



US 20140229478A1

(19) **United States**(12) **Patent Application Publication**
Knight et al.(10) **Pub. No.: US 2014/0229478 A1**(43) **Pub. Date: Aug. 14, 2014**(54) **COMPUTER-IMPLEMENTED SYSTEM AND METHOD FOR PROVIDING VISUAL CLASSIFICATION SUGGESTIONS FOR INCLUSION-BASED CONCEPT CLUSTERS**

(60) Provisional application No. 61/229,216, filed on Jul. 28, 2009, provisional application No. 61/236,490, filed on Aug. 24, 2009.

Publication Classification(71) Applicant: **FTI Consulting, Inc.**, Annapolis, MD (US)(51) **Int. Cl.**
G06F 17/30 (2006.01)(72) Inventors: **William C. Knight**, Bainbridge Island, WA (US); **Nicholas I. Nussbaum**, Seattle, WA (US); **John W. Conwell**, Tukwila, WA (US)(52) **U.S. Cl.**
CPC **G06F 17/3071** (2013.01); **G06F 17/30713** (2013.01)USPC **707/732; 707/737**(73) Assignee: **FTI Consulting, Inc.**, Annapolis, MD (US)(57) **ABSTRACT**(21) Appl. No.: **14/252,688**

A computer-implemented system and method for providing visual classification suggestions for inclusion-based concept clusters are provided. Reference concepts each associated with a classification code are designated. One or more of the reference concepts are grouped with a plurality of uncoded concepts into a grouped concept set. Clusters are generated, each including a portion of the uncoded concepts and the reference concepts of the grouped concept set. A visual suggestion for assigning one of the classification codes to one of the clusters including visually representing each of the reference concepts in that cluster is provided.

(22) Filed: **Apr. 14, 2014****Related U.S. Application Data**

(63) Continuation of application No. 12/844,810, filed on Jul. 27, 2010, now Pat. No. 8,700,627.

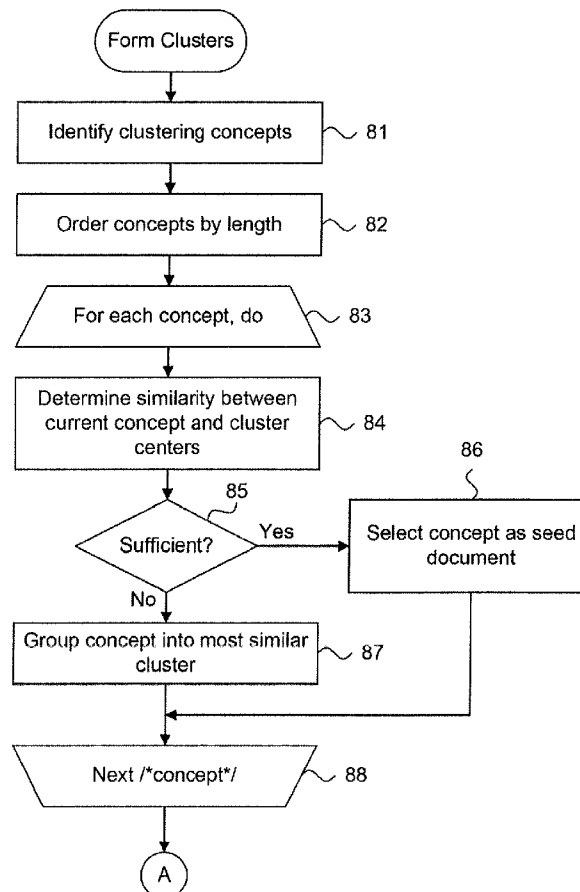
80

Fig. 1.




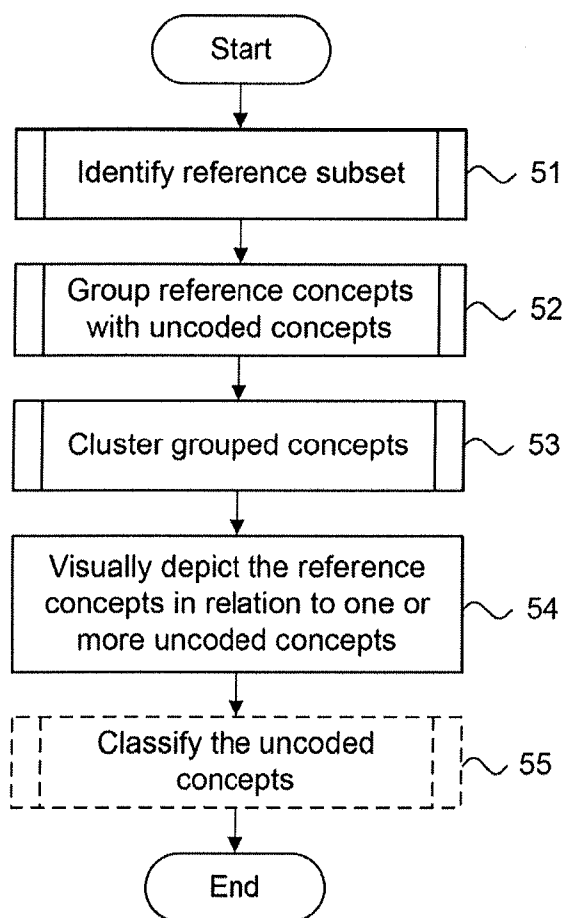
Fig. 2.50

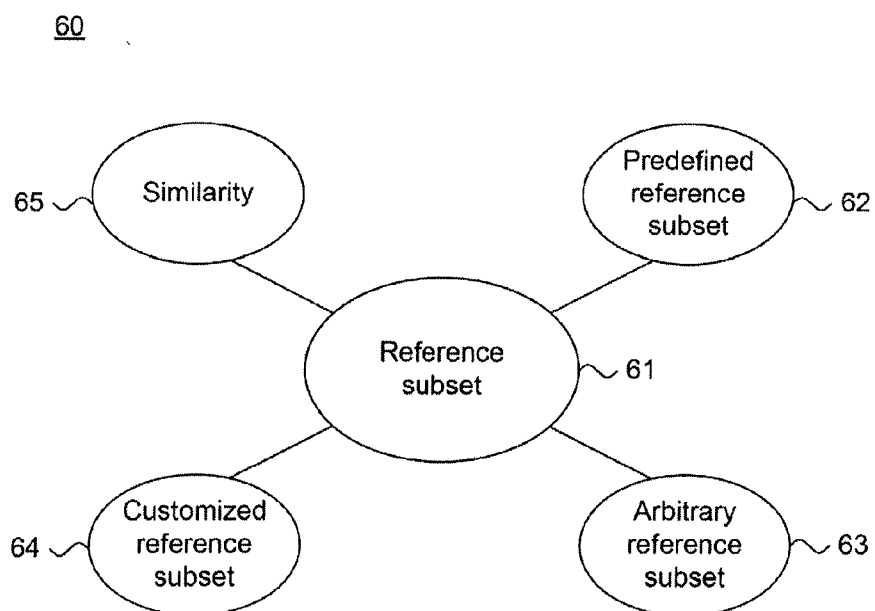
Fig. 3.

Fig. 4.

