are obtained from sentences conforming to the seed pattern. Then other sentences mentioning the same entities in the corpus are located and these serve as which serves as paraphrases for the initial sentence. (In some embodiments, the other sentences are restricted to those occurring within a narrow time window). Each one of these other sentences can then be treated as a seed template or pattern by removing the named entities and then repeating the search for other sentences that conform to this new seed pattern. The sentences can be ordered according to frequencies of component phrases and manually checked to generate gold data for the classifiers.

Conclusion

- The embodiments described above are intended only to illustrate and teach one or more ways of practicing or implementing the present invention, not to restrict its breadth or scope. The actual scope of the invention, which embraces all ways of practicing or implementing the teachings of the invention, is defined only by the issued claims and their equivalents.

CLAIMS

1. A computer system having at least one processor and at least one memory, the system comprising:
 - 5 means for automatically tagging entity names, monetary expressions, and temporal expressions within a text segment;
 - means for determining whether the automatically tagged text segment describes a financial event; and
 - 10 means, responsive to an affirmative determination of a financial event in the tagged text segment, for defining in memory a data record including data derived from the tagged text segment.
2. The system of claim 1, wherein the text segment is a grammatical sentence.
- 15 3. The system of claim 1, wherein the data record includes:
 - a company field including text identifying a named entity tagged in the text segment;
 - a company ID field including an alphanumeric code identifying the named
 - 20 entity; and
 - a time period field including an alphanumeric code identifying a financial reporting period.
4. The system of claim 1, wherein the data record includes a field indicating
25 whether a monetary expression tagged in the text segment is trending up or down.
5. The system of claim 1, wherein the data record includes a field indicating
30 that a monetary expression tagged in the text segment is a measure of earnings per share.

6. The system of claim 1, wherein the means for automatically tagging entity names, monetary expressions, and temporal expressions within a text segment includes:
 - 5 first means for tagging and resolving entity names;
 - second means for tagging monetary expression; and
 - third means for tagging temporal expressions.
7. The system of claim 1, wherein the means for determining whether the
10 automatically tagged text segment describes a financial event includes a mergers & acquisition (M&A) classifier for determining whether text segments describe an M&A event.
8. The system of claim 1, wherein the means for determining whether the
15 automatically tagged text segment describes a financial event includes a mergers & acquisition (M&A) classifier for determining whether or not text segments describe an M&A event.
9. The system of claim 8, wherein the M&A classifier is machine-learning
20 based classifier.
10. The system of claim 1, wherein the means for determining whether the
automatically tagged text segment describes a financial event includes a
an earnings event classifier for determining whether or not the text
25 segment describes an earnings event.
11. The system of claim 1, wherein the means for determining whether the
automatically tagged text segment describes a financial event includes a
guidance event classifier for determining whether or not the text segment
30 describes a financial guidance event.

12. A computer implemented method comprising:
automatically tagging entity names, monetary expressions, and temporal
expressions within a text segment;
5 determining whether the automatically tagged text segment describes a
financial event; and
in response to an affirmative determination of a financial event in the
tagged text segment, defining in memory a data record including
data derived from the tagged text segment.
- 10 13. The method of claim 12, further comprising displaying on a display
device at least a portion of the data record in association with a user
selectable command feature of a graphical user interface for causing
retrieval of a document including the text segment.
- 15 14. The method of claim 12, wherein the text segment is a grammatical
sentence.
- 20 15. The method of claim 12, wherein the data record includes:
a company field including text identifying a named entity tagged in the
text segment;
a company ID field including an alphanumeric code identifying the named
entity; and
a time period field including an alphanumeric code identifying a
25 financial reporting period.
- 30 16. The method of claim 12, wherein the data record includes a field
indicating whether a monetary expression tagged in the text segment is
trending up or down.
automatically tagging entity names within a text segment as being one of
a person, company, and location; and
automatically associating one or more of the tagged entity names with an
entry in a data set of named entities.

