

(51) International Patent Classification:
G06F 17/30 (2006.01) **G06F 17/21** (2006.01)(21) International Application Number:
PCT/US2009/034771(22) International Filing Date:
20 February 2009 (20.02.2009)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:

12/035,574	22 February 2008 (22.02.2008)	US
12/035,600	22 February 2008 (22.02.2008)	US
12/035,557	22 February 2008 (22.02.2008)	US
12/035,560	22 February 2008 (22.02.2008)	US
12/035,566	22 February 2008 (22.02.2008)	US
12/035,587	22 February 2008 (22.02.2008)	US
12/035,607	22 February 2008 (22.02.2008)	US
12/035,592	22 February 2008 (22.02.2008)	US
12/035,541	22 February 2008 (22.02.2008)	US
12/035,546	22 February 2008 (22.02.2008)	US
12/035,597	22 February 2008 (22.02.2008)	US

(71) Applicant (for all designated States except US): **TIGER-LOGIC CORPORATION** [US/US]; 25a Technology Drive, Irvine, CA 92618 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **DEXTER, Jeffrey, M.** [CA/US]; 176 Morning Glory, Rancho Santa Margarita, CA 92688 (US). **SMIK, Robert** [US/US]; 10831 Northoak Square, Cupertino, CA 95014 (US). **HYUN, Danny** [US/US]; 741 Wycliffe, Irvine, CA 92602 (US). **VEGERAJU, Srinivasa, R.** [IN/US]; 455 Crescent Avenue, #55, Sunnyvale, CA 94087 (US). **GARISH, Ilesh, H.** [IN/US]; 10669 Maplewood Road, #A, Cupertino, CA 95014 (US).(74) Agents: **SUN, Yalei, Y.** et al.; Morgan Lewis & Bockius Llp, 2 Palo Alto Square, 3000 El Camino Real, Suite 700, Palo Alto, CA 94306 (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,

[Continued on next page]

(54) Title: SYSTEMS AND METHODS OF IDENTIFYING CHUNKS WITHIN MULTIPLE DOCUMENTS




Figure 1

(57) Abstract: A computer identifies multiple resource identifiers, each resource identifier corresponding to a document at a respective data source. For at least one of the resource identifiers, the computer retrieves the corresponding document from the respective document source, identifies within the retrieved document a chunk that satisfies one or more user-specified search keywords, and displays the identified chunk and a link to the identified chunk within the document to the user.



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, 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 (Rule 48.2(g))*

Systems and Methods of Identifying Chunks within Multiple Documents

RELATED APPLICATIONS

[0001] This application relates to U.S. Patent Application, “Systems and methods of displaying document chunks in response to a search request,” filed on February 22, 2008 (attorney docket number 69218-5002-US), which is hereby incorporated by reference in its entirety.

[0002] This application relates to U.S. Patent Application, “Systems and methods of searching a document for relevant chunks in response to a search request,” filed on February 22, 2008 (attorney docket number 69218-5003-US), which is hereby incorporated by reference in its entirety.

[0003] This application relates to U.S. Patent Application, “Systems and methods of refining a search query based on user-specified search keywords,” filed on February 22, 2008 (attorney docket number 69218-5004-US), which is hereby incorporated by reference in its entirety.

[0004] This application relates to U.S. Patent Application, “Systems and methods of displaying and re-using document chunks in a document development application,” filed on February 22, 2008 (attorney docket number 69218-5005-US), which is hereby incorporated by reference in its entirety.

[0005] This application relates to U.S. Patent Application, “Systems and methods of performing a text replacement within multiple documents,” filed on February 22, 2008 (attorney docket number 69218-5006-US), which is hereby incorporated by reference in its entirety.

[0006] This application relates to U.S. Patent Application, “Systems and methods of refining chunks identified within multiple documents,” filed on February 22, 2008 (attorney docket number 69218-5007-US), which is hereby incorporated by reference in its entirety.

[0007] This application relates to U.S. Patent Application, “Systems and methods of pipelining multiple document node streams through a query processor,” filed on February 22,

2008 (attorney docket number 69218-5008-US), which is hereby incorporated by reference in its entirety.

[0008] This application relates to U.S. Patent Application, "Systems and methods of semantically annotating documents of different structures," filed on February 22, 2008
5 (attorney docket number 69218-5009-US), which is hereby incorporated by reference in its entirety.

[0009] This application relates to U.S. Patent Application, "Systems and methods of adaptively screening matching chunks within documents," filed on February 22, 2008
10 (attorney docket number 69218-5010-US), which is hereby incorporated by reference in its entirety.

[0010] This application relates to U.S. Patent Application, "Systems and methods of identifying chunks within inter-related documents," filed on February 22, 2008 (attorney docket number 69218-5011-US), which is hereby incorporated by reference in its entirety.

FIELD OF THE INVENTION

15 [0011] The present invention relates generally to the field of information retrieval in a computer system, in particular to systems and methods of locating information at different sources.

BACKGROUND OF THE INVENTION

20 [0012] The growth of information technology enables a user of a desktop or laptop computer to easily access information stored within a large number of documents at different locations such as the computer's local hard drive or a remote web server on the Internet. But quickly locating the information sought by the user within one or more documents remains a challenging task with today's information retrieval technologies.

25 [0013] In response to search keywords provided by a user, conventional web and desktop search engines typically return a list of document names with one or two sentences from each document that match the search keywords as search results. From the one or two matching sentences, the user often has trouble understanding the meaning of the search keywords in the context of the document. To determine whether the document has the user sought-after information, the user has no choice but to open the document using its native

application (e.g., the Microsoft Office application if the document is a Word document) and repeat the process if the document does not have the information sought by the user.

[0014] There are multiple issues with this approach. First, opening a document using its native application is a time-consuming operation. Second, and more importantly, the native application does not highlight any particular portion of the document that may contain the user-provided search keywords. To locate any search keywords within the document, the user has to do a new search of the document using a search tool of the native application. If the search tool can only look for multiple search keywords in exactly the same order (which is often the case), the user may end up finding nothing interesting in the document even if the document has a paragraph that contains the multiple search keywords but in a slightly different order. Alternatively, if the user limits the search to a subset of the multiple search keywords, many instances of the subset of search keywords may be in the document and the user could spend a significant effort before finding the document content of interest.

SUMMARY

[0015] The above deficiencies and other problems associated with conventional search tools are reduced or eliminated by the invention disclosed below. In some embodiments, the invention is implemented in a computer system that has a graphical user interface (GUI), one or more processors, memory and one or more modules, programs or sets of instructions stored in the memory for performing multiple functions. Instructions for performing these functions may be included in a computer program product configured for execution by one or more processors.

[0016] One aspect of the invention involves a computer-implemented method performed by a computer. The computer identifies multiple resource identifiers, each resource identifier corresponding to a document at a respective data source. For at least one of the resource identifiers, the computer retrieves the corresponding document from the respective document source, identifies within the retrieved document a chunk that satisfies one or more user-specified search keywords, and displays the identified chunk and a link to the identified chunk within the document to the user.

[0017] Another aspect of the invention involves a computer-implemented method performed by a client computer. After submitting one or more user-specified search keywords to a server computer, the client computer receives a set of search results from the

server computer, each search result identifying a document located at a respective document source that satisfies the search keywords in accordance with a first set of predefined criteria. For each identified document, the client computer retrieves the document from the corresponding document source, identifies a chunk within the document that satisfies the search query in accordance with a second set of predefined criteria, and displays the identified chunk and a link to the identified chunk within the document.

[0018] Another aspect of the invention involves a computer system. The computer system includes memory, one or more processors, and one or more programs stored in the memory and configured for execution by the one or more processors. The one or more programs include: instructions for identifying multiple resource identifiers, each resource identifier corresponding to a document at a respective data source; instructions for retrieving the corresponding document from the respective document source for at least one of the resource identifiers; instructions for identifying within the retrieved document a chunk that satisfies one or more user-specified search keywords; and instructions for displaying the identified chunk and a link to the identified chunk within the document to the user.

[0019] Another aspect of the invention involves a computer readable storage medium having stored therein instructions, which when executed by a computer system cause the computer system to: identify multiple resource identifiers, wherein each resource identifier corresponds to a document at a respective data source; retrieve the corresponding document from the respective document source for at least one of the resource identifiers; identify within the retrieved document a chunk that satisfies one or more user-specified search keywords; and display the identified chunk and a link to the identified chunk within the document to the user.

[0020] Another aspect of the invention involves a graphical user interface on a computer display. The graphical user interface includes one or more document links, each document link having one or more associated chunks identified within the corresponding document as satisfying one or more user-specified search keywords. Each identified chunk has an associated chunk link and includes terms matching each of the user-specified search keywords. In some embodiments, the terms are highlighted in the chunk in a visually distinguishable manner. In response to a user selection of a chunk's chunk link, the corresponding document is displayed in a window on the computer display, and at least a portion of the chunk is highlighted in the window.

[0021] Some embodiments may be implemented on either the client side or the server side of a client-server network environment.

BRIEF DESCRIPTION OF THE DRAWINGS

[0022] The aforementioned features and advantages of the invention as well as additional features and advantages thereof will be more clearly understood hereinafter as a result of a detailed description of preferred embodiments when taken in conjunction with the drawings.

[0023] Figure 1 is a block diagram of an exemplary computer system that includes a front end, a search server including a query engine, a cache engine, an index database, and a stream engine, and one or more data sources in accordance with some embodiments.

[0024] Figure 2 is a flowchart illustrative of how the front end processes user-provided search keywords in accordance with some embodiments.

[0025] Figure 3 is a flowchart illustrative of how the query engine generates search criteria for the user-provided search keywords in accordance with some embodiments.

[0026] Figure 4 is a flowchart illustrative of how the cache engine produces a set of candidate document identifiers for the user-provided search keywords in accordance with some embodiments.

[0027] Figure 5 is a flowchart illustrative of how the stream engine processes candidate documents retrieved from different data sources in accordance with some embodiments.

[0028] Figure 6 is a flowchart illustrative of how the cache engine processes the candidate documents coming out of the stream engine in accordance with some embodiments.

[0029] Figure 7 is a flowchart illustrative of how the query engine identifies relevant chunks within the candidate documents in accordance with some embodiments.

[0030] Figure 8A is a flowchart illustrative of how the stream engine generates semantic data models for different types of documents in accordance with some embodiments.

[0031] Figure 8B is a flowchart illustrating a first embodiment of how the query engine identifies a relevant chunk within a node stream representing a candidate document.

[0032] Figure 8C is a flowchart illustrating a second embodiment of how the query engine identifies a relevant chunk within a node stream representing a candidate document.

5 [0033] Figure 9A is a flowchart illustrative of how the stream engine processes multiple candidate documents to identify candidate chunks in accordance with some embodiments.

[0034] Figure 9B is an exemplary HTML document to be processed by the stream engine as shown in Figure 9A in accordance with some embodiments.

10 [0035] Figure 10A is a block diagram illustrative of how a query mediator coordinates the query engine and the stream engine to identify chunks within a node stream representing a candidate document in accordance with some embodiments.

[0036] Figure 10B is a flowchart illustrative of how the stream engine divides the node stream into multiple sub-streams using a filter model in accordance with some
15 embodiments.

[0037] Figure 11A is an exemplary XML document to be processed by the stream engine and the query engine in accordance with some embodiments.

[0038] Figure 11B is an exemplary XQuery to be applied to the XML document in accordance with some embodiments.

20 [0039] Figure 11C is a table of input sequences defined by the query engine in accordance with some embodiments.

[0040] Figure 11D is a flowchart illustrative of how the query engine processes node sub-streams at different input sequences in accordance with some embodiments.

[0041] Figure 11E is a block diagram illustrative of how a node stream corresponding
25 to the XML document is divided into multiple node sub-streams by a finite state machine in accordance with some embodiments.

[0042] Figure 11F is a block diagram illustrative of the input sequences and their associated node sub-streams after the first candidate chunk in the XML document is processed in accordance with some embodiments.

[0043] Figure 11G is the search result of applying the XQuery to the node sub-streams derived from XML document in accordance with some embodiments.

[0044] Figure 12A is a flowchart illustrative of a first process of identifying one or more documents, each document having one or more chunks that satisfy user-specified search keywords, in accordance with some embodiments.

[0045] Figure 12B is a flowchart illustrative of a second process of identifying one or more document, each document having one or more chunks that satisfy user-specified search keywords, in accordance with some embodiments.

[0046] Figures 12C through 12J are screenshots of a graphical user interface on a computer display illustrative of features associated with the processes as shown in Figures 12A and 12B in accordance with some embodiments.

[0047] Figure 13A is a flowchart illustrative of a first process of identifying within a document one or more chunks that satisfy user-specified search keywords in accordance with some embodiments.

[0048] Figure 13B is a flowchart illustrative of a second process of identifying within a document one or more chunks that satisfy user-specified search keywords in accordance with some embodiments.

[0049] Figures 13C through 13G are screenshots of a graphical user interface on a computer display illustrative of features associated with the processes as shown in Figures 13A and 13B in accordance with some embodiments.

[0050] Figure 14 is a flowchart illustrative of a process of modeling a document and identifying within the document one or more chunks that satisfy user-specified search keywords in accordance with some embodiments.

[0051] Figure 15 is a flowchart illustrative of a process of customizing a search query based on user-specified search keywords in accordance with some embodiments.

[0052] Figure 16A is a flowchart illustrative of a process of displaying and re-using search results based on user instructions in accordance with some embodiments.

[0053] Figures 16B through 16J are screenshots of a graphical user interface on a computer display illustrative of features associated with the process as shown in Figure 16A
5 in accordance with some embodiments.

[0054] Figure 17A is a flowchart illustrative of a process of finding and replacing text strings in connection with a set of search results based on user instructions in accordance with some embodiments.

[0055] Figure 17B is a flowchart illustrative of a process of finding and replacing text
10 strings within a set of documents based on user instructions in accordance with some embodiments.

[0056] Figures 17C through 17E are screenshots of a graphical user interface on a computer display illustrative of features associated with the processes as shown in Figures 17A and 17B in accordance with some embodiments.

[0057] Figure 18A is a flowchart illustrative of a first process of narrowing search
15 results based on user instructions in accordance with some embodiments.

[0058] Figure 18B is a flowchart illustrative of a second process of narrowing search results based on user instructions in accordance with some embodiments.

[0059] Figures 18C through 18D are screenshots of a graphical user interface on a
20 computer display illustrative of features associated with the processes as shown in Figures 18A and 18B in accordance with some embodiments.

[0060] Figure 19 is a flowchart illustrative of a process of alternatively processing document node streams in accordance with some embodiments.

[0061] Figure 20 is a flowchart illustrative of a process of semantically and
25 contextually annotating documents of different structures in accordance with some embodiments.

[0062] Figure 21A is a flowchart illustrative of a first process of screening matching chunks within a document based on predefined criteria in accordance with some embodiments.

[0063] Figure 21B is an exemplary HTML document illustrative of the process as shown in Figure 21A in accordance with some embodiments.

[0064] Figure 21C is a flowchart illustrative of a second process of screening matching chunks within a document based on predefined criteria in accordance with some embodiments.

[0065] Figure 21D is a screenshot of a graphical user interface on a computer display illustrative of features associated with the processes as shown in Figures 21A and 21B in accordance with some embodiments.

[0066] Figure 22A is a flowchart illustrative of a process of identifying contents matching a search request within a plurality of inter-related documents in accordance with some embodiments.

[0067] Figures 22B through 22D are screenshots of a graphical user interface on a computer display illustrative of features associated with the process as shown in Figure 22A in accordance with some embodiments.