70

		73						Documents (71)
		D ₁	D ₂	D ₃	.	.	.	
74	C ₁	2	12	4	77			
	C ₂	5	35	2				
	C ₃	21	0	2				
	.	18						
	.							
	C _m							
Concepts (72)		146						75

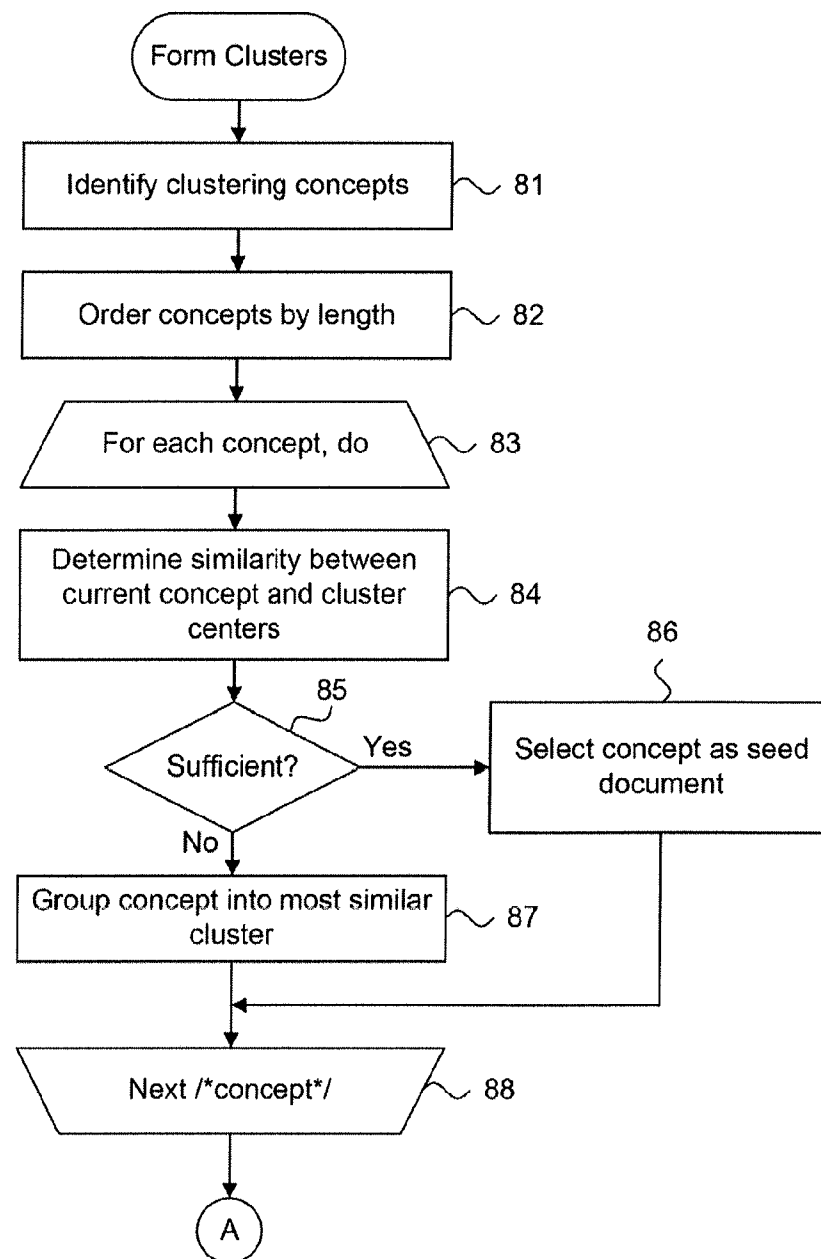
Fig. 5.80

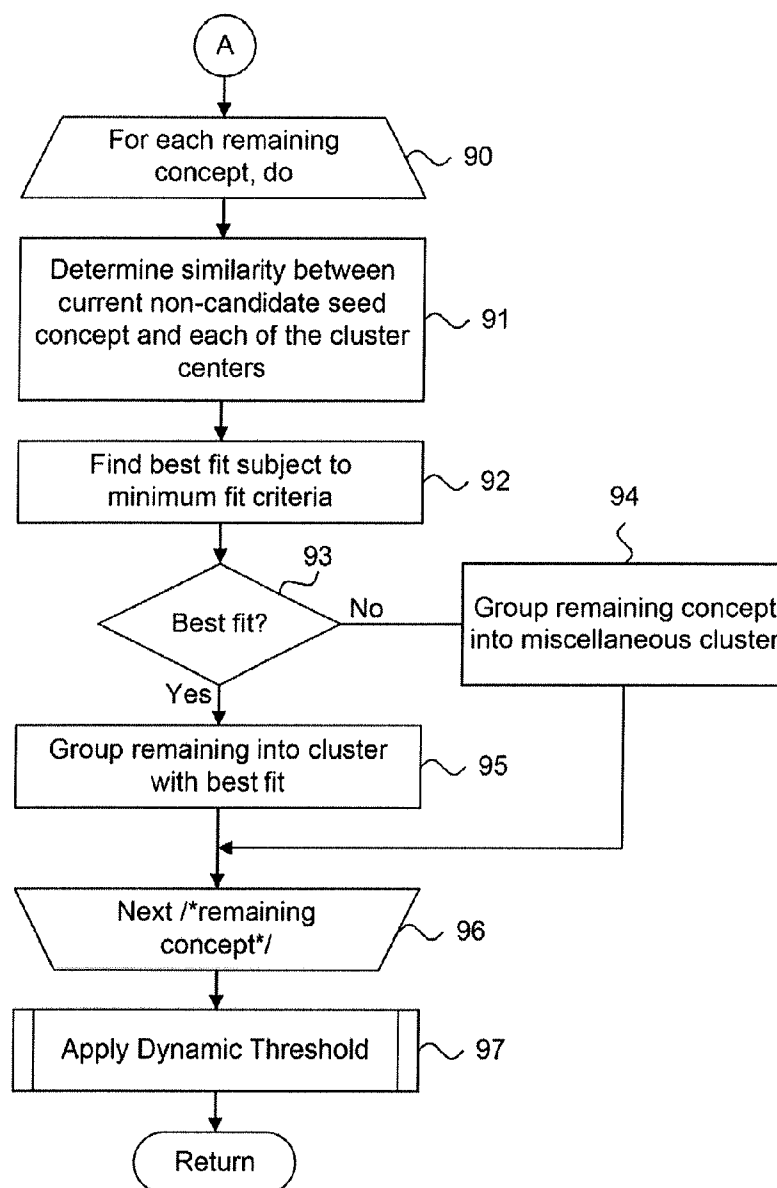
Fig. 5 (Cont).

Fig. 6.




Fig. 7A.




Fig. 7B.




Fig. 7C.




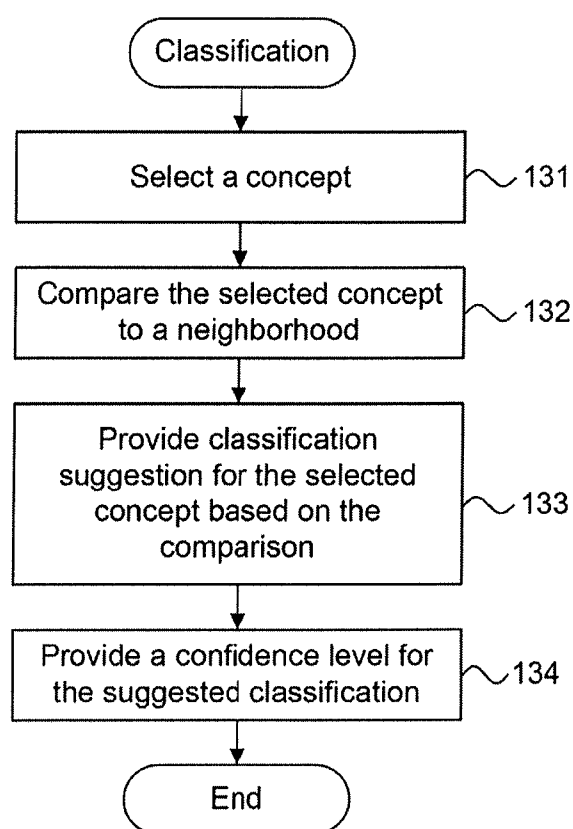
Fig. 8.130

Fig. 9.140

140

142

143

144

145

146

147

148

149

Inject...

Reference Source

Use Marked Documents in Matter ▼

Acme Training Set Binder

Edit Binders...

Reference Filter

Filter by Category ☒ ☒ ☒ ☒ ☒ ☒ ☒ ☒ ☒ ☒ ☒

Filter by Criteria ▼

Command Details

Target All Clusters ▼

Action Insert In Map until Recluster ▼

Count 0 + 0 %

Automatic ☒ On Assignment Open ▼

Document Suggest ☒

Cluster Suggest ☒

Exclude Duplicates ☐

Clean Prior Injections ☒

Add As References ☒

Advanced

Similarity 0 Cluster Sites Centroid ▼

Nearest Neighbor ▼

Classifier Parameters

ClassEstimator	DistanceWeightedMaximumCount
ProbabilityEstimator	Linear
ProbabilityEstimatorBeta	8

Commands

Inject Cancel

**COMPUTER-IMPLEMENTED SYSTEM AND
METHOD FOR PROVIDING VISUAL
CLASSIFICATION SUGGESTIONS FOR
INCLUSION-BASED CONCEPT CLUSTERS**

**CROSS-REFERENCE TO RELATED
APPLICATION**

[0001] This patent application is a continuation of U.S. patent application Ser. No. 12/844,810, filed Jul. 27, 2010, pending, which is a non-provisional patent application which claims priority under 35 U.S.C. §119(e) to U.S. Provisional Patent Application Ser. No. 61/229,216, filed Jul. 28, 2009, and U.S. Provisional Patent Application Ser. No. 61/236,490, filed Aug. 24, 2009, the disclosures of which are incorporated by reference.