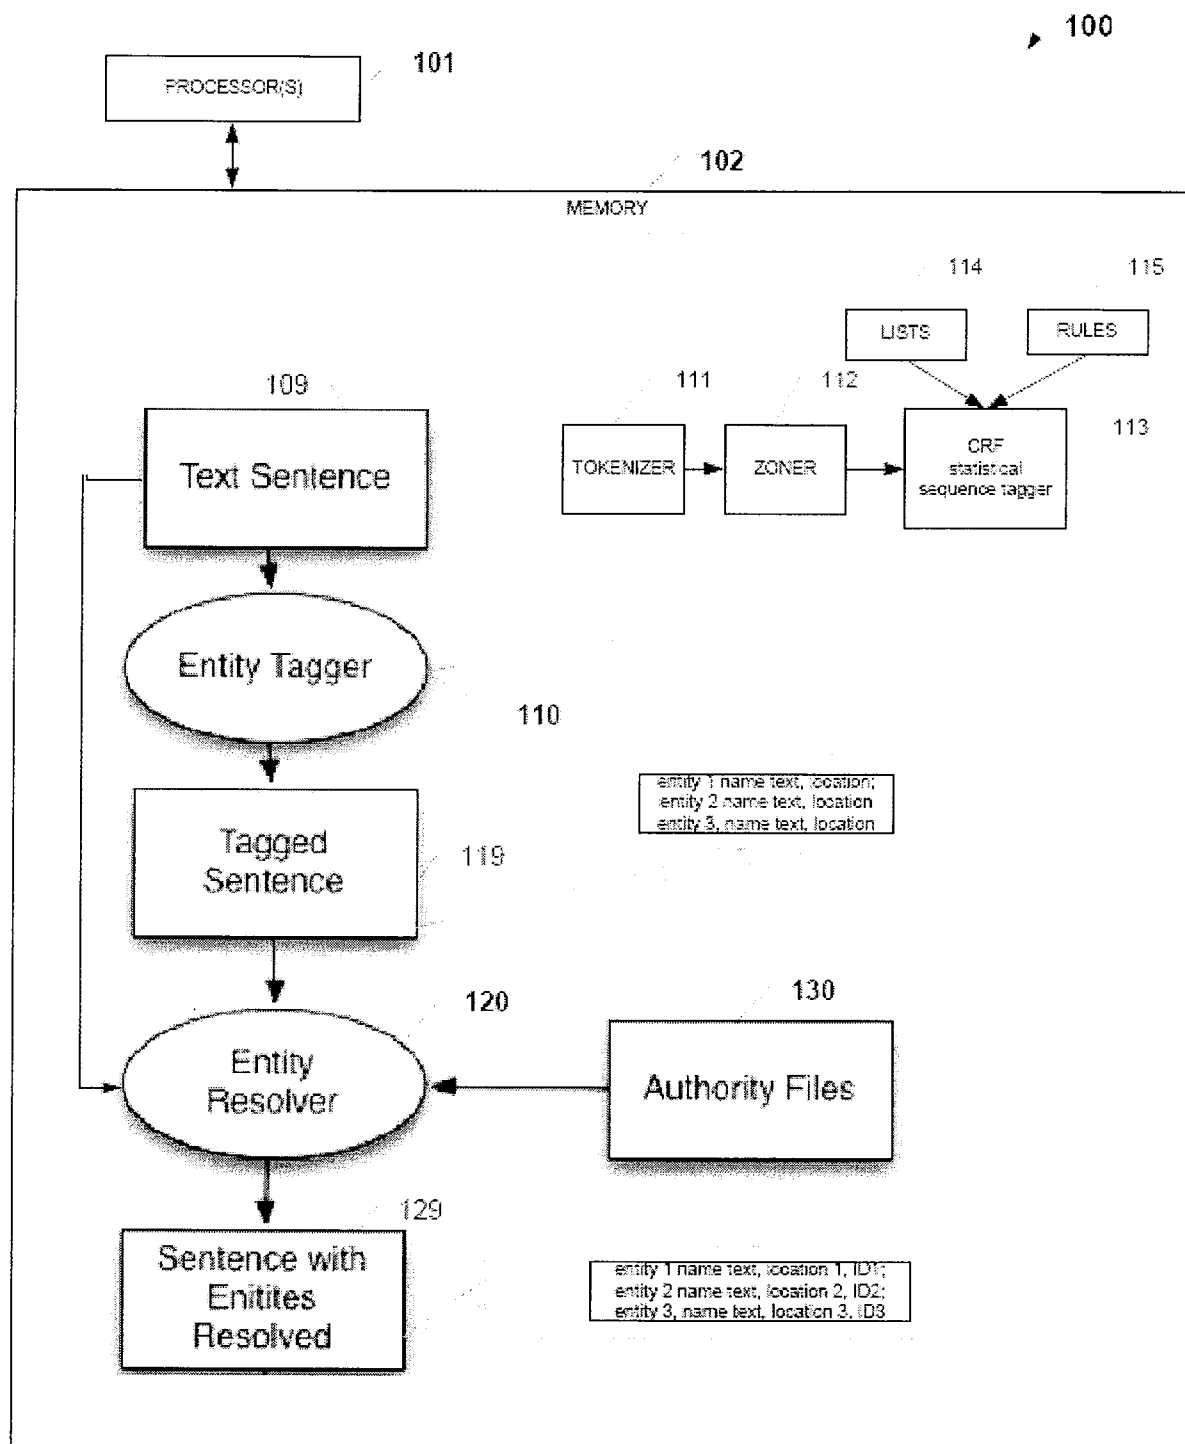