[0068] Figure 23 is a block diagram of an exemplary document search server computer in accordance with some embodiments.

[0069] Figure 24 is a block diagram of an exemplary client computer in accordance with some embodiments.

[0070] Like reference numerals refer to corresponding parts throughout the several views of the drawings.

DESCRIPTION OF EMBODIMENTS

[0071] Reference will now be made in detail to embodiments, examples of which are illustrated in the accompanying drawings. In the following detailed description, numerous specific details are set forth in order to provide a thorough understanding of the subject matter presented herein. But it will be apparent to one skilled in the art that the subject

matter may be practiced without these specific details. In other instances, well-known methods, procedures, components, and circuits have not been described in detail so as not to unnecessarily obscure aspects of the embodiments.

[0072] Figure 1 is a block diagram of an exemplary computer system 100 that includes a front end 15, a search server 20, and one or more data sources 70 in accordance with some embodiments. The front end 15 is a software application configured to receive and process input from a user 10 such as search keywords and present search results to the user 10. The search server 20 further includes a query engine 30, a cache engine 40, an index database 50, and a stream engine 60. The data sources 70 include storage devices such as file systems on hard drives accessible to the computer system 100 and remote web servers on the Internet.

[0073] At runtime, the front end 15 forwards the user-provided search keywords to the search server 20 in the form of a search query. In response, different components within the search server 20 work in concert to identify a set of candidate documents that matches the search keywords and retrieve the contents of the candidate documents from their respective locations at local and/or remote data sources 70. The different components within the search server 20 then search within the retrieved document contents for chunks that match the search keywords and return the identified chunks to the front end 15 in the form of search results.

[0074] In this application, a document is generally referred to as a data entity that has textual content, such as a Microsoft Office document, a plain-text document, a PDF document, an email or text message, a web page, etc. A “candidate chunk” within a document is a portion of the document that is semantically and contextually regarded as a unit of textual content by one skilled in the relevant art. For example, within a Word document, a sentence, a paragraph, a table, a figure’s caption, the document’s title, header, and footer are candidate chunks. Similarly, a slide within a PowerPoint document, a bullet point within the slide, and a cell or a row within an Excel spreadsheet are also candidate chunks. A “chunk” or more specifically a “relevant chunk” served as part of the search results is a candidate chunk that satisfies the search keywords in accordance with predefined search criteria, e.g., if the candidate chunk includes at least one instance of one of the search keywords.

[0075] Figure 2 is a flowchart illustrative of how the front end 15 processes user-provided search keywords in accordance with some embodiments. After receiving the search keywords (201), the front end 15 generates a search query using the search keywords (203). Depending on the type of data to be processed by the search server 20, the search query can be written in different query languages such as structured query language (SQL) for relational databases or XQuery for XML data sources. The front end 15 then submits the search query to the query engine 30 within the search server 20 (205) for further processing.

[0076] Figure 3 is a flowchart illustrative of how the query engine 30 generates search criteria for the user-provided search keywords in accordance with some embodiments. After receiving the search query (302), the query engine 30 analyzes the query (304) and generates optimized search criteria (306). In some embodiments, the query engine 30 also generates one or more path filters from the search query (308). The path filters are derived from the user-provided search keywords and search options. The stream engine 60 employs the path filters to exclude document content that is not part of any candidate chunks. A more detailed description is provided below in connection with Figures 10 and 11. The query engine 30 submits both the search criteria and the path filters to the cache engine 40 (310).

[0077] In some embodiments, the query engine 30 generates an optimized execution plan for the query according to the capabilities of other components within the search server 20. For example, if the search query contains a predicate limiting the search to documents at the local hard drive that have been updated within the last two days, the query engine 30 has two options. One option is that the query engine 30 keeps the predicate to itself and waits to apply the predicate to the candidate chunks. In this case, the search server 20 (especially the stream engine 60) may have processed more candidate documents than necessary. The other option is that the query engine 30 pushes the predicate down to the file system managing the local hard drive through the index database 50. In this case, only candidate documents that have been updated within the last two days are processed and the stream engine 60 is relieved from additional, unnecessary workload.

[0078] Figure 4 is a flowchart illustrative of how the cache engine 40 produces a set of candidate document identifiers for the user-provided search keywords in accordance with some embodiments. After receiving the search criteria from the query engine (401), the cache engine 40 submits the search criteria to the index databases 50. In some embodiments, the index databases include both a local index database and a remote index database. The

local index database manages index information of documents at the local hard drive and the remote index database manages index information of documents at a remote document server. In some embodiments, the remote index database refers to the index database of a third-party search engine.

5 **[0079]** For given user search criteria, the cache engine 40 may search the local index database (403) if the user is looking for documents at the local hard drive or the remote index database (405) if the user is submitting a search request to a third-party search engine or both. From these index databases 50, the cache engine 40 receives respective search results (407), e.g., in the form of a set of document references such as URIs, and identifies a candidate
10 document identifier within each search result (409). Note that a candidate document is a document that matches the search query at the document level, but may or may not at the chunk level. For example, a PowerPoint document that has slide A matching one search keyword and slide B matching another search keyword is a candidate document, but does not have any chunk matching the search query. In some embodiments, a universal resource
15 identifier (URI) is used as a document identifier. Thus, documents at the local hard drive and remote web servers can be referenced in the same manner.

[0080] In some embodiments, the search results returned by the index databases 50 are ordered by the corresponding candidate documents' relevancy to the search query. Many well-known algorithms for determining a document's relevancy to a search query can be
20 found in the classic book entitled "Automatic Information Organization and Retrieval" by G. Salton, McGraw-Hill, New York, 1968, which is incorporated here by reference in its entirety.

[0081] In some embodiments, a candidate document's relevancy is at least in part ranked by the past user activities on the candidate document. For example, a candidate
25 document that has been recently accessed by the user, such as browsing, copying and updating, is given a higher rank than another candidate document that has never been accessed by the user before. In one embodiment, a candidate document's ranking score is determined by combining the following two pieces of information:

- The frequency of a search keyword in the document – For each keyword, the
30 index database 50 may keep information such as a total count of occurrences of the keyword in a number of documents and a per-document count of the

occurrences of the keyword. By combining the frequencies of different search keywords within the same document, a basic ranking score of the document is computed using a generic inverse frequency algorithm.

- The personalized usage weight of the document – A respective number of points are assigned to each operation the user applies to the document. For example, the preview operation of a particular document is given two points, the re-use of content from the previewed document is given three points, and the re-use of a specific chunk within the document is given four points. The total points assigned to the document, when compared against the total points allocated for the corresponding document type, yields a personalized ranking score for the document, which may be combined with the aforementioned basic ranking score to generate a customized ranking score for the document.

[0082] In some embodiments, a document's relevancy to a search query is not solely determined at the document level but is, at least in part, determined at the chunk level. For example, the index database 50 may maintain information about the locations of candidate chunks within each document as well as the distinct ranking information of different candidate chunks within the same document relative to different search keywords, thereby making it possible to return the relevant chunks within the same document in an order consistent with their respective chunk-level ranking scores.

[0083] The cache engine 40 submits a set of candidate document identifiers and path filters, if any, generated by the query engine 30 to the stream engine 60 for further processing (411).

[0084] For illustration, the aforementioned processes in connection with Figures 2 through 4 are collectively referred to as the "downstream processes 25" as shown in Figure 1.

The input to the downstream processes is a search request including one or more search keywords and its output is a set of candidate document identifiers that identify candidate documents satisfying the search keywords. For example, a document is deemed to be a candidate document if it includes at least one instance of each search keyword. But the fact that each search keyword has a match in a candidate document does not guarantee that the candidate document has a chunk relevant to the search request.

[0085] As noted above, identifying a chunk within a document requires semantic information about the document. Such information is not available in the index database 50. To find out whether a candidate document has any relevant chunk, the search server 20 needs to retrieve the document and analyze the document's structure and content to understand the relationship between the document's structure and the document's content. The processes of retrieving a document, determining whether the document has any relevant chunks, and identifying and returning these relevant chunks to the requesting user are collectively referred to as the "upstream processes 35" as shown in Figure 1.

[0086] Figure 5 is a flowchart illustrative of how the stream engine 60 processes candidate documents retrieved from data sources 70 in accordance with some embodiments.

[0087] After receiving the candidate document identifiers and optional path filters from the cache engine (502), the stream engine 60 starts retrieving the candidate documents identified by the document identifiers from respective data sources, including retrieving some candidate documents from local data sources (504) and/or retrieving some candidate documents from remote data sources (506). In some embodiments, local data sources include any storage devices affiliated with the computer system 100, such as hard disk and CD/DVD drives, and remote data sources include any storage devices that can be accessed by the computer system 100 through wired and/or wireless network, such as a web server on the Internet and/or a network storage device on the Intranet.

[0088] In some embodiments, a specific user instruction may limit the document search to local or remote data sources. As shown in Figure 16B, if the user specifies that the type of the documents to be searched are Word documents, the stream engine 60 will retrieve only Word candidate documents from the local data source such as the local file system. For each candidate document identifier, the stream engine 60 submits a request for the corresponding candidate document to the file system and waits for the file system to return the candidate document. But if the user clicks the checkbox next to a web source such as "Source A," the stream engine 60 will retrieve the candidate documents identified by Source A from their respective remote web hosts. For example, if the candidate document is an HTML document hosted by a web server, the stream engine 60 submits an HTTP request to the web server for the HTML document and waits for an HTTP response including the HTML document from the web server. In some embodiments, the user instruction may

explicitly or implicitly request that candidate documents be retrieved from both local and remote data sources.

[0089] In some embodiments, the stream engine 60 submits multiple requests for different candidate documents in parallel or sequentially to the respective targeted data source(s). The candidate documents are then retrieved from the respective data sources and processed by the stream engine 60 in parallel or sequentially. For illustration, the following description in connection with Figure 5 focuses on a single candidate document. But this by no means limits the present application to processing documents one by one sequentially. As will become more apparent in connection with Figures 9A through 9B below, it is more efficient to process multiple candidate documents from different data sources in parallel in some embodiments.

[0090] Referring again to Figure 5, the stream engine 60 performs the following operations on each candidate document retrieved from a data source:

[0091] 1. Convert the candidate document into a node stream (508);

[0092] To reduce the computer system 100's response latency, the stream engine 60 starts converting immediately after receiving a portion of the candidate document, without waiting for the entire candidate document to be retrieved. A more detailed description of the conversion is provided below in connection with Figure 8A.

[0093] 2. Identify candidate chunks in the node stream (510);

[0094] As noted above, a candidate document includes one or more candidate chunks. A candidate chunk within the document, if identified as satisfying the user-specified search keywords, is usually more relevant to the user's search interest and therefore more useful to the user. A more detailed description of this operation is provided below in connection with Figures 8A and 9A.

[0095] 3. Apply the optional path filters to the node stream (512); and

[0096] For a user-specified search query, certain portions of the node stream are potentially relevant and other portions are completely irrelevant. It could be an efficiency gain if these irrelevant portions are excluded from further consideration. The path filters generated by the query engine (operation 308 in Figure 3) can be used to divide the node

stream into multiple node sub-streams, thereby eliminating the irrelevant portions of the node stream. In some embodiments, this procedure is optional if the query engine 30 generates no path filter. A more detailed description of the conversion is provided below in connection with Figures 10A-10B and 11A-11G.

5 **[0097]** 4. Submit the node stream (or sub-streams) to the cache engine (514).

[0098] After performing the operations above, the stream engine 60 submits the node stream or sub-streams to the cache engine 40. As will be explained below in connection with Figure 6, the cache engine 40 may or may not do anything depending on whether it needs to index the document represented by the node stream. If it does, the cache engine 40 invokes
10 the index database 50 to generate new indexes or update existing indexes for the document. Otherwise, the cache engine 40 simply forwards the node stream or sub-streams to the query engine 30 for further processing, which is provided below in detail in connection with Figures 8A-8C and 11A-11G.

[0099] Figure 6 is a flowchart illustrative of how the cache engine 40 processes the
15 candidate documents coming out of the stream engine 60 in accordance with some embodiments.

[00100] After receiving the node stream or sub-streams corresponding to a candidate document (601), the cache engine 40 performs different operations based on the type and status of the candidate document as well as its destination. For example, if the candidate
20 document is a Word document found in the local hard drive of the computer system 100 and has not been indexed or its indexes are deemed stale, the cache engine 40 will request that the index database 50 generate new indexes or update existing indexes for the candidate document (603). Depending on whether the document is requested by an end user through the front end 15 or a software agent monitoring the index database 50, the cache engine 40
25 may or may not return the node stream to the query engine 30 for further processing (605).

[00101] If the candidate document is an HTML document at a remote web server, which is identified through a third-party document source, it may be optional to index the HTML document. If so, the node stream or sub-streams will be returned to and further processed by the query engine 30 to determine whether it has any relevant chunk (605).

[00102] In sum, in some embodiments the cache engine 40 plays a relatively lightweight role in the upstream processes 35 especially if the candidate document is retrieved from a remote data source to satisfy an end user's search request and the computer system 100 does not intend to index the document. Therefore, some of the embodiments
5 below assume that the stream engine 60 directly communicates with the query engine 30 for clarity.

[00103] Figure 7 is a flowchart illustrative of how the query engine 30 identifies relevant chunks within the candidate documents in accordance with some embodiments.

[00104] Upon receipt of the node stream or sub-streams (e.g., path filtering at the
10 stream engine 60) (702), the query engine 30 traverses the node stream and compares the candidate document's content with the user-specified search keywords. If a match is found, the query engine 30 identifies the candidate chunk within the document corresponding to the match as one relevant chunk (704) and returns the identified chunk to the front end 15 to be displayed to the end user (706).

[00105] In some embodiments (as shown Figure 7), the query engine 30 returns any relevant chunk to the front end 15 as soon as the chunks are identified and this process repeats until the query engine 30 completely traverses the candidate document (708). In some other embodiments, the query engine 30 defers returning any chunk to the front end 15 until a more specific relevant chunk is found in the node stream. A more detailed description
15 of these two approaches is provided below in connection with Figures 8B and 8C, respectively.

[00106] As noted above, candidate documents arriving at the stream engine 60 are each converted into a node stream. The node stream is an instance of a data model of the corresponding candidate document. For example, the XQuery data model of an XML
25 document is a tree of nodes. The types of the nodes that appear at different hierarchical levels of the tree include: document, element, attribute, text, namespace, processing instruction, and comment. Any node in the tree has a unique node identity. The data model not only preserves the XML document's entire content but also has metadata derived from sources such as XML tags for identifying candidate chunks subsequently.

[00107] Unfortunately, not all candidate documents are structured like an XML document. For example, a plain-text document is completely unstructured such that it does

not have any metadata embedded therein defining a hierarchical structure for the document. Without any pre-processing, a node stream corresponding to the plain-text document loses the semantic information intrinsic in the content distribution of the document such that it is difficult to identify any chunk such as paragraph, headline, or title, within the node stream to satisfy a search query. PDF documents have similar problems that make it challenging to find relevant chunks within a PDF document.

[00108] Between the structured XML documents and the unstructured plain-text documents are semi-structured documents such as HTML documents. Unlike the plain-text document, an HTML document has a hierarchical structure defined by metadata embedded in the HTML document. But the metadata in the HTML document usually does not have a deterministic relationship with the content data as the metadata in an XML document has. The same HTML tag can be used purely for web page layout purpose at one location while carrying a semantic connotation at another location within the same document.

[00109] To expedite the upstream processes and accommodate more types of documents in the future, it is desired to have a unified approach such that different types of documents are processed in the same manner.