FIELD

[0002] This application relates in general to using documents as a reference point and, in particular, to a system and method for displaying relationships between concepts to provide classification suggestions via inclusion.

BACKGROUND

[0003] Historically, document review during the discovery phase of litigation and for other types of legal matters, such as due diligence and regulatory compliance, have been conducted manually. During document review, individual reviewers, generally licensed attorneys, are assigned sets of documents for coding. A reviewer must carefully study each document and categorize the document by assigning a code or other marker from a set of descriptive classifications, such as "privileged," "responsive," and "non-responsive." The classifications can affect the disposition of each document, including admissibility into evidence.

[0004] During discovery, document review can potentially affect the outcome of the underlying legal matter, so consistent and accurate results are crucial. Manual document review is tedious and time-consuming. Marking documents is solely at the discretion of each reviewer and inconsistent results may occur due to misunderstanding, time pressures, fatigue, or other factors. A large volume of documents reviewed, often with only limited time, can create a loss of mental focus and a loss of purpose for the resultant classification. Each new reviewer also faces a steep learning curve to become familiar with the legal matter, classification categories, and review techniques.

[0005] Currently, with the increasingly widespread movement to electronically stored information (ESI), manual document review is no longer practicable. The often exponential growth of ESI exceeds the bounds reasonable for conventional manual human document review and underscores the need for computer-assisted ESI review tools.

[0006] Conventional ESI review tools have proven inadequate to providing efficient, accurate, and consistent results. For example, DiscoverReady LLC, a Delaware limited liability company, custom programs ESI review tools, which conduct semi-automated document review through multiple passes over a document set in ESI form. During the first pass, documents are grouped by category and basic codes are assigned. Subsequent passes refine and further assign codings. Multiple pass review requires a priori project-specific knowledge engineering, which is only useful for the single

project, thereby losing the benefit of any inferred knowledge or know-how for use in other review projects.

[0007] Thus, there remains a need for a system and method for increasing the efficiency of document review that bootstraps knowledge gained from other reviews while ultimately ensuring independent reviewer discretion.

SUMMARY

[0008] Document review efficiency can be increased by identifying relationships between reference documents and uncoded documents and providing a suggestion for classification based on the relationships. The reference documents and uncoded documents are clustered based on a similarity of the documents. The clusters and the relationship between the uncoded documents and reference documents within the cluster are visually depicted. The visual relationship of the uncoded documents and reference documents provide a suggestion regarding classification for the uncoded documents.

[0009] One embodiment provides a computer-implemented system and method for providing visual classification suggestions for inclusion-based concept clusters. Reference concepts each associated with a classification code are designated. One or more of the reference concepts are grouped with a plurality of uncoded concepts into a grouped concept set. Clusters are generated, each including a portion of the uncoded concepts and the reference concepts of the grouped concept set. A visual suggestion for assigning one of the classification codes to one of the clusters including visually representing each of the reference concepts in that cluster is provided.

[0010] Still other embodiments of the present invention will become readily apparent to those skilled in the art from the following detailed description, wherein are described embodiments by way of illustrating the best mode contemplated for carrying out the invention. As will be realized, the invention is capable of other and different embodiments and its several details are capable of modifications in various obvious respects, all without departing from the spirit and the scope of the present invention. Accordingly, the drawings and detailed description are to be regarded as illustrative in nature and not as restrictive.

BRIEF DESCRIPTION OF THE DRAWINGS

[0011] FIG. 1 is a block diagram showing a system for displaying relationships between concepts to provide classification suggestions via inclusion, in accordance with one embodiment.

[0012] FIG. 2 is a process flow diagram showing a method for displaying relationships between concepts to provide classification suggestions via inclusion, in accordance with one embodiment.

[0013] FIG. 3 is a block diagram showing, by way of example, measures for selecting reference concept subsets for use in the method of FIG. 2.

[0014] FIG. 4 is a table showing, by way of example, a matrix mapping of uncoded concepts and documents.

[0015] FIG. 5 is a process flow diagram showing, by way of example, a method for forming clusters for use in the method of FIG. 2.

[0016] FIG. 6 is a screenshot showing, by way of example, a visual display of reference concepts in relation to uncoded documents.

[0017] FIG. 7A is a block diagram showing, by way of example, a cluster with “privileged” reference concepts and uncoded concepts.

[0018] FIG. 7B is a block diagram showing, by way of example, a cluster with “non-responsive” reference concepts and uncoded concepts.

[0019] FIG. 7C is a block diagram showing, by way of example, a cluster with uncoded concepts and a combination of differently classified reference concepts.

[0020] FIG. 8 is a process flow diagram showing, by way of example, a method for classifying uncoded concepts for use in the method of FIG. 2.

[0021] FIG. 9 is a screenshot showing, by way of example, a reference options dialogue box for entering user preferences for clustering concepts.

DETAILED DESCRIPTION

[0022] The ever-increasing volume of ESI underlies the need for automating document review for improved consistency and throughput. Token clustering via injection utilizes reference, or previously classified tokens, which offer knowledge gleaned from earlier work in similar legal projects, as well as a reference point for classifying uncoded tokens.

[0023] The tokens can include word-level, symbol-level, or character-level n-grams, raw terms, entities, or concepts. Other tokens, including other atomic parse-level elements, are possible. An n-gram is a predetermined number of items selected from a source. The items can include syllables, letters, or words, as well as other items. A raw term is a term that has not been processed or manipulated. Entities further refine nouns and noun phrases into people, places, and things, such as meetings, animals, relationships, and various other objects. Additionally, entities can represent other parts of grammar associated with semantic meanings to disambiguate different instances or occurrences of the grammar. Entities can be extracted using entity extraction techniques known in the field. Concepts are collections of nouns and noun-phrases with common semantic meaning that can be extracted from ESI, including documents, through part-of-speech tagging. Each concept can represent one or more documents to be classified during a review. Clustering of the concepts provides an overall view of the document space, which allows users to easily identify documents sharing a common theme.

[0024] The clustering of tokens, for example, concepts, differs from document clustering, which groups related documents individually. In contrast, concept clustering groups related concepts, which are each representative of one or more related documents. Each concept can express an idea or topic that may not be expressed by individual documents. A concept is analogous to a search query by identifying documents associated with a particular idea or topic.

[0025] A user can determine how particular concepts are related based on the concept clustering. Further, users are able to intuitively identify documents by selecting one or more associated concepts in a cluster. For example, a user may wish to identify all documents in a particular corpus that are related to car manufacturing. The user can select the concept “car manufacturing” or “vehicle manufacture” within one of the clusters and subsequently, the associated documents are presented. However, during document clustering, a user is first required to select a specific document from which other documents that are similarly related can then be identified.

Providing Suggestions Using Reference Concepts

[0026] Reference concepts are previously classified based on the document content represented by that concept and can be injected into clusters of uncoded, that is unclassified, concepts to influence classification of the uncoded concepts. Specifically, relationships between an uncoded concept and the reference concepts, in terms of semantic similarity or distinction, can be used as an aid in providing suggestions for classifying uncoded concepts. Once classified, the newly-coded, or reference, concepts can be used to further classify the represented documents. Although tokens, such as word-level or character-level n-grams, raw terms, entities, or concepts, can be clustered and displayed, the discussion below will focus on a concept as a particular token.