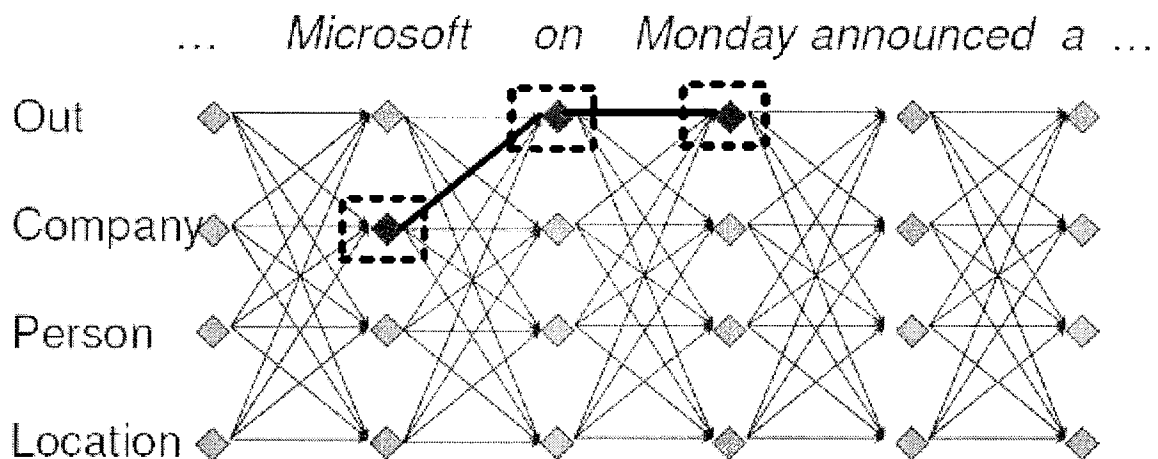