[00110] Figure 8A is a flowchart illustrative of how the stream engine 60 generates semantic data models for different types of documents in accordance with some embodiments.

[00111] After receiving the raw data of a candidate document, the stream engine 60 transforms the raw data into an instance of a data model of structured or semi-structured data (801). In some embodiments, this operation is straightforward if the candidate document is already a structured document like a Microsoft Office 2007 document. In some other embodiments, this operation is necessary if the candidate is a plain-text document without any structure-related metadata. In this case, the stream engine 60 may insert metadata into the document that defines a hierarchical structure for the document's content.

[00112] Based on the class of the raw data (803), the stream engine 60 then performs different sets of operations to the data model instance generated previously. For example, as noted above, the candidate document may be classified into one of three categories:

- unstructured documents (805-A) such as plain-text and PDF;

- semi-structured documents (805-B) such as HTML and RTF; and
- structured documents (805-C) such as XML and Office 2007.

[00113] For an unstructured document, the stream engine 60 applies a set of natural language and common formatting based heuristic rules to separate text within the document into separate candidate chunks (807). For example, one heuristic rule for identifying paragraphs stipulates that any two text segments separated by symbols such as an end-of-line (EOL) character or a blank line correspond to at least two separate paragraphs. Another heuristic rule stipulates that a text segment matching a predefined text pattern is deemed to be a candidate chunk. Consider the following text segment that has two hyphens, one at the start of a new line:

- This is a bullet list.
- What about a page number?

In this case, each line by itself may be a candidate chunk (although it may or may not be deemed to be a paragraph). The fact that the two lines have the same text pattern, i.e., a hyphen at the start of a new line followed by a text string, may indicate that the entire text segment is a candidate chunk at one level of the document's hierarchical structure and each line is also a candidate chunk at a lower level of the hierarchical structure.

[00114] Similarly, for a semi-structured document, the stream engine 60 has another set of heuristic rules based on the type of the semi-structured document (809). For a node stream corresponding to an HTML document, the stream engine 60 identifies candidate chunk break nodes within the node stream both dynamically and statically.

[00115] For example, the <p> tag defines a paragraph within the HTML document and it is deemed to be a candidate chunk break node. Whenever the <p> tag appears in an HTML document, the subsequent document segment following this <p> tag and before another <p> tag is identified as a candidate chunk.

[00116] Note that there are many ways of identifying chunk break nodes within a semi-structured document known to one skilled in the art. In some embodiments, the stream engine 60 applies different sets of customized heuristic rules to different types of documents. For a structured document or a semi-structured document for which there is no customized solution, the stream engine 60 assumes that there is a strongly-deterministic relationship

between the document's content and the document's metadata and a generic set of rules is applied to the data model instance to identify possible candidate chunks in the candidate document.

[00117] By traversing the node stream, the stream engine 60 generates a data model instance for the candidate document that includes semantic annotation (811). Subsequently, the semantically-annotated node stream is fed into the query engine 30. The query engine 30 then applies a search query to the node stream to identify among the candidate chunks any that satisfy the search query.

[00118] As noted above in connection with Figure 7, the query engine 30 does not have to wait until it traverses the entire node stream before returning any relevant chunk to the front end 15. Below are two embodiments of how the query engine 30 returns identified chunks after a respective condition is met and before the entire node stream is traversed.

[00119] Assume that the search query has two keywords, "raining" and "data," and the exemplary candidate document is as follows:

```
<c0>
    It's raining outside ...
    <c1>
        For XML-based data management,
        Raining Data is your choice.
    </c1>
</c0>
```

[00120] Figure 8B is a flowchart illustrating a first embodiment of how the query engine identifies a relevant chunk within a node stream representing a candidate document.

[00121] The query engine 30 starts the search after receiving a node stream corresponding to the candidate document above (821). If no more nodes are in the node stream (823, no), the query engine 30 assumes that it has completely traversed the node stream and the search stops (825). Otherwise, the query engine 30 processes the next node in the stream (827).

[00122] Before any further processing, the query engine 30 checks whether it is in the middle of accumulating nodes (829). In some embodiments, the query engine 30 begins accumulating nodes after it encounters the chunk break node of the first candidate chunk in

the node stream. In this example, the chunk break node of the first candidate chunk is the `<c0>` tag, which is the first node in the stream, and the accumulation has not started yet (829, no).

[00123] Next, the query engine 30 checks whether the node is a text node (839). Since the `<c0>` tag is not a text node (839, no), the query engine 30 updates the current node path to be `"/c0"` (841) and checks whether the current node is part of a candidate chunk (843). Because the `<c0>` tag is the chunk break node of the first candidate chunk (843, yes), the query engine 30 marks the current node path as corresponding to a candidate chunk (845) and then starts node accumulation immediately (847).

[00124] Following the `<c0>` tag node, the next node to be processed by the query engine 30 is a text node including "It's raining outside ...". In this case, because the accumulation has already started (829, yes), the query engine checks if the text node is part of a relevant chunk (831). But since no search keyword has been matched so far (831, no), the query engine 30 accumulates the text node (837). Because this is a text node (839, yes), the query engine 30 then checks whether it is in a candidate chunk (849).

[00125] In this case, the text node is in a candidate chunk (849, yes). The query engine applies the search query to the text node (851). But because only the keyword "raining" finds a match in the text string, which is a partial match of the search query, no relevant chunk has been found yet (853, no) and the query engine 30 returns to process the next node in the sub-stream (823). In some embodiments, the query engine 30 records the partial match result for subsequent use.

[00126] When the query engine 30 receives the second text node including the text string "For XML-based data management," it repeats the same processing steps 827 through 853 described above. In this case, because the two text nodes in combination match both keywords, a relevant chunk and its associated node path `"/c0/c1"` are identified (855). Next, the query engine 30 processes the third text node including the text string "Raining Data is your choice." Because the third node is already in a relevant chunk (831, yes), the query engine 30 checks whether the relevant chunk is completed (833). In some embodiments, a chunk is completed if the query engine encounters a node including the end tag of a candidate chunk, e.g., `</c0>` or `</c1>`.

[00127] In this case, because the query engine 30 has not seen any end tag yet (833, no), the process continues and the second and third text nodes in combination also match the two keywords because both the second and third text nodes are within the second candidate chunk (<c1>, </c1>), which is a descendent of the first candidate chunk (<c0>, </c0>). In some embodiments, if there is a hierarchical relationship between multiple relevant chunks, the query engine 30 first completes the relevant chunk at the lowest level, which is also referred to as the more specific relevant chunk, and then outputs this more specific relevant chunk to the front end 15 (835). In this example, the more specific relevant chunk is

<c1>

For XML-based **data** management,
Raining Data is your choice.

</c1>

[00128] Note that the query engine 30 does not necessarily stop after outputting the more specific relevant chunk (835). Rather, the query engine 30 proceeds to the next node that includes the </c0> tag. As a result, the less specific relevant chunk (as will be described below in connection with Figure 8C) is the next relevant chunk to be output.

[00129] In some embodiments, the query engine 30 outputs this relevant chunk to the front end 15. As a result, the front end 15 may ultimately display two relevant chunks to the end user. Alternatively, the front end 15 may compare the two relevant chunks before displaying them and choose only one of them, e.g., the more specific one above or the second broader one, to be displayed. In some other embodiments, the query engine 30 may choose not to output the second relevant chunk to the front end 15 if it determines that the first one is sufficient to satisfy the end user's search interest.

[00130] Figure 8C is a flowchart illustrating a second embodiment of how the query engine 30 identifies a relevant chunk within a node stream representing a candidate document. This embodiment is similar to the embodiment described above in connection with Figure 8C except that, after a relevant chunk is identified, the query engine 30 immediately starts outputting nodes in the identified chunk (895) without waiting for the completion of the relevant chunk (835 in Figure 8B). Moreover, the query engine 30 also outputs subsequent nodes within the same relevant chunk (877), if there are any, without waiting for the completion of the relevant chunk (835 in Figure 8B).

[00131] Using the same exemplary candidate document above, the query engine 30 outputs the relevant query when it encounters the second text node including the text string “For XML-based data management” because both search keywords have matches in the relevant chunk. Although this relevant chunk might not be as satisfactory as the more specific one, the response latency of this second embodiment is usually shorter than the response latency of the first embodiment.

[00132] As described above in connection with Figure 5, the stream engine 60 receives one or more candidate document identifiers such as URIs from the cache engine 40. For each URI, the stream engine 60 submits a request to a respective data source to retrieve the corresponding candidate document hosted by the data source. If multiple requests are submitted to different data sources within a short time period or even in parallel, the requested candidate documents may arrive at the stream engine 60 simultaneously or nearly so.

[00133] In some embodiments, a candidate document such as a web page at a remote web server is divided into multiple data packets at the respective data source and transmitted back to the stream engine 60 one packet at a time. But due to network traffic jams, the data packets from a single data source may arrive at the stream engine 60 out of their original transmission order and the data packets from different data sources may arrive at the stream engine 60 at different rates. The query engine 30, however, usually requires that the data packets of a particular candidate document be analyzed sequentially to identify relevant chunks therein and present them to the end user. This is especially true if a text node that satisfies a search query is larger than the maximum size of a packet and therefore has to be allocated into multiple data packets for network transmission.

[00134] As a result, such a deadlock situation often occurs: on the one hand, the stream engine 60 is waiting for a data packet from a first data source to support the query engine 30’s operation; on the other hand, the data packet cannot arrive at the stream engine 60 on time due to network delay. At the same time, multiple data packets from a second data source may have arrived at the stream engine 60, but they are postponed from further processing although they might contain a relevant chunk. If this issue is not appropriately resolved, it would significantly increase the computer system’s response latency, causing a less satisfactory user experience to the end user.

[00135] Figure 9A is a flowchart illustrative of how the stream engine 60 processes multiple candidate documents to identify candidate chunks in accordance with some embodiments. For illustration, assume that the stream engine 60 receives two URIs, UA and UB, from the cache engine 40, each identifying a candidate document at a respective data source. In reality, the stream engine 60 may receive N URIs and therefore process N node streams at the same time, N being an integer number varying from a few to a few hundred.

[00136] Initially, whenever it has bandwidth for processing more URIs (902, yes), the stream engine 60 checks whether there are any URI available for processing (904). If not (904, no), the stream engine 60 processes existing node streams (912). In this example, both UA and UB are available (904, yes). The stream engine 60 chooses one of them (906), e.g., UA, and checks the availability of the corresponding data source (908). If the data source is not available (908, no), the stream engine 60 then returns to process the next URI (902). Otherwise (908, yes), the stream engine 60 generates a node stream for UA (910) and then returns to process the next URI (902). At the end, for each candidate document, the stream engine 60 generates a node stream to manage incoming data packets corresponding to the document.

[00137] In some embodiments, the stream engine 60 checks the availability of a data source repeatedly until a predefined condition is met, e.g., the time elapsed from the first check to the last check is beyond a threshold level. If no, the stream engine 60 assumes that the corresponding document is not available and devotes its resources to processing other available candidate documents. Note that the stream engine 60 may perform the same or similar exercise repeatedly for each data source from which it has already received data packets. If the stream engine 60 fails to receive any data packet from a data source for a predefined time period, the stream engine 60 may assume that this data source is no longer available and free any resources allocated for this data source and the corresponding node stream. By doing so, the overall response latency is below a level of reasonable tolerance.

[00138] In this example, assume that the stream engine 60 chooses to work on one of the two available node streams (902), e.g., the UA node stream, and the first data packet has arrived (916). The stream engine 60 processes the data packet (920), such as verifying its accuracy, extracting the raw data of the candidate document from the data packet, and converting the raw data into one or more nodes in the UA node stream. Next, the stream

engine 60 parses the next node in the UA node stream (922) to identify candidate chunks within the node stream.

[00139] For each node in the UA node stream, the stream engine 60 determines if it corresponds to a new candidate chunk (926) or is within an existing candidate chunk (928) until finishing the last one in the node stream (924). In either case (926, yes) (928, yes), the stream engine 60 accumulates the node into the candidate chunk (930) and then determines whether it has reached the end of the corresponding candidate chunk. If so (932, yes), the stream engine 60 sends the candidate chunk to the query engine 30 for further processing (934), e.g., determining whether the candidate chunk is a chunk relevant to the user-specified search keywords.

[00140] In some embodiments, after sending the candidate chunk to the query engine 30, the stream engine 60 returns to parse the next one in the node stream (922) and repeats the aforementioned operations until it exhausts the last one in the node stream. In other words, the stream engine 60 and the query engine 30 may proceed in parallel and independently. This mechanism or the like can be very efficient if the computer system 100 has enough resources, e.g., multiple processors (including co-processors) and/or a large amount of memory, or if different components within the computer system 100, e.g., the stream engine 60 and the query engine 30, operate on separate threads and there is a carefully-maintained thread boundary between the two.

[00141] In some other embodiments, the stream engine 60 pauses after passing one candidate chunk to the query engine 30 (934) and resumes processing the node stream after it receives feedback from the query engine 30 (936). This mechanism or the like may be more feasible if the computer system 100 has limited resources, e.g., a single processor and/or limited memory. In this case, the stream engine 60 and the query engine 30 share the same thread. As a result, the computer system 100 may only need a small amount of memory to have a reasonably efficient performance. A more detailed description of this feedback-based scheme is provided below in connection with Figures 10A-10B and 11A-11G.

[00142] As noted above, a candidate chunk is semantically and contextually a unit within a candidate document. The process described above in connection with Figure 8A may annotate multiple nodes in a node stream, each annotated node corresponding to a candidate chunk. These candidate documents may be associated with different levels of a

hierarchical data model of the candidate document. In other words, a small candidate chunk can be a descendant of a large candidate chunk.

[00143] Figure 9B is an exemplary HTML document to be processed by the stream engine as shown in Figure 9A in accordance with some embodiments. From applying the corresponding heuristic rules to this HTML document, the stream engine 60 identifies nine candidate chunks 942 through 958. Note that the first node within each candidate chunk is highlighted in Figure 9B. For example, the first node of the candidate chunk 942 is the <table> tag 960 and the first node of the candidate chunk 956 is the <p> tag 974. The candidate chunk 956 and the candidate chunk 958 are at the same level in the hierarchical data model, both of which are descendants of the larger candidate chunks such as the candidate chunk 954.

[00144] When applying the process in Figure 9A to the HTML candidate document in Figure 9B, the stream engine 60 receives the node including the <table> tag 960 and a new candidate chunk 942 is found (924, yes). Subsequently, the stream engine 60 receives the node including the <td> tag 962 and another new candidate chunk 944 is found (924, yes). When the stream engine 60 receives the </p> tag 976, the first complete candidate chunk 956 is found (930, yes) and the stream engine 60 sends the candidate chunk 956 to the query engine 30 (932). Similarly, when the stream engine 60 reaches the </p> tag 980, the second complete candidate chunk 958 is found (930, yes) and sent to the query engine 30 (932). When the stream engine 60 reaches the </td> tag 982, the third complete candidate chunk 954 is found (930, yes) and sent to the query engine 30 (932). Note that the candidate chunk 954 is the parent of the two candidate chunks 956 and 958 and the candidate chunk 954 does not have any content outside the two descendant candidate chunks 956 and 958. As will be explained below, the query engine 30 identifies the candidate chunk 954 as the relevant chunk if the two descendant candidate chunks 956 and 958 in combination satisfy the user-specified search keywords.