[0027] Complete ESI review requires a support environment within which classification can be performed. FIG. 1 is a block diagram showing a system 10 for displaying relationships between concepts to provide classification suggestions via inclusion, in accordance with one embodiment. By way of illustration, the system 10 operates in a distributed computing environment, which includes a plurality of heterogeneous systems and ESI sources. Henceforth, a single item of ESI will be referenced as a “document,” although ESI can include other forms of non-document data, as described infra. A backend server 11 is coupled to a storage device 13, which stores documents 14a in the form of structured or unstructured data, a database 30 for maintaining information about the documents, a lookup database 38 for storing many-to-many mappings 39 between documents and document features, and a concept document index 40, which maps documents to concepts. The storage device 13 also stores classified documents 14b, concepts 14c, and reference concepts 14d. Concepts are collections of nouns and noun-phrases with common semantic meaning. The nouns and noun-phrases can be extracted from one or more documents in the corpus for review. Hereinafter, the terms “classified” and “coded” are used interchangeably with the same intended meaning, unless otherwise indicated. A set of reference concepts can be hand-selected or automatically selected through guided review, which is further discussed below. Additionally, the set of reference concepts can be predetermined or can be generated dynamically, as uncoded concepts are classified and subsequently added to the set of reference concepts.

[0028] The backend server 11 is coupled to an intranetwork 21 and executes a workbench suite 31 for providing a user interface framework for automated document management, processing, analysis, and classification. In a further embodiment, the backend server 11 can be accessed via an internet-work 22. The workbench software suite 31 includes a document mapper 32 that includes a clustering engine 33, similarity searcher 34, classifier 35, and display generator 36. Other workbench suite modules are possible.

[0029] The clustering engine 33 performs efficient document scoring and clustering of uncoded concepts and reference concepts, such as described in commonly-assigned U.S. Pat. No. 7,610,313, the disclosure of which is incorporated by reference. Clusters of uncoded concepts 14c and reference concepts 14d are formed and organized along vectors, known as spines, based on a similarity of the clusters. The similarity can be expressed in terms of distance. Concept clustering is further discussed below with reference to FIG. 5. The classifier 35 provides a machine-generated suggestion and confi-

dence level for classification of selected uncoded concepts **14c**, clusters, or spines, as further described below with reference to FIG. 8.

[0030] The display generator **36** arranges the clusters and spines in thematic relationships in a two-dimensional visual display space, as further described below beginning with reference to FIG. 2. Once generated, the visual display space is transmitted to a work client **12** by the backend server **11** via the document mapper **32** for presenting to a reviewer on a display **37**. The reviewer can include an individual person who is assigned to review and classify one or more uncoded concepts by designating a code. Hereinafter, the terms “reviewer” and “custodian” are used interchangeably with the same intended meaning, unless otherwise indicated. Other types of reviewers are possible, including machine-implemented reviewers.

[0031] The document mapper **32** operates on uncoded concepts **14a**, which can be retrieved from the storage **13**, as well as from a plurality of local and remote sources. As well, the local and remote sources can also store the reference documents **14b**, concepts **14c**, and reference concepts **14d**. The local sources include documents and concepts **17** maintained in a storage device **16** coupled to a local server **15**, and documents and concepts **20** maintained in a storage device **19** coupled to a local client **18**. The local server **15** and local client **18** are interconnected to the backend server **11** and the work client **12** over an intranetwork **21**. In addition, the document mapper **32** can identify and retrieve concepts from remote sources over an internetwork **22**, including the Internet, through a gateway **23** interfaced to the intranetwork **21**. The remote sources include documents and concepts **26** maintained in a storage device **25** coupled to a remote server **24**, and documents and concepts **29** maintained in a storage device **28** coupled to a remote client **27**. Other document sources, either local or remote, are possible.

[0032] The individual documents **14a**, **14b**, **17**, **20**, **26**, **29** include all forms and types of structured and unstructured ESI, including electronic message stores, word processing documents, electronic mail (email) folders, Web pages, and graphical or multimedia data. Notwithstanding, the documents could be in the form of structurally organized data, such as stored in a spreadsheet or database.

[0033] In one embodiment, the individual documents **14a**, **14b**, **17**, **20**, **26**, **29** include electronic message folders storing email and attachments, such as maintained by the Outlook and Outlook Express products, licensed by Microsoft Corporation, Redmond, Wash. The database can be an SQL-based relational database, such as the Oracle database management system, Release 8, licensed by Oracle Corporation, Redwood Shores, Calif.

[0034] Additionally, the individual concepts **14c**, **14d**, **17**, **20**, **26**, **29** include uncoded concepts **14c** and reference concepts **14d**. The uncoded concepts **14c**, which are unclassified, represent collections of nouns and noun-phrases that are semantically related and extracted from documents in a document review project. The reference concepts **14d** are initially uncoded concepts that can represent documents selected from the corpus or other sources of documents. The reference concepts **14d** assist in providing suggestions for classification of the remaining uncoded concepts representative of the document corpus based on visual relationships between the uncoded concepts and reference concepts. The reviewer can classify one or more of the remaining uncoded concepts by assigning a classification code based on the relationships. In

a further embodiment, the reference concepts can be used as a training set to form machine-generated suggestions for classifying the remaining uncoded concepts, as further described below with reference to FIG. 8.

[0035] The concept corpus for a document review project can be divided into subsets of uncoded concepts, which are each provided to a particular reviewer as an assignment. The uncoded documents are analyzed to identify concepts, which are subsequently clustered. A classification code can be assigned to each of the clustered concepts. To maintain consistency, the same codes can be used across all concepts representing assignments in the document review project. The classification codes can be determined using taxonomy generation, during which a list of classification codes can be provided by a reviewer or determined automatically. The classification code of a concept can be assigned to the documents associated with that concept.

[0036] For purposes of legal discovery, the list of classification codes can include “privileged,” “responsive,” or “non-responsive,” however, other classification codes are possible. The assigned classification codes can be used as suggestions for classification of associated documents. For example, a document associated with three concepts, each assigned a “privileged” classification can also be considered “privileged.” Other types of suggestions are possible. A “privileged” document contains information that is protected by a privilege, meaning that the document should not be disclosed or “produced” to an opposing party. Disclosing a “privileged” document can result in an unintentional waiver of the subject matter disclosed. A “responsive” document contains information that is related to the legal matter, while a “non-responsive” document includes information that is not related to the legal matter.

[0037] The system **10** includes individual computer systems, such as the backend server **11**, work server **12**, server **15**, client **18**, remote server **24** and remote client **27**. The individual computer systems are general purpose, programmed digital computing devices consisting of a central processing unit (CPU), random access memory (RAM), non-volatile secondary storage, such as a hard drive or CD ROM drive, network interfaces, and peripheral devices, including user interfacing means, such as a keyboard and display. The various implementations of the source code and object and byte codes can be held on a computer-readable storage medium, such as a floppy disk, hard drive, digital video disk (DVD), random access memory (RAM), read-only memory (ROM) and similar storage mediums. For example, program code, including software programs, and data are loaded into the RAM for execution and processing by the CPU and results are generated for display, output, transmittal, or storage.

[0038] Identifying relationships between the reference concepts and uncoded concepts includes clustering. FIG. 2 is a process flow diagram showing a method **50** for displaying relationships between concepts to provide classification suggestions via inclusion, in accordance with one embodiment. A subset of reference concepts is identified and selected (block **51**) from a representative set of reference concepts. The subset of reference concepts can be predefined, arbitrary, or specifically selected, as discussed further below with reference to FIG. 3. Upon identification, the reference concept subset is grouped with uncoded concepts (block **52**). The uncoded concepts can include all uncoded concepts in an assignment or in a corpus. The grouped concepts, including uncoded and reference concepts are organized into clusters

(block 53). Clustering of the concepts is discussed further below with reference to FIG. 5.

[0039] Once formed, the clusters can be displayed to visually depict relationships (block 54) between the uncoded concepts and the reference concepts. The relationships can provide a suggestion, which can be used by an individual reviewer for classifying one or more of the uncoded concepts, clusters, or spines. Based on the relationships, the reviewer can classify the uncoded concepts, clusters, or spines by assigning a classification code, which can represent a relevancy of the uncoded concept to the document review project. Further, machine classification can provide a suggestion for classification, including a classification code, based on a calculated confidence level (block 55). Classifying uncoded concepts is further discussed below with reference to FIG. 8.