FIGURE 1



- **Microsoft** is in the high precision list of company names
- **Monday** is in a list of tokens that should always be marked as Out
- A rule says that **on** should be marked as Out if it is followed by an Out

FIGURE 2

300

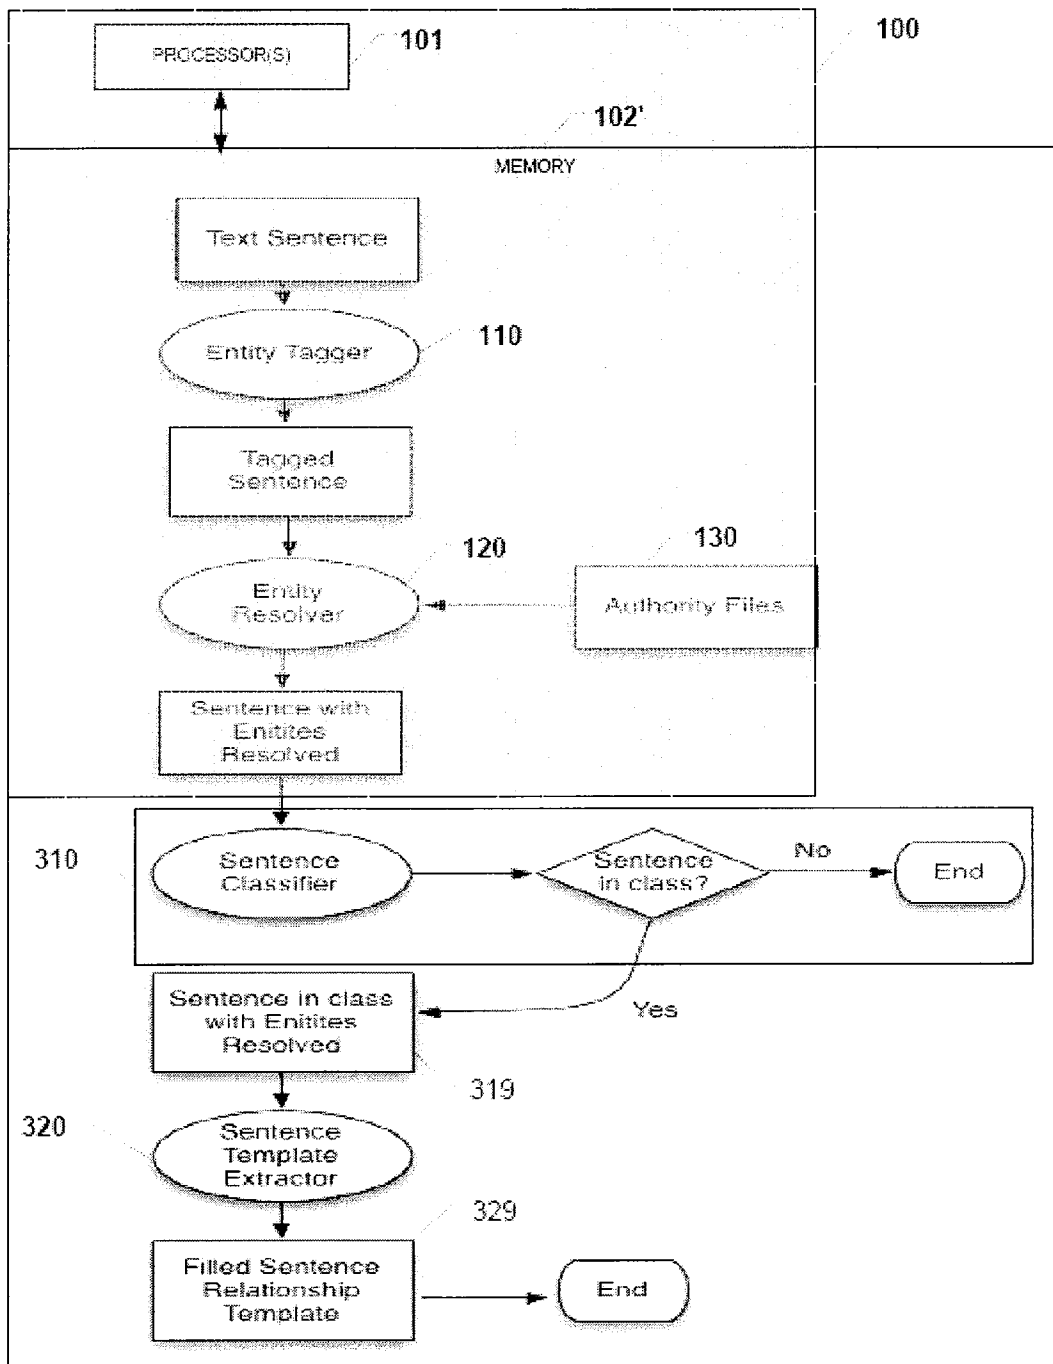