[00145] Assume that the stream engine 60 has processed the last node in the UA node stream, which is one of multiple data packets occupied by a large paragraph in the corresponding candidate document, and the stream engine 60 has not received the last of the multiple data packets yet. In this case, because there are no more nodes in the UA node stream (922, no), the stream engine 60 returns to process the next URI (902). But as noted above, there are no more URIs available (904, no) because the stream engine 60 receives only

two URIs from the cache engine 40 and it has already generated a node stream for each URI. The stream engine 60 then has to choose between the UA node stream and the UB node stream (912).

[00146] If the stream engine 60 chooses one of the two node streams, e.g., the UA node stream, and for some reason the next data packet associated with the UA node stream has not arrived at the stream engine 60 after a certain time (918, no), the stream engine 60 then returns to perform operation 902. In some embodiments, the stream engine 60 does not randomly choose the next available node stream. Rather, it compares the available node streams and selects one node stream that has one or more data packets waiting to be processed (912). By doing so, the stream engine 60 effectively reduces the risk of running into the deadlock situation described above, which blocks the query engine 30 from identifying and serving relevant chunks from a different node stream.

[00147] For example, after finishing the last node in the UA node stream, the stream engine 60 may choose the UB node stream (912) and start searching for candidate chunks within the UB node stream until (i) the UB node stream is completed (914, no) or (ii) there is a network traffic jam with the UB node stream (924, no). In either case, the stream engine 60 repeats the same process described above to work on the UA node stream if it has newly received data packets and there is still time for processing the node stream for a given response latency threshold.

[00148] In some embodiments, as noted above, a feedback mechanism (936, Figure 9A) is established between the stream engine 60 and the query engine 30. The description above illustrates the activities on the stream engine side. The description below in connection with Figures 10 and 11 focuses on the query engine side, in particular, how the query engine 30 works in concert with the stream engine 60 to identify relevant chunks in response to a search request.

[00149] Figure 10A is a block diagram illustrative of how a query mediator coordinates the query engine 30 and the stream engine 60 to identify chunks within a node stream representing a candidate document in accordance with some embodiments.

[00150] As described above, upon receiving a search query, the query engine 30 may generate one or more path filters (308, Figure 3), the path filters are passed down to the stream engine 60 by the cache engine 40 (411, Figure 4), and the stream engine 60 then

applies the path filters to a node stream (512, Figure 5). Figure 10A depicts these processing steps in a slightly different manner.

[00151] Upon receiving the search query, the query engine 30 performs query processing 1010 to define a set of input sequences 1015 for the search query. The set of input sequences 1015 further defines one or more path filters, which are used to build a filter model 1020. In some embodiments, as described below in connection with Figure 10B, the filter model 1020 is the same or similar to a deterministic finite state machine (FSM).

[00152] In addition to defining the path filters, the query engine 30 iterates the input sequences 1015 and their associated node sub-streams to identify relevant chunks. Initially, because the query engine 30 has not received anything from the stream engine 60, a data request is submitted to the query mediator 1025. The query mediator 1025 is a user-configurable tool through which the user can, e.g., specify the maximum number of nodes in memory at any given time and control the rate of node stream consumption by the query engine 30.

[00153] In some embodiments, as the query engine 30 iterates each input sequence 1015 and its associated node sub-stream, it determines whether a desired node is currently in memory. If not, the query engine 30 asks the query mediator 1025 for the desired node until one of the predefined conditions is met. These conditions include: (i) another context node for the input sequence is available; (ii) another fragment or content node of the current context node is available; and (iii) the current context node is complete. A more detailed description of context nodes is provided below in connection with Figures 11A through 11G.

[00154] In response to the data request, the query mediator 1025 triggers the stream engine 60 for further conversion of raw data into the node stream 1030. As a result, more nodes are submitted to the filter model 1020. The filter model 1020 feeds these nodes into the finite state machine it builds using the path filters to accumulate those nodes matching the path filters in their respective sub-streams until one of the predefined conditions is satisfied. By then, the query mediator 1025 passes the control back to the input sequences 1015 and therefore the query engine 30, which analyzes the node sub-streams to identify relevant chunks.

[00155] In sum, this feedback mechanism between the stream engine 60 and the query engine 30 ensures that a minimum number of nodes are stored in the computer system 100's

memory and processed by the query engine 30 to fulfill the search query, and that this process is accomplished without loss of any raw data.

[00156] Figure 10B is a flowchart illustrative of how the stream engine 60 divides the node stream into multiple sub-streams using a filter model in accordance with some
5 embodiments.

[00157] Using the path filters provided by the query engine 30, the stream engine 60 generates a finite state machine (1034). The input to the finite state machine is a node stream corresponding to the raw data of a candidate document and the output is one or more node sub-streams, each node sub-stream including a set of nodes that may be potentially relevant
10 to the search query. Thus, the finite state machine effectively filters out nodes that are deemed to be completely irrelevant to the search query and reduces the amount of data to be handled by the query engine 30. Next, the stream engine 60 receives the next node in the node stream (1036) and compares the node with the finite state machine (1038) to determine if the node belongs to one or more node sub-streams associated with the path filters.

[00158] In some embodiments, the finite state machine performs different operations in accordance with different comparison results. For example, the finite state machine may:
(i) perform a transition operation (1040-A) and move itself from the current state to a different one that is associated with the node (1042); (ii) perform a maintenance operation (1040-B) and stay at the current state (1044); or (iii) perform a null operation (1040-C) and
20 discard the node as irrelevant to the search query (1046). In the last case, the finite state machine may also stay at the current state.

[00159] After performing a transition/maintenance operation, the stream engine 60 accumulates the node into a respective node sub-stream (1048). Depending on whether the node is a context node (1050-A) or a content node (1050-B), the stream engine 60 may insert
25 the node into the context node sub-stream (1052) or insert the node into a node sub-stream that is associated with the current context node (1054). A more detailed description of this accumulation operation is provided below in connection with Figure 11E. Next, the stream engine 60 determines whether the node stream is completed (1056). If so (1056, yes), the stream engine 60 sends the node sub-streams to the query engine 30 for further process
30 (1058). Otherwise (1056, no), the stream engine 60 returns to process the next node in the node stream (1036).

[00160] To further explain the feedback mechanism between the stream engine 60 and the query engine 30, Figures 11A through 11G illustrate in detail how a candidate document is processed using the feedback mechanism.

[00161] Figure 11A is an exemplary XML document 1100 to be processed by the stream engine 60 and the query engine 30 in accordance with some embodiments. The XML document 1100 includes a list of books 1102 through 1108, each book being identified by its publication year (the “year” attribute in the <book> tag), its title (the pair of <title> and </title> tags), its author (the pair of <author> and </author> tags) including first name (the pair of <first> and </first> tags) and last name (the pair of <last> and </last> tags), its publisher (the pair of <publisher> and </publisher> tags), and price (the pair of <price> and </price> tags).

[00162] Figure 11B is an exemplary XQuery 1110 to be applied to the XML document 1100 in accordance with some embodiments. The XQuery 1110 searches for any book in the XML document “bib.xml” whose publisher is Addison-Wesley and whose publication year is later than 1991. The XQuery 1110 requires that the search results be returned in the form of a new XML document including a new list of the books matching the two search criteria, each book in the new XML document only including the book’s title and its publication year as an attribute in the <book> tag.

[00163] Figure 11C is a table 1115 of the five input sequences defined by the query engine 30 in accordance with some embodiments. Based on the XQuery 1110, the query engine 30 defines five input sequences, each input sequence corresponding to one tag or tag attribute within the XML document 1100. Note that the publication year attribute “@year” appears twice in the XQuery 1110, one in the where clause and the other in the return clause, and corresponds to two separate input sequences. The five input sequences each have an associated node sub-stream labeled “Node Sub-Stream (0)” through “Node Sub-Stream (4)” and correspond to a respective path filter as shown in the table 1115.

[00164] Different input sequences are associated with different portions of the XQuery 1110 and therefore have different modes. For example, the <book> tag associated with the input sequence “Node Sub-Stream (0)” appears in the for-in clause, but not the return clause. Thus, the input sequence “Node Sub-Stream (0)” is presumed to provide context for the

search process and serve in the “Context” mode, and the nodes in the corresponding node sub-stream are referred to as “context node sub-stream.”

[00165] Similarly, the content of the <publisher> tag associated with the input sequence “Node Sub-Stream (1)” is compared with “Addison-Wesley” in the where clause of the XQuery 1110. Thus, the input sequence “Node Sub-Stream (1)” is presumed to provide content for the search process and serve in the “Content” mode, and the nodes in the corresponding node sub-stream are therefore referred to as “content node sub-stream.” The <title> tag associated with the input sequence “Node Sub-Stream (4)” appears in the return clause. Thus, the input sequence “Node Sub-Stream (4)” is presumed to provide both context and content for the search process and serve in the “All” mode. In some embodiments, an input sequence in the “All” mode has two node sub-streams.

[00166] The “Parent” column in the table 1115 indicates whether an input sequence is a child of another input sequence. In this example, the input sequence associated with the for-in clause provides basis for the other input sequences associated with the other parts of the XQuery 1110. Any node in one of the other four input sequences corresponds to a specific node in the input sequence “Node Sub-Stream (0),” which is therefore deemed to be the parent input sequence of the other four input sequences.

[00167] Figure 11D is a flowchart illustrative of how the query engine 30 processes node sub-streams at different input sequences in accordance with some embodiments. This flowchart provides more details of the information flow shown in the block diagram of Figure 10A.

[00168] The query engine 30 initializes the stream engine 60 (1120) and processes the search query (1122) to define input sequences, path filters, and a finite state machine that is used for generating one or more node sub-streams. The query engine 30 then starts iterating the next node sub-stream (1124). In this example, the query engine 30 begins with the context node sub-stream of Node Sub-Stream (0).

[00169] If the context node sub-stream has no context node (1126, no), the query engine 30 then requests more context nodes from the stream engine 60 (1128, 1130). Consequently, more data packets are retrieved (1132) and parsed (1134) to provide more nodes, including context nodes and content nodes, to the query engine 30.

[00170] Once a new context node is present in the node sub-stream of Node Sub-Stream (0) (1126, yes), the query engine 30 applies the search query to the context node and its associated nodes in other node sub-streams (1136). If the search criteria are satisfied (1138, yes), a relevant chunk has been identified and there is no need to apply the search query to the remaining portion of the relevant chunk. Thus, the query engine 30 needs to quickly reach the end of the relevant chunk through round-tripping the content nodes in different node streams (1140). After finishing the content nodes, if the end of the chunk has not been reached (1142, no), the query engine 30 may request the stream engine 60 to process more data packets (1146).

[00171] If the search criteria are not met (1138, no), a relevant chunk has not been identified, and the query engine 30 sends a request to the query mediator to retrieve one or more content nodes and re-apply the search query. If the stream engine 60 has more nodes or node fragments (1144, yes), they will be parsed and submitted to the query engine 30 (1134). Otherwise (1144, no), the query engine 30 may request the stream engine 60 to process more data packets (1146).

[00172] As shown in Figure 11C, the XQuery 1110 defines five input sequences and therefore five path filters. The stream engine 60 uses these path filters to build a finite state machine, which, as shown in Figure 10B, is to divide the original node stream corresponding to the XML document 1100 into five node sub-streams. The finite state machine has an initial state, which can be any one of the five input sequences.

[00173] Figure 11E is a block diagram illustrative of how the node stream is divided into multiple node sub-streams by the finite state machine in accordance with some embodiments. From the start state (1150), the finite state machine receives a node including the <bib> tag. Because this tag is not relevant to any input sequence, the finite state machine discards the node. After receiving a node including the <book> tag, the finite state machine makes a transition to the state corresponding to Node Sub-Stream (0) and adds the node into the corresponding context node stream (1152). Next, the node including the publication year attribute is processed and added into the two respective node sub-streams corresponding to Node Sub-Stream (2) and Node Sub-Stream (3) (1154).

[00174] Upon receiving a node including the <title> tag, the finite state machine makes another transition to the state corresponding to Node Sub-Stream (4). As noted above in

connection with Figure 11C, the input sequence Node Sub-Stream (4) serves in the “All” mode. Thus, besides adding the node including the <title> tag into the corresponding node sub-stream (1156), the finite state machine adds everything enclosed by the pair of (<title>, </title>) tags into the same node sub-stream or a separate node sub-stream (1158). For example, if there is a pair of (<subtitle>, </subtitle>) tags within the pair of (<title>, </title>) tags, they will be added into the respective node sub-stream because, as explained above, the XQuery 1110 requires the return of each matching book’s title, including its subtitle, if present.

[00175] Similarly, the node including the <publisher> tag is added into the node sub-stream corresponding to Node Sub-Stream (1) (1160) and the textual portion within the pair of (<publisher>, </publisher>) tags is extracted by a text() function and added into the same or a separate node sub-stream (1162). This textual portion is required by the XQuery 1110 to check whether the book is published by the publisher Addison-Wesley.

[00176] Figure 11F is a block diagram illustrative of the input sequences and their associated node sub-streams after the first candidate chunk in the XML document is processed in accordance with some embodiments.

[00177] The node sub-stream “Node Sub-Stream (0)” is first populated with a context node “<book>” (1164). Next, the node sub-streams “Node Sub-Stream (2)” and “Node Sub-Stream (3)” are each populated with a content node “1994” (1166, 1168). For the node including the <title> tag, the stream engine 60 inserts into the node sub-stream “Node Sub-Stream (4)” both the <title> tag (1170) and the data descending from the <title> tag (1172). For the node including the <publisher> tag, the stream engine 60 is only interested in its content and therefore populates the node sub-stream “Node Sub-Stream (1)” with the textual portion of the <publisher> tag (1174).

[00178] Figure 11G is the search result of applying the XQuery 1110 to the node sub-streams derived from XML document 1100 in accordance with some embodiments. The search result is also an XML document 1180 that includes two books 1182 and 1184 that satisfy the XQuery 1110. As shown in Figure 11F, the node sub-streams corresponding to the five input sequences include all information necessary for generating this resulting XML document 1180.

[00179] Thus far, detailed descriptions of document-processing schemes in response to a search request, including the downstream processes 25 and the upstream processes 35, are provided above. These document-processing schemes can be used to implement various applications to satisfy different user needs. For illustration, embodiments of representative applications are provided below.

[00180] One application of the invention is to improve a user's experience with the search results generated by search engines. Although a document identified by the search results is relevant to the search keywords, the document may not include any relevant chunks because the search engines treat the entire document, not a chunk within the document, as the basic unit to be compared with the search keywords. Thus, one aspect of the invention is to identify and display relevant chunks within a document in response to a search request.

[00181] Figure 12A is a flowchart illustrative of a first process of identifying one or more documents, each document having one or more chunks that satisfy user-specified search keywords, in accordance with some embodiments.

[00182] A computer identifies multiple resource identifiers (1201), each resource identifier corresponding to a document at a respective data source. In some embodiments, a resource identifier is a URL, which identifies a web page at a remote web server. In some embodiments, the resource identifiers are part of search results produced by a server computer such as a search engine in response to one or more search keywords provided by an end user from a client computer.

[00183] For at least one of the resource identifiers, the computer retrieves the corresponding document from the respective document source (1203). If the document is a web page hosted by a web server, the computer submits an HTTP request to the web server and the web server returns the document in an HTTP response. Within the retrieved document, the computer identifies a chunk that satisfies the user-specified search keywords (1205) and displays the identified chunk and a link to the identified chunk within the document to the user (1207).

[00184] Figure 12B is a flowchart illustrative of a second process of identifying one or more document, each document having one or more chunks that satisfy user-specified search keywords, in accordance with some embodiments.

[00185] A client computer submits one or more user-specified search keywords to a server computer (1211). In some embodiments, the server computer is one or more third-party search engines. The client computer receives a set of search results from the server computer (1213), each search result identifying a document located at a respective document source that satisfies the search keywords in accordance with a first set of predefined criteria.