[0040] In one embodiment, the classified concepts can be used as suggestions for classifying those documents represented by that concept. For example, in a product liability lawsuit, the plaintiff claims that a wood composite manufactured by the defendant induces and harbors mold growth. During discovery, all documents within the corpus for the lawsuit and relating to mold should be identified for review. The concept for mold is clustered and includes a “responsive” classification code, which indicates that the noun phrase mold is related to the legal matter. Upon selection of the mold concept, all documents that include the noun phrase mold can be identified using the mapping matrix, which is described below with reference to FIG. 3. The responsive classification code assigned to the concept can be used as a suggestion for the document classification. However, if the document is represented by multiple concepts with different classification codes, each different code can be considered during classification of the document.

[0041] In a further embodiment, the concept clusters can be used with document clusters, which are described in commonly-owned in U.S. Patent Application Publication No. 20110029526, published Feb. 3, 2011, pending, and U.S. Pat. No. 8,515,957, issued Aug. 20, 2013, the disclosures of which is incorporated by reference. For example, selecting a concept in the concept cluster display can identify one or more documents with a common idea or topic. Further selection of one of the documents represented by the selected cluster in the document concept display can identify documents that are similarly related to the content of the selected document. The identified documents can be the same or different as the other documents represented by the concept.

[0042] Similar documents can also be identified as described in commonly-assigned U.S. Pat. No. 8,572,084, issued Oct. 29, 2013, the disclosure of which is incorporated by reference.

[0043] In an even further embodiment, the documents identified from one of the concepts can be classified automatically as described in commonly-assigned U.S. Pat. No. 8,635,223, issued Jan. 21, 2014, pending, the disclosure of which is incorporated by reference.

[0044] Identifying a Set and Subset of Reference Concepts

[0045] Prior to clustering, the uncoded concepts and reference concepts are obtained. The reference concepts used for clustering can include a particular subset of reference concepts, which are selected from a general set of reference concepts. Alternatively, the entire set of reference concepts can be clustered with the uncoded concepts. The set of reference concepts is representative of document in the corpus for

a document review project in which data organization or classification is desired. The reference concept set can be previously defined and maintained for related concept review projects or can be specifically generated for each review project. A predefined reference set provides knowledge previously obtained during the related concept review project to increase efficiency, accuracy, and consistency. Reference sets newly generated for each review project can include arbitrary or customized reference sets that are determined by a reviewer or a machine.

[0046] The set of reference concepts can be generated during guided review, which assists a reviewer in building a reference concept set. During guided review, the uncoded concepts that are dissimilar to the other uncoded concepts are identified based on a similarity threshold. Other methods for determining dissimilarity are possible. Identifying a set of dissimilar concepts provides a group of uncoded concepts that is representative of the corpus for the document review project. Each identified dissimilar concept is then classified by assigning a particular classification code based on the content of the concept to collectively generate a set of reference concepts. Guided review can be performed by a reviewer, a machine, or a combination of the reviewer and machine.

[0047] Other methods for generating a reference concept set for a document review project using guided review are possible, including clustering. For example, a set of uncoded concepts to be classified is clustered, as described in commonly-assigned U.S. Pat. No. 7,610,313, the disclosure of which is incorporated by reference. A plurality of the clustered uncoded concepts are selected based on selection criteria, such as cluster centers or sample clusters. The cluster centers can be used to identify uncoded concepts in a cluster that are most similar or dissimilar to the cluster center. The identified uncoded concepts are then selected for classification by assigning classification codes. After classification, the concepts represent a reference set. In a further embodiment, sample clusters can be used to generate a reference concept set by selecting one or more sample clusters based on cluster relation criteria, such as size, content, similarity, or dissimilarity. The uncoded concepts in the selected sample clusters are then assigned classification codes. The classified concepts represent a concept reference set for the document review project. Other methods for selecting concepts for use as a reference set are possible.

[0048] Once generated, a subset of reference concepts is selected from the reference concept set for clustering with uncoded concepts. FIG. 3 is a block diagram showing, by way of example, measures 60 for selecting reference concept subsets 61 for use in the method of FIG. 2. A reference concept subset 61 includes one or more reference concepts selected from a set of reference concepts associated with a document review project for use in clustering with uncoded concepts. The reference concept subset can be predefined 62, customized 64, selected arbitrarily 63, or based on similarity 65.

[0049] A subset of predefined reference concepts 62 can be selected from a reference set, which is associated with another document review project that is related to the current document review project. An arbitrary reference subset 63 includes reference concepts randomly selected from a reference set, which can be predefined or newly generated for the current document review project or a related document review project. A customized reference subset 64 includes reference concepts specifically selected from a current or

related reference set based on criteria, such as reviewer preference, classification category, document source, content, and review project. Other criteria are possible. The number of reference concepts in a subset can be determined automatically or by a reviewer based on reference factors, such as a size of the document review project, an average size of the assignments, types of classification codes, and a number of reference concepts associated with each classification code. Other reference factors are possible. In a further embodiment, the reference concept subset can include more than one occurrence of a reference concept. Other types of reference concept subsets and methods for selecting the reference concept subsets are possible.

[0050] Forming Clusters

[0051] Once identified, the reference concept subset can be used for clustering with uncoded concept representative of a corpus for a particular document review project. The corpus of uncoded concepts for a review project can be divided into assignments using assignment criteria, such as custodian or source of the uncoded concept, content, document type, and date. Other criteria are possible. In one embodiment, each assignment is assigned to an individual reviewer for analysis. The assignments can be separately clustered with the reference concept subset or alternatively, all of the uncoded concepts in the corpus can be clustered with the reference concept subset. The assignments can be separately analyzed or alternatively, analyzed together to determine concepts for the one or more document assignments. The content of each document within the corpus can be converted into a set of concepts. As described above, concepts typically include nouns and noun phrases obtained through part-of-speech tagging that have a common semantic meaning. The concepts, which are representative of the documents can be clustered to provide an intuitive grouping of the document content.

[0052] Clustering of the uncoded concepts provides groupings of related uncoded concepts and is based on a similarity metric using score vectors assigned to each uncoded concept. The score vectors can be generated using a matrix showing the uncoded concepts in relation to documents that contain the concepts. FIG. 4 is a table showing, by way of example, a matrix mapping 70 of uncoded concepts 74 and documents 73. The uncoded documents 73 are listed along a horizontal dimension 71 of the matrix, while the concepts 74 are listed along a vertical dimension 72. However, the placement of the uncoded documents 73 and concepts 74 can be reversed. Each cell 75 within the matrix 70 includes a cumulative number of occurrences of each concept within a particular uncoded document 73. Score vectors can be generated for each document by identifying the concepts and associated weights within that document and ordering the concepts along a vector with the associated concept weight. In the matrix 70, the score vector 76 for a document 73 can be identified as all the concepts included in that document and the associated weights, which are based on the number of occurrences of each concept. Score vectors can also be generated for each concept by identifying the documents that contain that concept and determining a weight associated with each document. The documents and associated weights are then ordered along a vector for each concept, as the concept score vector. In the matrix 70, the score vector 77 for a concept can be identified as all the documents that contain that concept and the associated weights.

[0053] Clustering provides groupings of related uncoded concepts and reference concepts. FIG. 5 is a flow diagram

showing a routine 80 for forming clusters for use in the method 40 of FIG. 2. The purpose of this routine is to use score vectors associated with the concepts, including uncoded and reference concepts, to form clusters based on relative similarity. Hereinafter, the term “concept” is intended to include uncoded concepts and reference concepts selected for clustering, unless otherwise indicated. The score vector associated with each concept includes a set of paired values of documents and associated weights, which are based on scores. The score vector is generated by scoring the documents, as described in commonly-assigned U.S. Pat. No. 7,610,313, the disclosure of which is incorporated by reference.

[0054] As an initial step for generating score vectors, each document within a concept is individually scored. Next, a normalized score vector is created for the concept by identifying paired values, consisting of a document represented by that concept and the scores for that document. The paired values are ordered along a vector to generate the score vector. The paired values can be ordered based on the documents, as well as other factors. For example, assume a normalized score vector for a first Concept A is $\vec{S}_A = \{(5, 0.5), (120, 0.75)\}$ and a normalized score vector for another Concept B is $\vec{S}_B = \{(3, 0.4), (5, 0.75), (47, 0.15)\}$. Concept A has scores corresponding to tokens ‘5’ and ‘120’ and Concept B has scores corresponding to tokens ‘3,’ ‘5’ and ‘47.’ Thus, these concepts only have token ‘5’ in common. Once generated, the score vectors can be compared to determine similarity or dissimilarity between the corresponding concepts during clustering.