FIGURE 3

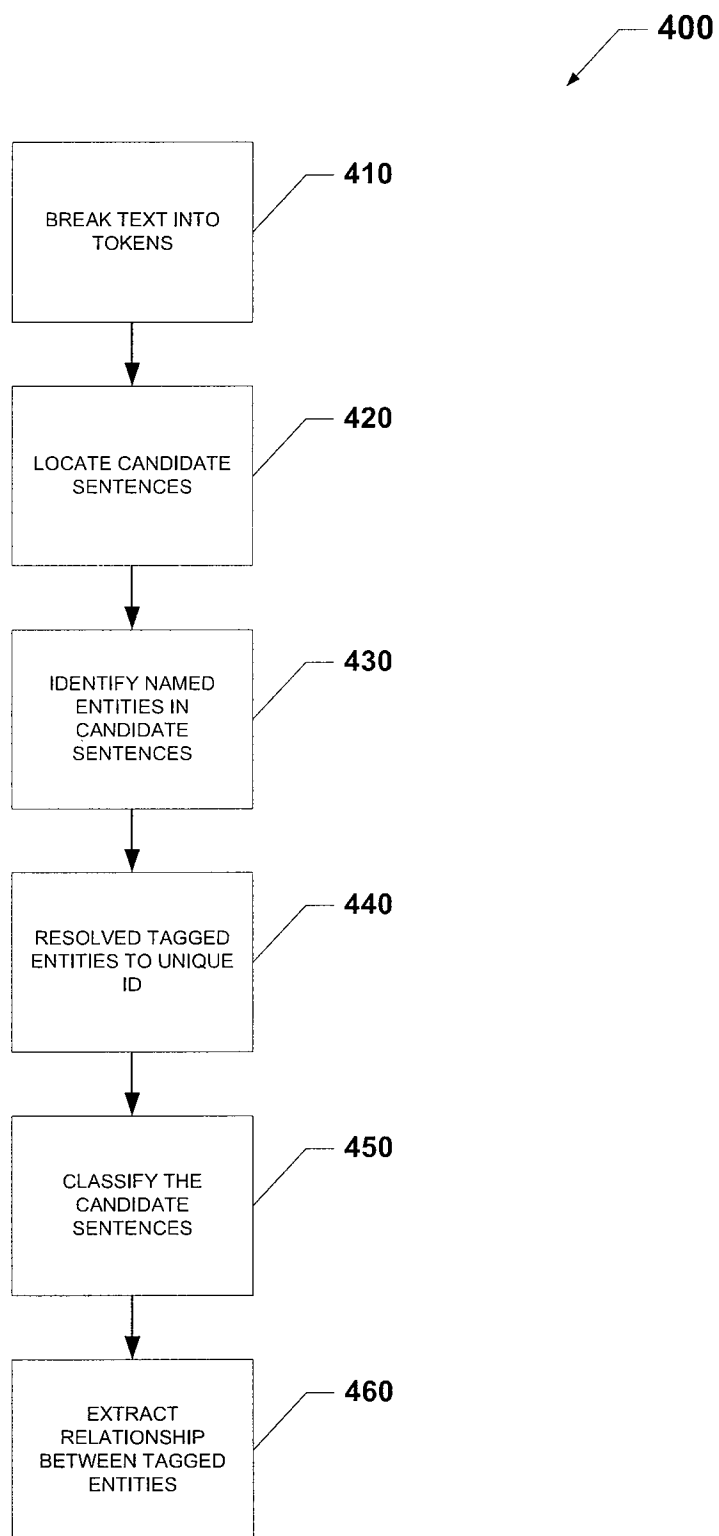


FIGURE 4

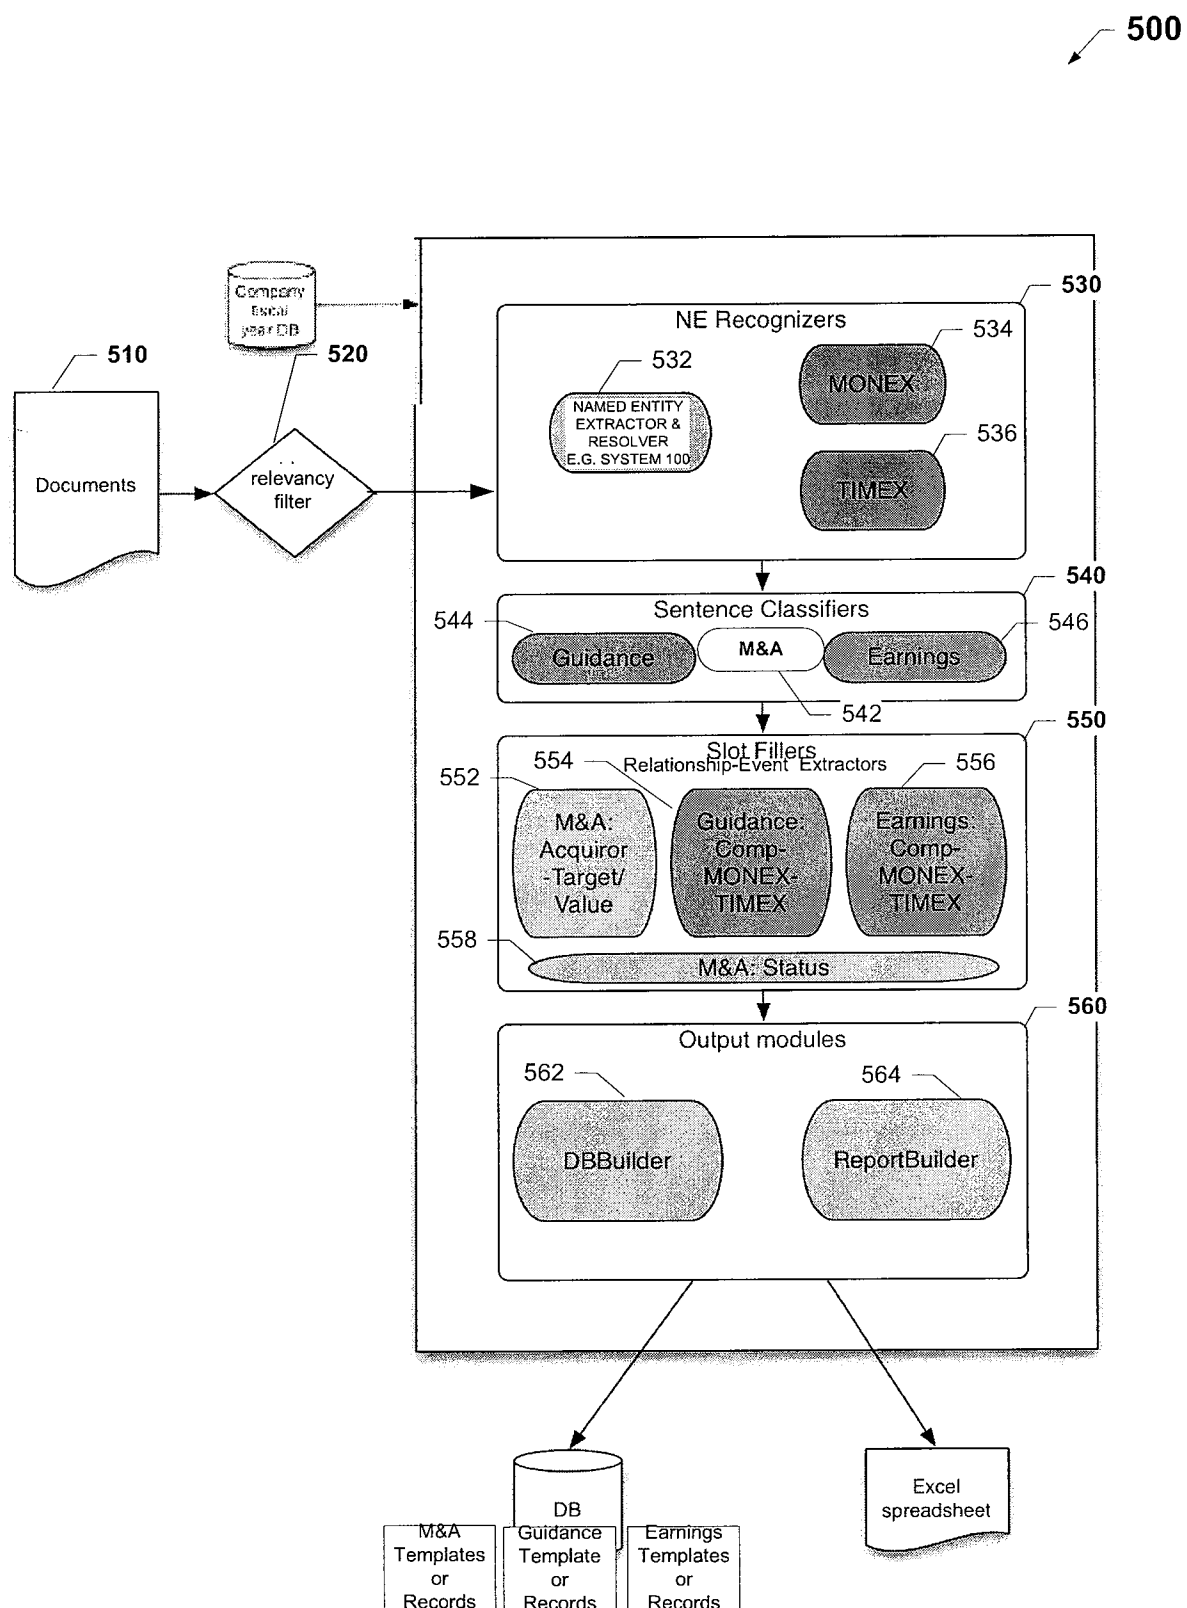


FIGURE 5