[00186] For each identified document, the client computer retrieves the document from the corresponding document source (1215), identifies a chunk within the document that satisfies the search query in accordance with a second set of predefined criteria (1217), and displays the identified chunk and a link to the identified chunk within the document (1219).

In some embodiments, the two sets of predefined criteria are different. For example, the first set of criteria requires that all the search keywords be found within a document, but not necessarily within a chunk. In contrast, the second set of criteria is satisfied only if all the search keywords are found within a chunk.

[00187] Figures 12C through 12J are screenshots of a graphical user interface on a computer display illustrative of features associated with the processes as shown in Figures 12A and 12B in accordance with some embodiments.

[00188] The graphical user interface includes one or more document links, each document link having one or more associated chunks identified within the corresponding document as satisfying one or more user-specified search keywords. In some embodiments, each chunk has an associated chunk link and includes terms matching each of the user-specified search keywords. The matching terms may also be highlighted in the chunk in a visually distinguishable manner (such as in different colors, font types or combination thereof). In response to a user selection of a chunk's chunk link, the corresponding document is displayed in a window on the computer display and at least a portion of the chunk is highlighted in the window.

[00189] In some embodiments, each document link has an associated chunk page-link icon for searching chunks within documents that are referenced by the corresponding document. In response to a user selection of a document link's associated chunk page-link icon, one or more referenced document links are displayed on the computer display, with each referenced document link having one or more associated chunks identified within a corresponding referenced document as satisfying the user-specified search keywords.

[00190] In some embodiments, each document link has an associated hide-chunk icon. In response to a user selection of a document link's associated hide-chunk icon, the chunks associated with the document link and their associated chunk links disappear from the computer display.

5 [00191] In some embodiments, chunks associated with a respective document link are displayed in an order consistent with their relative relevancy to the user-specified search keywords. In some other embodiments, chunks associated with a respective document link are displayed in an order consistent with their relative locations within the corresponding document.

10 [00192] As shown in Figure 12C, through an application 1220 (e.g., a web browser window), a user submits three search keywords 1221 from a client computer to a content provider such as a search engine. In this example, the application 1220 provides four different search options for the user to choose. They are:

- 15 • “Best Match” option 1226-A – This search option allows the application 1220 to adaptively select one or more chunks satisfying one or more of the user-specified search keywords according to predefined criteria. In some embodiments, the “Best Match” option is the default option if the user does not expressly choose a different one. A more detailed description of this search option is provided below in connection with Figures 21A through 21D.
- 20 • “Match All” option 1226-B – This search option limits the search results to relevant chunks that satisfy each of three user-specified search keywords. Thus, a candidate chunk that only includes “einstein” and “bohr,” but not “debate,” should not be in the search results responsive to the “Match All” option, but may be in the search results responsive to the “Best Match” option.
- 25 As shown in Figure 12C, the user expressly chooses the “Match All” option.
- “‘Exact’ Match” option 1226-C – This search option further limits the search results to relevant chunks that not only satisfy each of three user-specified search keywords, but also include an exact match of the search keywords as a string. Examples of “exact”-matching chunks are shown in Figures 12F and
- 30 12G, respectively. Note that this option is different from the string-match approach, which is variant-sensitive. Under the string-match approach,

“einstein bohr debates” does not match “Einstein-Bohr debate.” But according to the “‘Exact’ Match” option, the two sets of terms do match each other as this search option ignores any non-word characters such as white space, punctuation, etc., and only requires that the three terms appear in the same order and have no intervening terms.

- “Match Any” option 1226-D – This search option allows the application 1220 to identify and display any chunk that satisfies at least one of the user-specified search keywords. Thus, the search results responsive to any of the three options above are a subset of the search results responsive to the “Match Any” option, an example of which is depicted in Figures 12I and 12J.

[00193] The content provider returns a search result 1225 to the client computer and the search result 1225 includes an abbreviated document segment identified by the search engine as satisfying the search keywords. The client computer retrieves a document identified in the search result 1225 (an HTML web page in this example) from a web server and identifies one or more chunks 1229-A through 1229-C therein that satisfy the search keywords 1221, each chunk having an associated link 1231 to the chunk in the original web page. In some embodiments, each of the chunks 1229-A through 1229-C are different from the abbreviated document segment because it is a semantically and contextually consistent unit within the document without abbreviation.

[00194] In some embodiments, after retrieving a candidate document, the application 1220 generates a search query using the search keywords and applies the search query to the retrieved document to identify relevant chunks within the document.

[00195] In some embodiments, the terms that match the search keywords in the identified chunk are ordered differently from the user-specified search keywords. For example, the term “debate” appears between “Bohr” and “Einstein” in the chunk 1229-B of Figure 12C.

[00196] In some embodiments, the terms that match the search keywords in the identified chunk are separated from one another by at least one term not matching any of the search keywords. For example, the three terms appearing in the last sentence of the chunk 1229-A are separated from one another by many other words. Unlike the conventional string search, an identified chunk may or may not include an exact match of the search keywords as

a string. Rather, the search process according to some embodiments of the invention includes tokenization of the search keywords in a text string into atoms and subsequent search in the token space according to the atoms, which is variant-agnostic by, e.g., ignoring things like grammatical tense, punctuation, white space, casing, diacritics, etc. in the search keywords.

5 For example, in the screenshot of Figure 12C, “Einstein Bohr debate” and “einstein bohr debating” are deemed to be identical according to some embodiments of the invention.

[00197] In some embodiments, an identified chunk includes an identical instance of the search keywords appearing as a phrase. But, as noted above, although the instance is the same as the result of a string search, the search keywords are not identified collectively as a
10 text string within the chunk.

[00198] In some embodiments, different terms matching different search keywords in the identified chunk are highlighted in different manners such as different colors, different foreground/background patterns, different font types/sizes, or a combination thereof. In Figure 12C, the three terms appearing in each chunk are highlighted using large, italic, and
15 underlined font. In some embodiments, the three terms are further distinguished from one another using a unique style for each term. For example, the three terms may have three completely different styles such as Courier New for “**Einstein**,” Arial for “**Bohr**,” and Monotype Corsiva for “*debate*.” In some other embodiments, the three terms may have different background colors, such as gray for “Einstein,” green for “Bohr,” and yellow for
20 “debate.” In yet some other embodiments, the different manners may be combined to further distinguish different search keywords appearing in the same chunk.

[00199] In some embodiments, one or more sponsored links (1227, Figure 12C) are identified to be associated with at least one of the search keywords and displayed adjacent the identified chunk.

25 [00200] As shown in Figure 12C, there are a chunk page-link icon 1223 and a hide-chunk icon 1224 below the search result 1225. In response to a user selection of the chunk page-link icon 1223, the computer retrieves documents that are referenced by the document identified by the search result 1225 and therefore have page links in the document. For each retrieved document, the computer identifies chunks within the document that satisfy the
30 search keywords 1221 by apply the same “chunking” process that has been applied to the document identified by the search result 1225.

[00201] Figure 12D is a screenshot illustrative of the search results after a user selection of the chunk page-link icon 1251, including a link 1253-A, 1253-B to a respective document and a set of relevant chunks 1255-A, 1255-B identified within the corresponding document. The terms that match the search keywords are similarly highlighted in the relevant chunks. Note that a user can repeat this process by clicking the chunk page-link icons 1254-A, 1254-B associated with the respective documents. In some embodiments, the application 1220 applies its default search option, e.g., “Best Match” option 1226-A, for performing the task associated with the chunk page-link icon 1251. In some other embodiments, the user can override the default search option by expressly selecting another option.

[00202] Figure 12E is a screenshot illustrative of the search results after the user clicks the hide-chunk icons (1224, Figure 12C) associated with the respective search results. In this example, the relevant chunks associated with a search result disappear from the web browser window and the hide-chunk icons become show-chunk icons 1257A, 1257B. The relevant chunks are displayed again in the web browser window after a user selection of the show-chunk icons.

[00203] In some embodiments, multiple relevant chunks are identified within a candidate document and these chunks are displayed in an order consistent with their relative locations within the document. Figure 12F is a screenshot that includes multiple relevant chunks, each one satisfying the two search keywords “Bohr-Einstein” and “debates.” These chunks are listed in the web browser window in accordance with their relative locations in the web page such that the first chunk 1233-A that appears first in the web page is displayed above the other ones and the last chunk 1233-B that appears below the other chunks is displayed at the bottom of the web browser window.

[00204] In some embodiments, multiple relevant chunks are identified within a candidate document and these chunks are displayed in an order consistent with their relative relevancy to the search keywords. Figure 12G is another screenshot that includes the same set of relevant chunks. Assume that the chunk 1233-B is more relevant than the chunk 1233-A. The more relevant chunk 1233-B is displayed above the other less relevant chunks including the chunk 1233-A. For illustrative purposes, the two screenshots in Figures 12F and 12G are generated using the “‘Exact’ Match” option 1226-C. Each chunk 1233-A, 1233-B includes at least one instance of the two search keywords as a string (ignoring the casing

difference). The aforementioned chunk-ordering schemes or the like are equally applicable to the other search options.

[00205] In some embodiments, in response to a user selection of the link to an identified chunk, at least a portion of the identified document is displayed in a document view window and the displayed portion includes, at least partially, the identified chunk. Figure 12H is a screenshot of the web browser window after a user click of the chunk link 1235. A document view window 1237 is displayed next to the search results. The document view window 1237 displays a portion of the document that includes the relevant chunk and the displayed portion includes at least part of the relevant chunk 1239 within the document. In this example, the relevant chunk 1239 is highlighted in the document view window. Sometimes, the terms matching the search keywords in the relevant chunk 1239 are processed such that they are visually distinguishable over the rest of the identified chunk, e.g., using different colors or font types.

[00206] In some embodiments, for each relevant chunk in the identified document, the computer inserts a pair of unique chunk virtual delimiters into the identified document. This pair of chunk virtual delimiters uniquely defines the scope of the relevant chunk within the identified document, but is invisible to the user when the identified document is being rendered by an application. In response to a user request to view the relevant chunk 1239 in the document view window 1237, the computer can quickly locate the scope of the relevant chunk 1239 within the document by looking for the corresponding pair of chunk virtual delimiters and then highlight the chunk in the document view window appropriately.

[00207] In some embodiments, the HTML tag `<span>` can be introduced into a candidate document for forming chunk virtual delimiters. For example, the following chunk in an HTML document

`<p>This is a candidate chunk.</p>`

can be re-defined as:

`<span id="chunk-1"><p>This is a candidate chunk.</p></span>`

[00208] The HTML tag `<span>` has no effect on the appearance of the chunk in a web browser window because it has no associated style information. But the pair of chunk virtual delimiters (`<span id="chunk-1">`, `</span>`) uniquely identifies the chunk's location in the

document, which a web browser application can rely upon to highlight the chunk's existence by, e.g., altering its background color. Note that the HTML tag is not the only choice of a suitable invisible anchor element. In some other embodiments, it is possible to use one or more document-unique, chunk-unique identifiers or the like within the document as chunk virtual delimiters to achieve the same or similar effect.

[00209] In some embodiments, for at least one of the resource identifiers, after the corresponding document is retrieved from the respective document source, no relevant chunk that satisfies each of the search keywords is identified therein. This scenario happens if the terms matching the search keywords are distributed in different chunks within the document. In this case, the web browser window displays a link to search for chunks that satisfy any of the search keywords within the document. In response to a user selection of the link to search for chunks that satisfy any of the search keywords within the document, the retrieved document is re-processed, and as a result, one or more chunks that satisfy at least one of the search keywords is identified in the document. Accordingly, these chunks are displayed to the end user.

[00210] Figure 12I is a screenshot that includes a search result 1241 that satisfies all the search keywords "Einstein" and "big bang." Because no relevant chunk is found in the web page, the web browser window provides a link 1243 to "re-chunk" the web page to search for any chunk matching any search keywords. Figure 12J is another screenshot after the user click of the link 1243. Note that at least five chunks are identified in the document, three chunks 1245 including the keyword "Einstein" and two other chunks 1247 including the keywords "big bang." But no chunk satisfies all the search keywords. In some embodiments, the same set of chunks can be identified in the document through a user selection of the "Match Any" option 1226-D.

[00211] Another application of the invention is to identify and display within a document relevant chunks satisfying user-specified search keywords while the user is browsing the document. Conventionally, a user visiting a web page may be only interested in the content of a particular paragraph therein. To find the paragraph, the user-specified text string has to exactly match the one in the paragraph. Otherwise, the paragraph can not be easily located in the document if the user can provide a few search keywords but is mistaken about their exact sequence in the paragraph. Such issues with the conventional approach have been solved by the application described below.

[00212] Figure 13A is a flowchart illustrative of a first process of identifying within a document one or more chunks that satisfy user-specified search keywords in accordance with some embodiments.

[00213] A computer displays a portion of a document to a user (1302). Upon receiving
5 a user-specified text string that includes multiple search keywords, the computer identifies a chunk within the document that satisfies the search keywords (1304) and displays the identified chunk to the user (1306). In some embodiments, the identified chunk is not within the initially displayed portion of the document. To locate the chunk, the computer generates a search query using the search keywords and applies the search query to the document to
10 identify the chunk. In some embodiments, the terms that match the search keywords are either ordered differently from the search keywords in the user-specified text string or separated from one another by at least one term not matching any of the search keywords.

[00214] Figure 13B is a flowchart illustrative of a second process of identifying within
15 a document one or more chunks that satisfy user-specified search keywords in accordance with some embodiments.

[00215] While a user is browsing a document through a computer, the computer receives multiple user-specified search keywords (1312). The search keywords have a first sequence. Within the document, the computer identifies at least one chunk that satisfies the search keywords (1314) and displays a portion of the document including the identified
20 chunk (1316). In some embodiments, the search keywords are highlighted in the identified chunk and have a second sequence that is different from the first sequence.

[00216] Figures 13C through 13G are screenshots of a graphical user interface on a computer display illustrative of features associated with the second process as shown in Figures 13A and 13B in accordance with some embodiments.

[00217] Figure 13C is a screenshot of a web page covering Bohr-Einstein debates at
25 www.wikipedia.org. Assuming that a visitor of this web page is interested in learning about the experimental apparatus developed by George Gamow, the visitor can enter a few search keywords relating to this topic in the input field 1322 and then click the “Chunk Page” icon 1323.

[00218] Figure 13D is a screenshot of the web page including the identified chunk 1326 that satisfies the user-specified search keywords 1324, i.e., “gamow” and “experiment.” In this example, the relevant chunk 1326 is actually not a paragraph, but a caption of a figure in the document. The sentence 1327 including the two keywords is read as follows: “George
5 *Gamow*’s make-believe experimental apparatus for validating the thought *experiment* ...” Although the two keywords are separated from each other by other terms, the figure caption is identified nonetheless because the two keywords happen to be within the same chunk.

[00219] Figure 13E is a screenshot illustrative of another embodiment of the invention in response to a user selection of the “Chunk Page” icon at the top of the web browser
10 window. In this example, the left side of the web browser window displays the relevant chunks 1325 identified within the web page. If the web page has multiple relevant chunks, the user can easily get an overview of these chunks from the left side of the web browser. The right side of the web browser is a document view window that displays the portion of the document including the relevant chunk 1326. Thus, this document view window provides
15 more contexts for each relevant chunk to the user.

[00220] In some embodiments, like the examples described above in connection with Figures 12C through 12J, different terms in the identified chunk that match different search keywords are highlighted in different manners such as different colors, different foreground/background styles, different font types, or a combination thereof.