[0055] The routine for forming clusters of concepts, including uncoded concepts and reference concepts, proceeds in two phases. During the first phase (blocks 83-88), the concepts are evaluated to identify a set of seed concepts, which can be used to form new clusters. During the second phase (blocks 90-96), any concepts not previously placed are evaluated and grouped into the existing clusters based on a best-fit criterion.

[0056] Initially, a single cluster is generated with one or more concepts as seed concepts and additional clusters of concepts are added, if necessary. Each cluster is represented by a cluster center that is associated with a score vector, which is representative of all the documents associated with concepts in that cluster. The cluster center score vector can be generated by comparing the score vectors for the individual concepts in the cluster and identifying common documents shared by the concepts. The most common documents and associated weights are ordered along the cluster center score vector. Cluster centers and thus, cluster center score vectors may continually change due to the addition and removal of concepts during clustering.

[0057] During clustering, the concepts are identified (block 81) and ordered by length (block 82). The concepts can include all reference concepts in a subset and one or more assignments of uncoded concepts. Each concept is then processed in an iterative processing loop (blocks 83-88) as follows. The similarity between each concept and a center of each cluster is determined (block 84) as the cosine (cos) of the score vectors for the concept and cluster being compared. The cos a provides a measure of relative similarity or dissimilarity between the concepts associated with the documents and is equivalent to the inner products between the score vectors for the concept and cluster center.

[0058] In the described embodiment, the $\cos \sigma$ is calculated in accordance with the equation:

$$\cos \sigma_{AB} = \frac{\langle \vec{S}_A, \vec{S}_B \rangle}{|\vec{S}_A| |\vec{S}_B|}$$

where $\cos \sigma_{AB}$ comprises the similarity metric between Concept A and cluster center B, \vec{S}_A comprises a score vector for the Concept A, and \vec{S}_B comprises a score vector for the cluster center B. Other forms of determining similarity using a distance metric are feasible, as would be recognized by one skilled in the art. An example includes using Euclidean distance.

[0059] Only those concepts that are sufficiently distinct from all cluster centers (block 85) are selected as seed concepts for forming new clusters (block 86). If the concept being compared is not sufficiently distinct (block 85), the concept is then grouped into a cluster with the most similar cluster center (block 87). Processing continues with the next concept (block 88).

[0060] In the second phase, each concept not previously placed is iteratively processed in an iterative processing loop (blocks 90-96) as follows. Again, the similarity between each remaining concept and each of the cluster centers is determined based on a distance (block 91), such as the $\cos \sigma$ of the normalized score vectors for each of the remaining concepts and the cluster centers. A best fit between a remaining concept and a cluster center can be found subject to a minimum fit criterion (block 92). In the described embodiment, a minimum fit criterion of 0.25 is used, although other minimum fit criteria could be used. If a best fit is found (block 93), the remaining concept is grouped into the cluster having the best fit (block 95). Otherwise, the remaining concept is grouped into a miscellaneous cluster (block 94). Processing continues with the next remaining concept (block 96). Finally, a dynamic threshold can be applied to each cluster (block 97) to evaluate and strengthen concept membership in a particular cluster. The dynamic threshold is applied based on a cluster-by-cluster basis, as described in commonly-assigned U.S. Pat. No. 7,610,313, the disclosure of which is incorporated by reference. The routine then returns. Other methods and processes for forming clusters are possible.

[0061] Alternatively, clusters can be generated by injection as further described in commonly-owned U.S. Pat. No. 8,635,223, issued Jan. 21, 2014, the disclosure of which is incorporated by reference.

[0062] Once clustered, similar concepts can be identified as described in commonly-assigned U.S. Pat. No. 8,645,378, issued Feb. 4, 2014, the disclosure of which is incorporated by reference.

[0063] Displaying the Reference Concepts

[0064] Once formed, the clusters of concepts can be organized to generate spines of thematically related clusters, as described in commonly-assigned U.S. Pat. No. 7,271,804, the disclosure of which is incorporated by reference. Each spine includes those clusters that share one or more concepts, which are placed along a vector. Also, the cluster spines can be positioned in relation to other cluster spines based on a theme shared by those cluster spines, as described in commonly-assigned U.S. Pat. No. 7,610,313, the disclosure of which is incorporated by reference. Each theme can include

one or more concepts defining a semantic meaning. Organizing the clusters into spines and groups of cluster spines provides an individual reviewer with a display that presents the concepts according to a theme while maximizing the number of relationships depicted between the concepts.

[0065] FIG. 6 is a screenshot 100 showing, by way of example, a visual display 81 of reference concepts 105 in relation to uncoded concepts 104. Clusters 103 can be located along a spine, which is a straight vector, based on a similarity of the concepts 104, 105 in the clusters 103. Each cluster 103 is represented by a circle; however, other shapes, such as squares, rectangles, and triangles are possible, as described in U.S. Pat. No. 6,888,548, the disclosure of which is incorporated by reference. The uncoded concepts 104 are each represented by a smaller circle within the clusters 103, while the reference concepts 105 are each represented by a circle having a diamond shape within the boundaries of the circle. The reference concepts 105 can be further represented by their assigned classification code. The classification codes can include “privileged,” “responsive,” and “non-responsive” codes, as well as other codes. Each group of reference concepts associated with a particular classification code can be identified by a different color. For instance, “privileged” reference concepts can be colored blue, while “non-responsive” reference concepts are red and “responsive” reference concepts are green. In a further embodiment, the reference concepts for different classification codes can include different symbols. For example, “privileged” reference concepts can be represented by a circle with an “X” in the center, while “non-responsive” reference concepts can include a circle with striped lines and “responsive” reference concepts can include a circle with dashed lines. Other classification representations for the reference concepts are possible. Each cluster spine 86 is represented as a straight vector along which the clusters are placed.

[0066] The display 101 can be manipulated by an individual reviewer via a compass 102, which enables the reviewer to navigate, explore, and search the clusters 103 and spines 106 appearing within the compass 102, as further described in commonly-assigned U.S. Pat. No. 7,356,777, the disclosure of which is incorporated by reference. Visually, the compass 102 emphasizes clusters 103 located within the compass 102, while deemphasizing clusters 103 appearing outside of the compass 102.

[0067] Spine labels 109 appear outside of the compass 102 at an end of each cluster spine 106 to connect the outermost cluster of a cluster spine 106 to the closest point along the periphery of the compass 102. In one embodiment, the spine labels 109 are placed without overlap and circumferentially around the compass 102. Each spine label 109 corresponds to one or more documents represented by the clustered concepts that most closely describe the cluster spines 106. Additionally, the documents associated with each of the spine labels 109 can appear in a documents list (not shown) also provided in the display. Additionally, the cluster concepts for each of the spine labels 109 can appear in a documents list (not shown) also provided in the display. Toolbar buttons 107 located at the top of the display 101 enable a user to execute specific commands for the composition of the spine groups displayed. A set of pull down menus 108 provide further control over the placement and manipulation of clusters 103 and cluster spines 106 within the display 101. Other types of controls and functions are possible.

[0068] A concept guide 110 can be placed within the display 101. The concept guide 110 can include a “Selected” field, a “Search Results” field, and details regarding the numbers of uncoded concepts and reference concepts provided in the display. The number of uncoded concepts includes all uncoded concepts selected for clustering, such as within a corpus of uncoded concepts for a review project or within an assignment. The number of reference concepts includes the reference concept subset selected for clustering. The “Selected” field in the document guide 110 provides a number of concepts within one or more clusters selected by the reviewer. The reviewer can select a cluster by “double clicking” the visual representation of that cluster using a mouse. The “Search Results” field provides a number of uncoded concepts and reference concepts that include a particular search term identified by the reviewer in a search query box 112.

[0069] In one embodiment, a garbage can 111 is provided to remove documents, from consideration in the current set of clusters 103. Removed cluster documents prevent those documents from affecting future clustering, as may occur when a reviewer considers a document irrelevant to the clusters 103.

[0070] The display 101 provides a visual representation of the relationships between thematically-related concepts, including the uncoded concepts and reference concepts. The uncoded concepts and reference concepts located within a cluster or spine can be compared based on characteristics, such as the assigned classification codes of the reference concepts, a number of reference concepts associated with each classification code, and a number of different classification codes to identify relationships between the uncoded concepts and reference concepts. The reviewer can use the displayed relationships as suggestions for classifying the uncoded concepts. For example, FIG. 7A is a block diagram showing, by way of example, a cluster 120 with “privileged” reference concepts 122 and uncoded concepts 121. The cluster 120 includes nine uncoded concepts 121 and three reference concepts 122. Each reference concept 122 is classified as “privileged.” Accordingly, based on the number of “privileged” reference concepts 122 present in the cluster 120, the absence of other classifications of reference concepts, and the thematic relationship between the uncoded concepts 94 and the “privileged” reference concepts 122, the reviewer may be more inclined to review the uncoded concepts 121 in that cluster 120 or to classify one or more of the uncoded concepts 121 as “privileged” without review.

[0071] Alternatively, the three reference concepts can be classified as “non-responsive,” instead of “privileged” as in the previous example. FIG. 7B is a block diagram showing, by way of example, a cluster 123 with “non-responsive” reference concepts 124 and uncoded concepts 121. The cluster 123 includes nine uncoded concepts 121 and three “non-responsive” concepts 124. Since the uncoded concepts 121 in the cluster are thematically related to the “non-responsive” reference concepts 124, the reviewer may wish to assign a “non-responsive” code to one or more of the uncoded concepts 121 without review, as they are most likely not relevant to the legal matter associated with the document review project. In making a decision to assign a code, such as “non-responsive,” the reviewer can consider the number of “non-responsive” reference concepts in the cluster, the presence or absence of other reference concept classification codes, and the thematic relationship between the “non-responsive” reference concepts and the uncoded concepts. Thus, the pres-

ence of the three “non-responsive” reference concepts 124 in the cluster provides a suggestion that the uncoded concepts 121 may also be “non-responsive.” Further, the label 109 associated with the spine 106 upon which the cluster is located can also be used to influence a suggestion.

[0072] A further example can include a cluster with combination of “privileged” and “non-responsive” reference concepts. For example, FIG. 7C is a block diagram showing, by way of example, a cluster 125 with uncoded concepts 121 and a combination of differently classified reference concepts 122, 124. The cluster 125 can include one “privileged” reference concept 122, two “non-responsive” reference concepts 124, and nine uncoded concepts 121. The “privileged” 122 and “non-responsive” 124 reference concepts can be distinguished by different colors or shape, as well as other identifiers. The combination of “privileged” 122 and “non-responsive” 124 reference concepts within the cluster 98 can suggest to a reviewer that the uncoded reference concepts 121 should be reviewed before classification or that one or more uncoded reference concepts 121 should be classified as “non-responsive” based on the higher number of “non-responsive” reference concepts 124 in the cluster 125. In making a classification decision, the reviewer may consider the number of “privileged” reference concepts 122 versus the number of “non-responsive” reference concepts 124, as well as the thematic relationships between the uncoded concepts 121 and the “privileged” 122 and “non-responsive” 124 reference concepts. Additionally, the reviewer can identify the closest reference concept to an uncoded concept and assign the classification code of the closest reference concept to the uncoded concept. Other examples, classification codes, and combinations of classification codes are possible.

[0073] Additionally, the reference concepts can also provide suggestions for classifying clusters and spines. The suggestions provided for classifying a cluster can include factors, such as a presence or absence of classified concepts with different classification codes within the cluster and a quantity of the classified concepts associated with each classification code in the cluster. The classification code assigned to the cluster is representative of the concepts in that cluster and can be the same as or different from one or more classified concepts within the cluster. Further, the suggestions provided for classifying a spine include factors, such as a presence or absence of classified concepts with different classification codes within the clusters located along the spine and a quantity of the classified concepts for each classification code. Other suggestions for classifying concepts, clusters, and spines are possible.

Classifying Uncoded Concepts

[0074] The display of relationships between the uncoded concepts and reference concepts can provide suggestions to an individual reviewer. The suggestions can indicate a need for manual review of the uncoded concepts, when review may be unnecessary, and hints for classifying the uncoded concepts. Additional information can be generated to assist the reviewer in making classification decisions for the uncoded concepts, such as a machine-generated confidence level associated with a suggested classification code, as described in common-assigned U.S. Pat. No. 8,515,958, issued on Aug. 20, 2013, the disclosure of which is incorporated by reference.

[0075] The machine-generated suggestion for classification and associated confidence level can be determined by a

classifier. FIG. 8 is a process flow diagram 130 showing, by way of example, a method for classifying uncoded concepts by a classifier for use in the method of FIG. 2. An uncoded concept is selected from a cluster within a cluster set (block 131) and compared to a neighborhood of x-reference concepts (block 132), also located within the cluster, to identify those reference concepts that are most relevant to the selected uncoded concept. In a further embodiment, a machine-generated suggestion for classification and an associated confidence level can be provided for a cluster or spine by selecting and comparing the cluster or spine to a neighborhood of x-reference concepts determined for the selected cluster or spine.

[0076] The neighborhood of x-reference concepts is determined separately for each selected uncoded concept and can include one or more reference concepts within that cluster. During neighborhood generation, an x-number of reference concepts is first determined automatically or by an individual reviewer. Next, the x-number of reference concepts nearest in distance to the selected uncoded concept are identified. Finally, the identified x-number of reference concepts are provided as the neighborhood for the selected uncoded concept. In a further embodiment, the x-number of reference concepts are defined for each classification code, rather than across all classification codes. Once generated, the x-number of reference concepts in the neighborhood and the selected uncoded concept are analyzed by the classifier to provide a machine-generated classification suggestion (block 133). A confidence level for the suggested classification is also provided (block 134).

[0077] The analysis of the selected uncoded concept and x-number of reference concepts can be based on one or more routines performed by the classifier, such as a nearest neighbor (NN) classifier. The routines for determining a suggested classification code include a minimum distance classification measure, also known as closest neighbor, minimum average distance classification measure, maximum count classification measure, and distance weighted maximum count classification measure. The minimum distance classification measure includes identifying a neighbor that is the closest distance to the selected uncoded concept and assigning the classification code of the closest neighbor as the suggested classification code for the selected uncoded concept. The closest neighbor is determined by comparing the score vectors for the selected uncoded concept with each of the x-number of reference concepts in the neighborhood as the cos a to determine a distance metric. The distance metrics for the x-number of reference concepts are compared to identify the reference concept closest to the selected uncoded concept as the closest neighbor.

[0078] The minimum average distance classification measure includes calculating an average distance of the reference concepts in a cluster for each classification code. The classification code with the reference concepts having the closest average distance to the selected uncoded concept is assigned as the suggested classification code. The maximum count classification measure, also known as the voting classification measure, includes counting a number of reference concepts within the cluster for each classification code and assigning a count or "vote" to the reference concepts based on the assigned classification code. The classification code with the highest number of reference concepts or "votes" is assigned to the selected uncoded concept as the suggested classification. The distance weighted maximum count classification

measure includes identifying a count of all reference concepts within the cluster for each classification code and determining a distance between the selected uncoded concept and each of the reference concepts. Each count assigned to the reference concepts is weighted based on the distance of the reference concept from the selected uncoded concept. The classification code with the highest count, after consideration of the weight, is assigned to the selected uncoded concept as the suggested classification.

[0079] The machine-generated classification code is provided for the selected uncoded concept with a confidence level, which can be presented as an absolute value or a percentage. Other confidence level measures are possible. The reviewer can use the suggested classification code and confidence level to assign a classification to the selected uncoded concept. Alternatively, the x-NN classifier can automatically assign the suggested classification. In one embodiment, the x-NN classifier only assigns an uncoded concept with the suggested classification code if the confidence level is above a threshold value, which can be set by the reviewer or the x-NN classifier.