20 [00221] In some embodiments, multiple relevant chunks are identified within a document, each one appearing at a respective location in the document. In this case, the web browser window displays, at least partially, the chunk that appears above the other chunks in the document and its associated context.

[00222] Figure 13F is a screenshot of another web page at www.wikipedia.org. In
25 response to the user-specified search keywords 1328 “cosmic,” “background,” and “radiation,” the first relevant chunk 1330 in the web page that matches the three search keywords is identified and displayed in a visually distinguishing manner. A scroll down of the web page displays additional relevant chunks identified in the web page.

[00223] Sometimes, the first relevant chunk shown in Figure 13F is not necessarily the
30 most relevant one. In some embodiments, after identifying multiple chunks within the document, the web browser assigns to each chunk a ranking metric indicative of its relevancy

to the search keywords and displays in a prominent location, at least partially, the chunk that has the highest ranking metric.

[00224] Figure 13G is a screenshot of the same web page shown in Figure 13F. But the relevant chunks are now displayed in an order consistent with their relevancy to the search keywords. In this case, the relevant chunk 1332 is a section heading, which is presumably more relevant than the chunk 1330 shown in Figure 13F.

[00225] In some embodiments, if there is no chunk within the document that satisfies each of the search keywords, the web browser, or more specifically, the “Chunk Page” toolbar application, may relax its search criteria to look for any chunks in the document that satisfy any of the search keywords and display them to the user. In other words, this feature is similar to the one described above in connection with Figures 12I and 12J.

[00226] Another application of the invention is to identify relevant chunks within unstructured or semi-structured documents. It has been a particular challenge to identify chunks within an HTML web page because the HTML syntax allows its user to produce the same or similar web page layout using very different metadata.

[00227] Figure 14 is a flowchart illustrative of a process of modeling a document and identifying within the document one or more chunks that satisfy user-specified search keywords in accordance with some embodiments.

[00228] A computer identifies a document in response to a search request from a user (1401). The document includes content data and metadata, and the search request includes one or more search keywords. In some embodiments, the document is a semi-structured document, e.g., an HTML web page. The content data refers to the document’s content such as a paragraph, a table, or a list of bullet items, etc. The metadata specifies how the content data should be rendered through an application, e.g., a web browser window.

[00229] The computer generates a hierarchical semantic model of the content data of the document by applying heuristics to the metadata of the document (1403). In some embodiments, the generation of the hierarchical semantic model includes identifying one or more candidate chunks in the document, each candidate chunk corresponding to a respective subset of the document. As noted above, the HTML web page shown in Figure 9B has a hierarchical semantic model, which includes a set of HTML tags at different levels.

[00230] In some embodiments, a first subset of the document associated with a first candidate chunk encompasses a second subset of the document associated with a second candidate chunk. For example, as shown in Figure 9B, both the candidate chunks 956 and 958 are within the candidate chunk 954, which is, in turn, within the candidate chunk 952.

5 There is no overlapping between the candidate chunk 956 and the candidate chunk 958.

[00231] In some embodiments, the heuristics stipulates that a subset of the document is identified as a candidate chunk if the subset of the document has at least one instance of predefined metadata. For example, the candidate chunks 956 and 958 are identified because each begins with the <p> paragraph tag.

10 [00232] In some embodiments, the heuristics stipulates that a subset of the document is deemed to be a candidate chunk if the subset of the document has at least two instances of predefined metadata. For example, two or more instances of the tag appearing in a web page one after another are collectively identified as a candidate chunk.

[00233] The computer identifies a chunk within the document by sequentially scanning
15 the hierarchical semantic model (1405). The identified chunk includes a subset of the content data that satisfies the search keywords and the corresponding metadata. The computer returns the identified chunk to the requesting user (1407).

[00234] In some embodiments, assume that there are two search keywords, a first search keyword and a second search keyword. While sequentially scanning the semantic
20 model, the computer first identifies some content data that is in the first candidate chunk and precedes the second candidate chunk as satisfying the first search keyword (e.g., "It's **raining** outside ...") and then identifies content data in the second candidate chunk that satisfies the second search keyword (e.g., "For XML-based **data** management"). Because both search keywords are matched, the first candidate chunk is chosen to be the identified chunk and
25 returned to the requesting client.

[00235] In some embodiments, the computer does not return the first chunk immediately after finding a match for the search keyword. Rather, the computer continues scanning the model until identifying content data in the second candidate chunk that also satisfies the first search keyword (e.g., "Raining Data is your choice"). In this case, the
30 second candidate chunk is returned as the relevant chunk that is more specific than the first one.

[00236] In some embodiments, while sequentially scanning the hierarchical semantic model, the computer identifies content data that satisfies the first search keyword in one candidate chunk and content data that satisfies the second search keyword in another candidate chunk. For example, assume that the search keywords are “CAD” and “job listings.” As shown in Figure 9B, the candidate chunk 956 includes the search keyword “CAD” and the candidate chunk 958 includes the search keyword “job listings.” In this case, the computer chooses the candidate chunk 954, which is the parent of the chunks 956 and 958 in the hierarchical semantic mode, as the identified chunk. Note that there is no other content data or metadata within the candidate chunk 954 besides the two candidate chunks 956 and 958.

[00237] Another application of the invention is to transform the user-specified search keywords into a finely-tuned query. Sometimes, the user-specified search keywords may include a special character (e.g., “%”) or sequence of characters (e.g., “Jan 22 2008”). This special character or sequence of characters, if interpreted appropriately, can help to find the relevant chunks more efficiently.

[00238] Figure 15 is a flowchart illustrative of a process of customizing a search query based on user-specified search keywords in accordance with some embodiments.

[00239] After receiving a search keyword provided by a user (1502), the computer selects an archetype for the search keyword (1504). The computer identifies one or more search results in accordance with the archetype (1506) and returns at least one of the search results to the user (1508).

[00240] In some embodiments, the archetype has an enumerable set of instances and the search keyword is associated with one of the instances. For example, if the user-specified search keyword is “Tuesday,” a possible archetype would be “week,” of which “Tuesday” represents one of the seven members in the archetype.

[00241] In some embodiments, after selecting the archetype, the computer identifies at least one query operator for the selected archetype, constructs a search query using the query operator, and then executes the search query against one or more data sources. For example, for the “week” archetype, the computer may generate a search query that looks for chunks including not only the keyword “Tuesday,” but any of the seven days within a week such as “Sunday,” “Monday,” etc.

[00242] In some embodiments, the query operator has a scope and the search query is constructed to limit search results within the scope. For example, assume that the search phrase is “discount of 10%.” It is likely that the user is not only interested in chunks having the phrase “discount of 10%,” but also chunks having similar phrases, e.g., “discount of 15%.” Alternatively, the user may be mistaken when entering the phrase and the candidate document actually has no chunk including the phrase “discount of 10%,” but does have chunks including the phrase “discount of 20%.” In this case, the computer may generate a search query for discount within the scope of 0% to 100%. As a result, more relevant chunks are identified.

[00243] In some embodiments, the query operator has a pattern and the search query is constructed to limit search results including the pattern. For example, the user-specified phrase “Jan 22 2008” indicates a date pattern. If so, the computer may generate a search query accordingly to search for any chunk having the date pattern.

[00244] In some embodiments, after selecting the archetype and before identifying the search results, the computer solicits user instructions in connection with the archetype, constructs the search query in accordance with the user instructions, and executes the search query against the data sources. For example, if the user-specified search keyword includes the special character “%,” the computer may display a user interface through which the user may specify the scope or range associated with that special character, which is then built into the search query.

[00245] In some embodiments, based on the user instructions, the computer may generate feedback to the user instructions and then receive more user instructions in connection with the archetype and the feedback. Note that this process may repeat for multiple loops until the user submits a search query execution request, which suggests that the user is satisfied with the customized search query.

[00246] Another application of the invention is not only to display relevant chunks identified within a document but also to re-use them for different purposes. For example, when a user composes a Word document using Microsoft Office, the user may like to view a slide in a PowerPoint document and, if needed, generate a copy of the slide in the Word document. Currently, there is no convenient way to do so other than opening the PowerPoint

document in a separate window, manually searching for the slide in the window, and manually copying the slide and pasting it into the Word document.

[00247] Figure 16A is a flowchart illustrative of a process of displaying and re-using search results based on user instructions in accordance with some embodiments.

5 [00248] A computer displays an application user interface (1601). The application user interface includes a document authoring window and a search results window. In response to a search request including one or more user-specified search keywords, the computer displays in the search results window a set of search results in a text-only display format (1603). In some embodiments, each search result includes a chunk within a respective
10 document that satisfies the search keywords. In response to a user request to view a chunk, the computer launches a document display window in the application user interface and displays therein a portion of the corresponding document that includes the chunk in its native display format (1605). In response to a user request to duplicate a segment of the
15 corresponding document in the document authoring window, the computer generates therein an instance of the segment of the corresponding document in its native display format (1607).

[00249] Figures 16B through 16J are screenshots of a graphical user interface on a computer display illustrative of features associated with the process as shown in Figure 16A in accordance with some embodiments.

[00250] The application user interface includes a document authoring window and a
20 search results window. A set of search results associated with one or more user-specified search keywords is displayed in the search results window in a text-only display format and each search result includes one or more chunks identified within a respective document as satisfying the user-specified search keywords. In response to a user request to duplicate a chunk within a document in the document authoring window, an instance of the chunk is
25 displayed in the document authoring window in the document's native display format. In some embodiments, two chunks identified within two different documents have different native display formats.

[00251] In some embodiments, each chunk in the search results window has an associated chunk link. In response to a user selection of a respective chunk link, a document
30 display window is displayed in the application user interface and a portion of the

corresponding document that includes the corresponding chunk is displayed in the document display window in the document's native display format.

[00252] In some embodiments, each chunk includes terms that match the user-specified search keywords an associated chunk link. Different terms matching different search keywords are highlighted in the search results window in a visually distinguishable manner.

[00253] In some embodiments, the chunks identified within a document are displayed in the search results window in an order consistent with their relative relevancy to the user-specified search keywords. In some other embodiments, the chunks identified within a document are displayed in the search results window in an order consistent with their relative locations within the corresponding document.

[00254] Figure 16B is a screenshot of the Microsoft Office 2007 Word application user interface 1611. The main region of the user interface is occupied by a document authoring window 1613. Above the document authoring window 1613 is an add-in 1615 to Microsoft Office 2007. The add-in 1615 includes a keyword(s) input field 1615-A into which the user enters search keywords, a document type selection field 1615-B through which the user selects the types of candidate documents to be searched, and a web source field 1615-C including multiple document sources through which the user can search and re-use documents identified by the respective document sources.

[00255] In some embodiments, the set of search results includes a first chunk within a first document having a first content type and a second chunk within a second document having a second content type, wherein the first content type is different from the second content type. Different search keywords in the search results window are highlighted in different manners.

[00256] Figure 16C is a screenshot including a search results window 1625 and the search phrases 1621 "Einstein general relativity." In this example, the user limits the document search to two types of documents 1623, Word and PowerPoint. As described above in connection with Figure 1, this search limit is passed down from the front end 15 (the add-in 1615 in this example) to the query engine 30 and then to the cache engine 40. Thus, the cache engine 40 only looks for Word and PowerPoint documents in the index database

50. In this example, one chunk 1627 from a PowerPoint document and another chunk 1629 from a Word document are shown in the search results window 1625.

[00257] Note that each chunk in the search results window has an associated content type, which may be different from the document type of the corresponding document that includes the chunk. For example, a Word document may include a PowerPoint slide or an Excel spreadsheet. If the PowerPoint slide is identified to be the relevant chunk, the content type of the relevant chunk is PowerPoint, not Word, although the PowerPoint slide is within a Word document. Similarly, if a row in the Excel spreadsheet is identified to be the relevant chunk and the content type of the relevant chunk is therefore Excel, not Word. These chunks may or may not be displayed depending upon the embodiment.

[00258] In some embodiments, in response to a user request to duplicate the first chunk from the search results window into the document authoring window, the computer generates therein an instance of a first segment of the first document, including the first chunk, in its native display format. In response to a user request to duplicate the second chunk from the search results window into the document authoring window, the computer generates therein an instance of a second segment of the second document, including the second chunk, in its native display format. Sometimes, the first document and the second document have different native display formats.

[00259] Figure 16D is a screenshot including a PowerPoint slide 1633 in the document authoring window and the slide 1633 corresponds to the relevant chunk 1627 in Figure 16C. To duplicate this slide 1633 in the document authoring window, the user first selects the checkbox 1631 next to the text-only version of the slide in the search results window and then clicks the duplicate icon 1635 at the top of the search results window.

[00260] Figure 16E is another screenshot including not only the PowerPoint slide 1633 but also a paragraph 1643, which corresponds to the relevant chunk 1629 in Figure 16C. To duplicate this paragraph 1643 in the document authoring window, the user first selects the checkbox 1641 next to the text-only version of the paragraph in the search results window and then clicks the duplicate icon 1645 at the top of the search results window.

[00261] Note that a PowerPoint document and a Word document are deemed to have different native display formats. But relevant chunks in the search results window are displayed in a text-only format regardless of whether these chunks are identified within a

PowerPoint document, a Word document, a plain-text document or even a PDF document. But when a chunk is duplicated into the document authoring window, the computer tries to display the chunk in its native format. Note that a chunk found in a plain-text or PDF document will be customized to a native display format associated with the document authoring window. In other words, if the document authoring window is a Word document authoring window, the chunk is displayed in the Word document's native display format.

[00262] In some embodiments, the user may like to display a relevant chunk in its native display format before re-producing the chunk in the document authoring window. For example, in response to a first user selection of the first chunk, the computer launches a first document display window in the application user interface and displays therein a first document that includes the first chunk in its native display format. In response to a second user selection of the second chunk, the computer launches a second document display window in the application user interface and displays therein a second document that includes the second chunk in its native display format.

[00263] In some embodiments, the application user interface allows multiple document display windows associated with different document types to exist simultaneously. In some other embodiments, at one time, the application user interface only allows one document display window associated with a document type, e.g., by closing the first document display window before launching the second document display window in response to the second user selection of the second chunk.

[00264] In some embodiments, in response to a user request to view the chunk, the computer generates an empty region in the application user interface by shrinking the document authoring window and then occupies the empty region with the document display window in the application user interface.

[00265] In some embodiments, the portion of the corresponding document in the document display window includes more information about the search keywords than the chunk in the search results window, such as the location of the search keywords in the corresponding document or the textual contents adjacent to the search keywords in the corresponding document.

[00266] Figure 16F is a screenshot including a document display window 1653 in the process of being rendered within the application user interface in response to a user selection

of the link 1651. Note that the link 1651 is next to a chunk identified within a PowerPoint document. As shown in Figure 16G, the corresponding slide 1657 is displayed in the document display window and its location 1659 is highlighted in the document display window.

5 [00267] After viewing a chunk in the document display window, the author may want to duplicate the chunk in the document authoring window as well. As shown in Figures 16H-16J, respectively, in response to a user request to copy and paste a segment 1657 of the first document from the first document display window into the document authoring window, the computer generates therein an instance 1661 of the segment of the first document in its native display format; in response to a user request to copy and paste a segment 1663 of the second document display window into the document authoring window, the computer generates therein an instance 1665 of the segment 1663 of the second document in its native display format. This process is similar to the process described above in connection with Figures 16D and 16E.

15 [00268] In some embodiments, the document display window is a preview-only window of the corresponding document (e.g., a PDF document). The user cannot modify the document through the preview-only window. In some other embodiments, the document display window itself is a document authoring window, which may be another instance of the document authoring window (see, e.g., Figure 16I) or may be different from the original document authoring window (see, e.g., Figure 16G). Sometimes, the search keywords in the document display window are also highlighted.