[0080] Classification can also occur on a cluster or spine level. For instance, for cluster classification, a cluster is selected and a score vector for the center of the cluster is determined as described above with reference to FIG. 5. A neighborhood for the selected cluster is determined based on a distance metric. The x-number of reference concepts that are closest to the cluster center can be selected for inclusion in the neighborhood, as described above. Each reference concept in the selected cluster is associated with a score vector and the distance is determined by comparing the score vector of the cluster center with the score vector of each reference concept to determine an x-number of reference concepts that are closest to the cluster center. However, other methods for generating a neighborhood are possible. Once determined, one of the classification measures is applied to the neighborhood to determine a suggested classification code and confidence level for the selected cluster.

[0081] During classification, either by an individual reviewer or a machine, the reviewer can retain control over many aspects, such as a source of the reference concepts and a number of reference concepts to be selected. FIG. 9 is a screenshot 140 showing, by way of example, an options dialogue box 141 for entering user preferences for clustering and display of the uncoded concepts and reference concepts. The dialogue box 141 can be accessed via a pull-down menu as described above with respect to FIG. 6. Within the dialogue box 141, the reviewer can utilize user-selectable parameters to define a reference source 142, category filter 143, command details 144, advanced options 145, classifier parameters 146, and commands 147. Each user-selectable option can include a text box for entry of a user preference or a drop-down menu with predetermined options for selection by the reviewer. Other user-selectable options and displays are possible.

[0082] The reference source parameter 142 allows the reviewer to identify one or more sources of the reference concepts. The sources can include all reference concepts for which the associated classification has been verified, all reference concepts that have been analyzed, and all reference concepts in a particular binder. The binder can include reference concepts particular to a current document review project or that are related to a prior document review project. The category filter parameter 143 allows the reviewer to generate

and display the subset of reference concepts using only those reference concepts associated with a particular classification code. Other options for generating the reference set are possible, including custodian, source, and content. The command parameters **144** allow the reviewer to enter instructions regarding actions for the uncoded and reference concepts, such as indicating counts of the concepts, and display of the concepts. The advanced option parameters **145** allow the reviewer to specify clustering thresholds and classifier parameters. The parameters entered by the user can be compiled as command parameters **146** and provided in a drop-down menu on a display of the clusters. Other user selectable parameters, options, and actions are possible.

[0083] In a further embodiment, once the uncoded concepts are assigned a classification code, the newly-classified uncoded concepts can be placed into the concept reference set for use in providing classification suggestions for other uncoded concepts.

[0084] In yet a further embodiment, each document can be represented by more than one concept. Accordingly, to determine a classification code for the document, the classification codes for each of the associated concepts can be analyzed and compared for consideration in classifying the document. In one example, a classification code can be determined by counting the number of associated concepts for each classification code and then assigned the classification code with the most associated concepts. In a further example, one or more of the associated concepts can be weighted and the classification code associated with the highest weight of concepts is assigned. Other methods for determining a classification code for uncoded documents based on reference concepts are possible.

[0085] Although clustering and displaying relationships has been described above with reference to concepts, other tokens, such as word-level or character-level n-grams, raw terms, and entities, are possible.

[0086] While the invention has been particularly shown and described as referenced to the embodiments thereof, those skilled in the art will understand that the foregoing and other changes in form and detail may be made therein without departing from the spirit and scope.

What is claimed is:

1. A computer-implemented system for providing visual classification suggestions for inclusion-based concept clusters, comprising the steps of:

- a computer processor configured to execute modules, comprising:
 - a designation module configured to designate reference concepts each associated with a classification code;
 - a grouping module configured to group one or more of the reference concepts with a plurality of uncoded concepts into a grouped concept set;
 - a generation module configured to generate clusters, each comprising a portion of the uncoded concepts and the reference concepts of the grouped concept set; and
 - a suggestion module configured to provide a visual suggestion for assigning one of the classification codes to one of the clusters comprising visually representing each of the reference concepts in that cluster.

2. A system according to claim 1, further comprising:

- a presence module configured to provide at least one of a presence and an absence of the concepts with each of the classification codes in that cluster; and

- a number module configured to provide a number of the concepts with each of the classification codes in that cluster,

- wherein the suggestion includes at least one of the number and the presence and the absence.

3. A system according to claim 2, further comprising:

- a concept module configured to provide a suggestion for classifying one of the uncoded concepts in that cluster comprising at least one of the number and the presence and the absence.

4. A system according to claim 1, further comprising:

- a selection module configured to receive a user selection of one or more sources, custodians, content, and the classification codes of the reference concepts, wherein the reference concepts are designated in accordance with the selection.

5. A system according to claim 1, further comprising:

- a parameter module configured to receive a user selection of one or more parameters for generating the clusters, wherein the clustering is performed in accordance with the selection.

6. A system according to claim 1, further comprising:

- a spine module configured to determine a similarity between the clusters and organizing the clusters along one or more spines based on the similarity, each of the spines comprising a vector; and

- a spine classification module configured to provide a visual classification suggestion for assigning one of the classification codes to one of the spines comprising visually representing the clusters along that spine and the reference concepts in the clusters along that spine.

7. A system according to claim 6, further comprising:

- a user selection module configured to receive a user selection of one or more of the clusters and spines;

- a display module to display one or more of the clusters and spines within a compass;

- an emphasis module configured to emphasize those of the clusters and spines displayed within the compass; and

- a deemphasizing module configured to deemphasize those of the clusters and spines displayed outside of the compass.

8. A system according to claim 6, further comprising:

- a label module configured to associate a label with each of the spines, each label associated with one or more documents from which one or more of the concepts in the clusters along that spine were extracted; and

- a list module configured to display a list of the documents associated with one of the spine labels.

9. A system according to claim 1, wherein the visual representation of the each classification codes comprises at least one of a symbol, shape, and color different from the visual representations of the remaining classification codes.

10. A system according to claim 1, further comprising:

- an identification module configured to identify one of the concepts as a center of the cluster;

- a neighborhood module configured to identify a neighborhood of similar reference concepts for the cluster based on the cluster center; and

- an assignment module configured to assign one of the classification codes to the cluster based on the neighborhood.

11. A computer-implemented method for providing visual classification suggestions for inclusion-based concept clusters, comprising the steps of:

designating reference concepts each associated with a classification code;

grouping one or more of the reference concepts with a plurality of uncoded concepts into a grouped concept set;

generating clusters, each comprising a portion of the uncoded concepts and the reference concepts of the grouped concept set; and

providing a visual suggestion for assigning one of the classification codes to one of the clusters comprising visually representing each of the reference concepts in that cluster,

wherein the steps are performed by a suitably programmed computer.

12. A method according to claim **11**, further comprising: providing at least one of a presence and an absence of the concepts with each of the classification codes in that cluster; and

providing a number of the concepts with each of the classification codes in that cluster,

wherein the suggestion includes at least one of the number and the presence and the absence.

13. A method according to claim **12**, further comprising: providing a suggestion for classifying one of the uncoded concepts in that cluster comprising at least one of the number and the presence and the absence.

14. A method according to claim **11**, further comprising: receiving a user selection of one or more sources, custodians, content, and the classification codes of the reference concepts,

wherein the reference concepts are designated in accordance with the selection.

15. A method according to claim **11**, further comprising: receiving a user selection of one or more parameters for generating the clusters,

wherein the clustering is performed in accordance with the selection.

16. A method according to claim **11**, further comprising: determining a similarity between the clusters and organizing the clusters along one or more spines based on the similarity, each of the spines comprising a vector; and providing a visual classification suggestion for assigning one of the classification codes to one of the spines comprising visually representing the clusters along that spine and the reference concepts in the clusters along that spine.

17. A method according to claim **16**, further comprising: receiving a user selection of one or more of the clusters and spines;

displaying one or more of the clusters and spines within a compass;

emphasizing those of the clusters and spines displayed within the compass; and

deemphasizing those of the clusters and spines displayed outside of the compass.

18. A method according to claim **16**, further comprising: associating a label with each of the spines, each label associated with one or more documents from which one or more of the concepts in the clusters along that spine were extracted; and

displaying a list of the documents associated with one of the spine labels.

19. A method according to claim **11**, wherein the visual representation of the each classification codes comprises at least one of a symbol, shape, and color different from the visual representations of the remaining classification codes.

20. A method according to claim **11**, further comprising: identifying one of the concepts as a center of the cluster; identifying a neighborhood of similar reference concepts for the cluster based on the cluster center; and assigning one of the classification codes to the cluster based on the neighborhood.

* * * * *