[00269] Another application of the invention is to replace one text string with another text string among a set of documents without having to open any of them. For example, a user may like to change the name of a subject from A to B within many documents of different types that cover the subject. In some cases, the user may like to limit the change to certain types of documents or certain locations within the documents. Currently, the user has to open each document one by one and manually apply the change. This is not only time-consuming but also error-prone.

[00270] Figure 17A is a flowchart illustrative of a process of finding and replacing text strings in connection with a set of search results based on user instructions in accordance with some embodiments.

[00271] A computer receives a user request to replace a first text string with a second text string in a first document and a second document (1702). The first text string in the first document has a first content type and the first text string in the second document has a second content type, which is different from the first content type. The computer substitutes the second text string for the first text string in the first document and the second document (1704). The replacing second text string in the first document has the first content type and the replacing second text string in the second document has the second content type.

[00272] Figure 17B is a flowchart illustrative of a process of finding and replacing text strings within a set of documents based on user instructions in accordance with some embodiments.

[00273] After receiving a search request that includes one or more user-specified search keywords (1710), a computer identifies a first document and a second document (1712), each document having at least one chunk that satisfies the search keywords. A first text string in the first document has a first content type and the first text string in the second document has a second content type, which is different from the first content type. After receiving a user request to replace the first text string with a second text string (1714), the computer substitutes the second text string for the first text string in the first document and the second document (1716). The replacing second text string in the first document has the first content type and the replacing second text string in the second document has the second content type.

[00274] Figures 17C through 17E are screenshots of a graphical user interface on a computer display illustrative of features associated with the processes as shown in Figures 17A and 17B in accordance with some embodiments.

[00275] Figure 17C is a screenshot including a search assistant window 1722, which occupies the space in the application user interface previously occupied by the document display window (see, e.g., Figure 16J). In some embodiments, the search assistant window 1722 is activated by a user selection of the search assistant icon 1720. The search assistant window 1722 includes three tabs, "Search Options," "History," and "Replace." The "Replace" tab allows a user to replace one text string 1724 ("Albert Einstein" in this example) with another text string 1726 ("A. Einstein" in this example) by clicking the "Update Content" button 1727.

[00276] In some embodiments, the “Replace” tab provides additional options 1728 for the user to choose. For example, the user can limit the replacement to the selected search results in the search results window or relevant chunks in the identified documents, which documents result from a search and display of chunks that satisfy user-specified search keywords. Note that the text string 1724 to be replaced does not have to be related to the user-specified search keywords. They can be the same or overlapping (as is the case in Figure 17C) or completely different.

[00277] In some embodiments, the user can broaden the scope of the replacement to be the identified documents including, but not limited to, the relevant chunks. In some other embodiments, the user can further expand the scope to cover all the documents whether or not they have a relevant chunk.

[00278] In some embodiments, the “Replace” tab also allows the user to specify the locations within a document at which the replacement may happen. For example, Figure 17C depicts target options 1729 that include multiple locations, each having an associated checkbox. Thus, the user can stipulate that the first text string at one or more user-specified locations in the first and second documents be replaced by the second text string by checking the respective checkboxes. As a result, the computer substitutes the second text string for the first text string at the user-specified locations in the first document and the second document, respectively. Possible locations within a document include one or more selected from the group consisting of title, paragraph, table, header, footer, slide, spreadsheet, and all.

[00279] In some embodiments, after identifying the first document and the second document, the computer displays a first chunk from the first document and a second chunk from the second document, each chunk including at least one instance of the first text string. The instances of the first text string within the first and second chunks are displayed in a text-only display format. As described above, a PowerPoint document and a Word document are identified as having chunks satisfying the search phrase “Einstein general relativity.” The two relevant chunks are displayed in a text-only display format and different matching terms therein are highlighted in different colors.

[00280] In some embodiments, the first and second documents may have different document type. Note that a document’s document type is relevant to the document’s distinct appearance when the document is rendered through its native application. For example, the

first text string in the first document may have a first appearance when the first document is rendered by its native application and the first text string in the second document may have a second appearance that is different from the first appearance when the second document is rendered by its native application.

5 **[00281]** In this example, the Word document and the PowerPoint document have different document types because their contents have different appearances when rendered by Microsoft Office. Sometimes, a document's suffix may uniquely identify its document type, e.g., a document with the suffix ".docx" is a Microsoft Office 2007 Word document. Sometimes, a document's suffix cannot uniquely identify its document type, e.g., documents
10 like "hello.c" and "hello.java" are probably both plain-text documents and therefore have the same document type.

[00282] Figure 17D is a screenshot after the update is completed 1730. In some embodiments, replacing one text string with another text string does not trigger an update of the chunks in the search results window. Thus, the instances 1732, 1734 of the old text string
15 "Albert Einstein" still appear in the search results window. To view the replacing text string, the user has to perform a new search for the replacing text string.

[00283] As shown in Figure 17E, in response to a new search request including search keywords 1740 "Einstein general relativity," the computer updates the chunks in the search results window, and as a result, "Albert Einstein" is replaced with "A. Einstein." Note that
20 the instances 1742, 1744 of the replacing second text string within the first and second chunks are also displayed in the text-only display format.

[00284] In some embodiments, after substituting the second text string for the first text string, the computer also replaces the displayed instances of the first text string within the first and second chunks in the search results window with respective instances of the second
25 text string.

[00285] In some embodiments, the first document includes an original second text string that has a content type different from the replacing second text string. For example, the Word document may include a PowerPoint slide that has the phrase "A. Einstein," but not the phrase "general relativity." Assuming that the user limits the replacement to the chunks in
30 the search results window, after the update, when the Word document is rendered by

Microsoft Office, the second text string has at least two different appearances, one being a Word appearance and the other being a PowerPoint appearance.

[00286] Note that the methodology enabling the application of text string finding-and-replacement can be used for implementing other document-editing features such as undoing or reversing last N editing operations (including addition, deletion, and modification) applied to a set of documents and redoing or repeating N editing operations (including addition, deletion, and modification) applied to the set of documents. The set of documents may be located at the same data source or distributed across multiple data sources.

[00287] Another application of the invention is to refine search results using different search keywords. For example, after conducting one search using a set of search keywords, a user may like to conduct another search among the documents (or chunks) identified by the first search using another set of search keywords.

[00288] Figure 18A is a flowchart illustrative of a first process of narrowing search results based on user instructions in accordance with some embodiments.

[00289] After receiving a first user request including a first set of search keywords (1801), a computer identifies a first set of chunks within multiple documents (1803). Each chunk includes terms matching the first set of search keywords. The computer displays at least a portion of the first set of chunks (1805), including highlighting the terms matching the first set of search keywords in the displayed portion in a first manner. After receiving a second user request to search among the documents for documents that satisfy a second set of search keywords (1807), the computer identifies a second set of chunks within the documents (1809). Each chunk includes terms matching the second set of search keywords. The computer displays at least a portion of the second set of chunks (1811), including highlighting the terms matching the second set of search keywords in the displayed portion in a second manner that is different from the first manner.

[00290] Figure 18B is a flowchart illustrative of a second process of narrowing search results based on user instructions in accordance with some embodiments.

[00291] After receiving a first user request including a first set of search keywords (1821), a computer identifies multiple documents (1823). Each document includes at least one chunk that satisfies the first set of search keywords. After receiving a second user

request to search among the chunks in the identified documents for chunks that satisfy a second set of search keywords (1825), the computer identifies a subset of the chunks (1827). Each chunk in the subset satisfies the second set of search keywords.

[00292] Note that a user can repeat any of the two processes above for many times by providing different sets of search keywords for each search step until a predefined condition is met, e.g., the chunks of the user's interest have been found or no chunk is identified. At any time, the user can roll back the search process to a previously-identified set of chunks and try a different set of search keywords that has not been used previously.

[00293] Figures 18C through 18D are screenshots of a graphical user interface on a computer display illustrative of features associated with the processes as shown in Figures 18A and 18B in accordance with some embodiments.

[00294] The graphical user interface includes a first set of search results displayed in a text-only display format, each search result including one or more chunks identified within a respective document as satisfying a first set of search keywords. In response to a user request to search among the identified chunks for chunks that satisfy a second set of search keywords, the first set of search results is replaced by a second set of search results. Each search result in the second set includes one or more chunks identified within a respective document as satisfying both the first set of search keywords and the second set of search keywords. In some embodiments, two chunks identified within two different documents have different native display formats. In some embodiments, the second set of search keywords includes at least one search keyword that is not present in the first set of search keywords.

[00295] In some embodiments, terms matching the first set of search keywords and terms matching the second set of search keywords within a respective chunk are highlighted in a visually distinguishable manner.

[00296] In some embodiments, the chunks identified within a respective document as satisfying the first set of search keywords are displayed in an order consistent with their relative relevancy to the first set of search keywords, and the chunks identified within a respective document as satisfying both the first set of search keywords and the second set of search keywords are displayed in an order consistent with their relative relevancy to the second set of search keywords. In some other embodiments, the chunks identified within a respective document as satisfying any of the first and second sets of search keywords are

displayed in an order consistent with their relative locations within the corresponding document.

[00297] Figure 18C is a screenshot including a first set of relevant chunks 1833 identified within a PowerPoint document as satisfying the search keyword 1831 “A.

5 Einstein.” In some embodiments, the chunks 1833 are ordered by their respective relevancy to the search keywords 1831. In this example, the chunk 1835-B has a relative lower ranking metric when compared with the other chunks above (e.g., 1835-A) and is therefore displayed at the bottom of the search results window. In some embodiments, if the subset of chunks includes a first chunk and a second chunk, the computer displays the first chunk ahead of the
10 second chunk in response to the first user request and displays the second chunk ahead of the first chunk in response to the second user request.

[00298] Figure 18D is a screenshot including a second set of relevant chunks 1843 identified within the PowerPoint document as satisfying the search keyword 1841 “gravitation.” Note that the second set of search keywords 1841 can be completely different
15 from the first set of search keywords 1831. In this example, the user has selected the checkbox next to the “Search Within Results” icon 1847. Accordingly, the search for the second set of chunks is limited to the documents identified as having chunks that satisfy the search keywords 1831. In this case, it is possible that the second set of chunks includes at least one chunk that is not included in the first set of chunks. In some embodiments, the
20 search for the second set of chunks is further limited to the chunks 1833 that are identified by the first search.

[00299] In some embodiments, the second set of chunks includes at least one chunk that is included in the first set of chunks. For example, the chunks 1845-A, 1845-B in Figure 18D are the same as the respective chunks 1835-A, 1835-B in Figure 18C. In some
25 embodiments, the chunks 1835-A, 1835-B are displayed in an order consistent with their relevancy to the first set of search keywords 1831 in the first set of chunks and the chunks 1845-A, 1845-B are displayed in an order consistent with their relevancy to the second set of search keywords 1841 in the second set of chunks

[00300] In some embodiments, the terms in the chunks 1843 matching the first set of
30 search keywords 1831 and the terms in the chunks 1843 matching the second set of search

keywords are highlighted in different manner (e.g., different colors, font type, etc.). In this example, the matching terms are displayed using larger, italic, and underlined font.

[00301] At any time, if the user is unsatisfied with the identified chunks 1843, the user can bring back the previously-identified chunks by clicking the “Previous” link 1849-A and
5 restart the search process by entering a different set of search keywords. Similarly, the user can skip some search results by clicking the “Next” link 1849-B.

[00302] Another application of the invention is to minimize the response latency by alternatively processing different node streams to identify the relevant chunk within a node stream as quickly as possible.

10 [00303] Figure 19 is a flowchart illustrative of a process of alternatively processing document node streams in accordance with some embodiments.

[00304] The computer identifies a first candidate document at a first data source and a second candidate document at a second data source in response to a request from a user (1902). The request includes one or more keywords. In some embodiments, the request is a
15 search including one or more search keywords. The computer generates a first node stream for the first candidate document and a second node stream for the second candidate document using data packets received from the respective first and second data sources (1904). The computer alternatively processes the first node stream and the second node stream until a candidate chunk is identified therein (1906). In some embodiments, the candidate chunk
20 includes a set of nodes within a respective data source. Optionally, the computer returns the candidate chunk as a relevant chunk to the user if the candidate chunk satisfies the keywords (1908). Note that the first data source and the second data source may or may not be the same one. For example, they may be two different web servers. Thus, each candidate document can be an HTML web page.

25 [00305] In some embodiments, the computer submits an HTTP request to the first data source and receives an HTTP response from the first data source. The HTTP response may include multiple data packets corresponding to the first candidate document. After receiving one of the data packets from the first data source, the computer extracts one or more nodes from the data packet and inserts the one or more nodes into the first node stream. Sometimes,
30 the computer may extract only a node fragment from the data packet if the node is too large to fit in a single data packet. In this case, the computer then forms a node by combining the

node fragment with another node fragment, which may be extracted from a previous data packet, and insert the formed node (if the node is now complete) into the first node stream.

[00306] In some embodiments, after processing nodes currently in the first node stream, the computer waits for more nodes to appear in the first node stream. If no new node appears in the first node stream for a first amount of time, the computer may suspend processing the first node stream and switch to process nodes currently in the second node stream and identify the candidate chunk in the second node stream, if there is any one.

[00307] In some embodiments, after processing the nodes currently in the second node stream, the computer may switch back to process nodes currently in the first node stream if no new node appears in the second node stream for the first amount of time and identify the candidate chunk in the first node stream, if there is any one.

[00308] In some embodiments, the computer may discard processing results associated with one of the first node stream and the second node stream if no new node appears in the node stream for a second amount of time, which should be no less than and preferably longer than the first amount of time. For example, if there is a network traffic jam and the computer has not received any data packet from a remote data source for a relatively long period of time, the computer can stop working on the corresponding node stream and use the resources associated with the node stream for other purposes, e.g., processing another node stream.

[00309] Note that the HTTP-related example above is for illustrative purposes. The process is equally applicable to any communication protocol in which response latency is a concern, such as other TCP/IP based network protocols, file transfer protocol (FTP), or the like.

[00310] Another application of the invention is to provide a unified data model for documents having different structure types such as a strictly-structured XML document, a semi-structured HTML web page, and an unstructured plain-text document. This unified data model simplifies the process of identifying relevant chunks therein in response to a set of search keywords.

[00311] Figure 20 is a flowchart illustrative of a process of semantically annotating documents of different structures in accordance with some embodiments.

[00312] After retrieving a document from a data source (2001), the computer generates a customized data model (e.g., a hierarchical data mode) for the document in accordance with its structure type (2003). In some embodiments, the structure type can be structured, semi-structured, and unstructured. The computer identifies one or more candidate chunks within
5 the customized data model in accordance with a set of heuristic rules associated with the structure type (2005). Optionally, the computer selects one of the candidate chunks that satisfies one or more search keywords and returns it to an end user as a relevant chunk (2007).

[00313] In some embodiments, the data source is a web server and the document is an
10 HTML web page that includes multiple pairs of HTML tags. In this case, the computer identifies a first subset of the HTML web page between a first pair of HTML tags as a first candidate chunk if the first pair of HTML tags satisfies one of the set of heuristic rules. If necessary, the computer recursively identifies a second subset of the HTML web page within the first subset of the HTML web page between a second pair of HTML tags as a second
15 candidate chunk if the second pair of HTML tags satisfies one of the set of heuristic rules.

[00314] In some embodiments, for a plain-text document, the computer generates the data model by heuristically inserting metadata such as XML tags into the data model. The document contents following different XML tags are identified to be different candidate
20 chunks if they have predefined textual patterns. For example, a paragraph separated by blank lines is a candidate chunk and a sentence following a hyphen is also a candidate chunk if it is deemed to be part of a list of items.

[00315] Another application of the invention is to adaptively select matching chunks from a plurality of candidate chunks identified within a candidate document in response to a search request so as to improve the usability of the chunks to the end user.

[00316] As noted above in connection with Figure 12C, because the “‘Exact’ Match” and “Match All” options require all the search keywords find their matches in a chunk, they may ignore a chunk that, although highly relevant, fails to satisfy one of the search keywords. Alternatively, these two search options may return a chunk that, although satisfying all the search keywords, is too long to retain the benefits an ideal chunk should offer, e.g., being
30 both precise and efficient in locating the information of the user’s search interest. The latter

case is especially true if the candidate document has a hierarchical data model and the search keywords spread over multiple layers of the data model.

[00317] On the other hand, the “Match Any” option accepts any chunk that satisfies at least one search keyword. This could end up with returning too many short chunks to a user, which is equally frustrating because the user has to review many remotely matching chunks before locating the information of the user’s search interest or concluding that no such information is in the document.

[00318] Fortunately, the “Best Match” option, as will be described below, can successfully avoid the issues associated with these more polarized search options by screening out chunks that are potentially more distractive and presenting only chunks that satisfy a set of carefully-chosen criteria to the user.

[00319] Figure 21A is a flowchart illustrative of a first process of screening matching chunks within a candidate document based on predefined criteria in accordance with some embodiments. In this application, a “matching chunk” is defined as a candidate chunk that matches at least one search keyword. Certainly, a matching chunk could be an all-match if it matches all the search keywords and even an exact-match if it matches the search keywords in exactly the same order.

[00320] Assume that a set of matching chunks within the candidate document have been identified and they are fed into a computer in an order consistent with their respective locations in the document. The computer begins the adaptive process by checking if there is any more matching chunk to be further processed (2102). If so (2102, yes), the computer receives the next matching chunk (2104) and checks if the matching chunk meets the current minimum matching level set for the document (2106).

[00321] In some embodiments, a matching chunk is characterized by one or more attributes such as its matching level to the corresponding search request and its length. For example, the matching level of a matching chunk may be the total count of unique search keywords found within the chunk and the chunk’s length may be the total count of words or non-white-space characters in the chunk. Initially, the computer assigns a minimum matching level, e.g., one unique keyword per chunk, and a range of accepted chunk length, e.g., 50-70 words per chunk, to the candidate document.

[00322] If the matching level of the next matching chunk is below the minimum matching level (2106, no), the computer invalidates the matching chunk (2110) and proceeds to the next one in the pipeline. If the matching level of the next matching chunk is above the minimum matching level (2106, yes), the computer checks whether the chunk's length is within the range of accepted chunk length (2108). If the length of the chunk is outside the scope of accepted chunk length (2108, no), either too long or too short, the computer repeats the same procedure of invalidating the matching chunk (2110) and proceeds to the next one in the pipeline.

[00323] Otherwise (2108, yes), the computer inserts the matching chunk into a respective queue in accordance with the chunk's match level (2112). In some embodiments, matching chunks having different total counts of unique search keywords are put into separate queues. In some other embodiments, matching chunks having different sets of unique search keywords are grouped into separate queues. In either case, the computer calculates the current total count of matching chunks within the different queues (2113).

[00324] If the total count of matching chunks is greater than a predefined threshold, e.g., 10 chunks per document, the computer updates the document's current minimum matching level (2114) by, e.g., increasing the minimum matching level by one. As a result, at least one queue of matching chunks has a matching level less than the updated minimum matching level. In some embodiments, the computer invalidates the entire queue of matching chunks, re-determines the current total count of matching chunks, and repeats this procedure until the total count of matching chunks is less than the threshold. Certainly, the computer should not invalidate any matching chunk if the total count of matching chunks is less than the predefined threshold.

[00325] After updating the current minimum matching level, the computer checks whether the current minimum matching level has reached the maximum matching level associated with the search request (2116). In some embodiments, the maximum matching level is defined by identifying a best-matching chunk such as an all-match chunk or an exact-match chunk. If true (2116, yes), the computer outputs all the best-matching chunks it has accumulated in one or more queues to the user (2118). By doing so, the computer effectively reduces the latency by serving the presumably most relevant chunks to the user while continuously processing the other matching chunks. Otherwise (2116, no), the computer proceeds to the next one in the pipeline. In some embodiments, the operations 2116, 2118

are optional and the computer postpones returning any chunk to the user until after processing all the matching chunks.

[00326] At the end of the aforementioned process, the computer should filter out most, if not all, the distractive chunks that are presumably of little interest to the user and is now ready to serve the remaining matching chunks in the queues to the user. Assuming that the computer has queued multiple groups of matching chunks (2120, yes), it begins with serving a group of currently best-matching chunks to the user (2122). After that, the computer checks if the total count of matching chunks that have been served exceeds the predefined threshold or not (2124). If so (2124, yes), the computer stops the process of serving any more matching chunks even if there are additional queues of un-served matching chunks. By keeping the total count of served matching chunks below the threshold, the computer can avoid overwhelming the user with too many chunks in the search results view. Otherwise (2124, no), the computer repeats the process of serving the group of second best-matching chunks until the predefined threshold is met. In some embodiments, the computer stops serving any matching chunk if no more matching chunks are left in any queue (2120, no). This may occur even if the total count of served matching chunks has not reached the predefined threshold.

[00327] In some embodiments, the matching chunks identified within a document having a hierarchical data model are queued in an order such that a descendant matching chunk always precedes its ancestor matching chunks if they appear in the same queue. This ordering guarantees that the computer first serve the more refined descendant matching chunk before encountering any of the ancestor matching chunks because, as noted above, the serving process proceeds from perfect-matching chunks to less perfect ones. After serving the more refined descendant matching chunk, the computer also invalidates all the ancestor matching chunks in the same queue since none of them are presumably more relevant than the descendant chunk.

[00328] According to the aforementioned process, the matching chunks are served in an order consistent with their relevancy to the search request, which may be different from the order of the chunks' locations in the document. For example, a best-matching chunk served before the other matching chunks may be located at the end of the document and vice versa. In some embodiments, the computer may apply a two-phase process to ensure that the

matching chunks be served in an order consistent with their locations in the candidate document:

- Phase One – The computer screens the matching chunks as described above, including assigning a monotonically increasing chunk identifier to each matching chunk based on the matching chunk's location in the document and invalidating any chunk and its ancestors that fail to meet any of the predefined criteria, without serving any chunk to an end user.
- Phase Two – The computer sorts the surviving matching chunks within different queues in accordance with their respective chunk identifiers such that the first matching chunk to be served is located above the other matching chunks in the same document and outputs the matching chunks in this new sorted order.

[00329] Note that there are many other approaches of outputting chunks in an order consistent with their locations in the document. For example, the computer may generate a chunk linked-list during initial data model generation or matching chunk screening process such that each chunk includes a reference to the next chunk in the document. After the screening process, the computer can output the result matching chunks in an order consistent with their locations in the document by navigating the chunk linked-list and skipping invalidated chunks.

[00330] Figure 21B is an exemplary HTML document 2130 illustrative of the process as shown in Figure 21A in accordance with some embodiments. For illustration, the HTML document 2130 includes five matching chunks, each chunk having a unique chunk ID "cid."

[00331] Assume that there are seven user-specified search keywords, "*Scintillating Examples of the Best Match Algorithm*." Further assume that the predefined threshold of total chunk count is two (2), the range of accepted chunk length is 30-200 characters, and the initial minimum matching level is one keyword per chunk. The five matching chunks, each satisfying at least one of the seven search keywords, are fed into the computer in the order (as represented by their chunk IDs) of #2, #3, #1, #5, #4.

[00332] According to the flow chart shown in Figure 21A, chunks #2 and #3 are both placed in Queue 4, which contains the chunks matching four search keywords, although the

two chunks do not have the same four search keywords. Chunk #1 is placed in Queue 6, which contains the chunks matching six search keywords. Since three chunks have been placed into different queues, exceeding the threshold, the computer updates the current minimum matching level from “one keyword per chunk” to “four keywords per chunk.”

5 **[00333]** Although containing four matching keywords, chunk #5 is nonetheless invalidated because its length (26 characters) is outside the range of accepted chunk length. In contrast, chunk #4, which is a parent of chunk #5, is placed in Queue 4 for containing the same four matching keywords and being longer than 30 characters.

10 **[00334]** After processing all the matching chunks, the computer begins outputting the matching chunks within different queues. In this example, the computer outputs the chunks in an order consistent with their respective relevancy to the search request. Thus, chunk #1 in Queue 6 is first served to the user. As noted above, the export of chunk #1 also causes the invalidation of chunks #2 and #3 in Queue 4 because they are descendants of chunk #1. Because Queue 5 is empty, the computer proceeds to Queue 4, which has only chunk #5 left
15 for output. Finally, the computer stops the process after examining the queues of matching chunks with a matching level no less than the current minimum matching level.

[00335] Figure 21C is a flowchart illustrative of a second process of screening matching chunks within a document based on predefined criteria in accordance with some embodiments.

20 **[00336]** A computer identifies within a document multiple matching chunks in response to a search request from a user (2142). In some embodiments, the search request includes one or more search keywords and each of the multiple matching chunks matches at least one of the search keywords. The computer partitions the matching chunks into multiple groups (2144). The matching chunks within a respective group have an associated matching
25 level to the search request. In some embodiments, the partition is a queuing process wherein chunks containing the same number of matching keywords are placed in the same queue. The computer returns one or more groups of the matching chunks to the user in an order consistent with their respective matching levels to the search request (2136). In some embodiments, the computer displays a respective relevancy indicator adjacent each of the
30 returned matching chunks, indicating the relevancy between the corresponding matching chunk and the search request. The relevancy indicator can be formed using image, text,

number or the like to give the user an intuitive impression as to the matching chunk's proximity to the search keywords.

[00337] In some embodiments, each of the search keywords has an associated weight indicative of its relevance to the user's search interest. Different search keywords may have the same weight or different weights. Some of the search keywords may even have an associated weight of zero. For instance, in the example described above in connection with Figure 21B, the keyword "the" may be given a weight of zero and therefore have no impact on the search results.

[00338] In some embodiments, the matching level of a respective group of matching chunks is, at least partially, determined by summing the weights of unique search keywords within one of the matching chunks. For example, the matching level of a respective group of matching chunks may be, at least partially, determined by the number of unique search keywords within one of the matching chunks. If all the search keywords (including "the") are given the same weight, chunks #2 and #3 would have the same matching level and therefore be put in the same group.

[00339] In some embodiments, to partition the matching chunks into multiple groups, the computer selects one of the matching chunks, determining the chunk's matching level and length, and invalidates the chunk if its matching level is less than a minimum matching level or if its length is outside a predefined range of acceptable chunk length. If the selected matching chunk satisfies all the criteria including the minimum matching level and the predefined range of acceptable chunk length, the computer inserts the chunk into one of the groups of matching chunks. As noted above, the length of the matching chunk can be the total word count of the textual content of the matching chunk, or alternatively, the total character count of the textual content of the matching chunk after white-space normalization.

[00340] In some embodiments, after selecting a matching chunk that satisfies all the criteria, the computer compares the chunk's matching level to the matching level of a respective group of matching chunks until identifying a group of matching chunks whose matching levels are the same or similar to the selected chunk's matching level and then adds the chunk to the group of matching chunks.

[00341] In some embodiments, after placing a matching chunk within a group or exporting a matching chunk to the end user, the computer checks whether there are any

chunks within the same group that are descendants of the newly-placed or newly-exported matching chunk in a hierarchical data model of the document. If so, the computer then invalidates the descendant matching chunks from the group of matching chunks because they are redundant chunks from the user's perspective.

5 **[00342]** In some embodiments, after inserting one matching chunk into a group of matching chunks, the computer determines a total count of matching chunks whose matching levels are no less than the minimum matching level and updates the current minimum matching level if the total count of matching chunks is greater than a predefined count threshold. Additionally, the computer may invalidate at least a subset of one of the groups of
10 matching chunks whose matching levels are less than the updated minimum matching level.

[00343] In some embodiments, if there are multiple groups of matching chunks (e.g., Queue 6 and Queue 4 in the example shown in Figure 21B), the computer selects among the groups of matching chunks a group of matching chunks that has a highest matching level (e.g., Queue 6) and returns the selected group of matching chunks to the user. If there are
15 still groups of matching chunks left, the computer then returns to select a group of matching chunks having a next highest matching level (e.g., Queue 4) until the total count of the returned matching chunks is not less than a predefined count threshold.

[00344] Figure 21D is a screenshot of a graphical user interface on a computer display illustrative of features associated with the processes as shown in Figures 21A and 21B in
20 accordance with some embodiments. In this example, the search keywords box includes five search keywords 2150, "*distance between earth and moon*," and the "Best Match" search option is chosen for selecting matching chunks.

[00345] Based on these search keywords, it is not difficult to appreciate that the user is probably interested in knowing the spatial distance between the earth and the moon. But as
25 shown in Figure 21D, the search result 2154 provided by a generic search tool is not satisfactory because it has nothing to do with the answer expected by the user although all the four search keywords are present in the search result (note that the term "and" is treated as a stop-word with no weight).

[00346] In contrast, a process according to some embodiments of the invention
30 identifies multiple matching chunks within the same document, 2152-A through 2152-C, different chunks having different numbers of search keywords. In this example, the matching

chunks are ordered by their matching levels to the search keywords. Therefore, the matching chunk 2152-A appears before the other two chunks because it includes at least one instance of each of the four search keywords, which is essentially an all-match chunk. But this chunk does not have the answer to the user's question either. Actually, it is the second matching
5 chunk 2152-B that, although having no match for the search keyword "between," has the answer to the user's question, that is, the phrase 2156 "*distance from the Earth to the Moon is 384,403 km.*" Thus, the user receives a satisfactory answer to his or her question from the matching chunks without visiting any of the candidate documents. Note that the same matching chunk 2152-B would have been ignored by the "Match All" and "Exact Match"
10 options because it does not have the keyword "between."

[00347] Another application of the invention is to search a set of inter-related documents for contents matching a search request. This application is different from the conventional search tools, which always treat the Internet as the search space and perform all the searches in the entire search space no matter how irrelevant most of the documents in the
15 space are to the user-specified search keywords. Consequently, many documents identified by the conventional search tool, although have nothing to do with the user's search interest, end up occupying prominent spots in the search results window. If a user is allowed to narrow the search space to a small set of user-specified documents, it is possible for a computer to produce more relevant search results at a fraction of the cost wasted by the
20 conventional search tools.

[00348] Figure 22A is a flowchart illustrative of a process of identifying contents matching a search request within a plurality of inter-related documents in accordance with some embodiments. In this application, a first document is inter-related to a second document if the first document includes a document link that either directly references the
25 second document or indirectly references the second document by referencing a third document that directly or indirectly references the second document. The first document is also inter-related to the second document if they are both directly or indirectly referenced by a third document. As such, different documents referenced by respective document links within an HTML web page are referred to as "inter-related documents." In this case, the
30 HTML web page is called "primary document" and the documents referenced by the web page are called "secondary documents."

[00349] A computer receives a request to search one or more secondary documents (2201). At least one of the secondary documents is associated with a primary document. The computer searches at least a subset of the secondary documents for documents that satisfy the search request (2203) and identifies at least one secondary document that satisfies the search request (2205).

[00350] In some embodiments, the computer first displays the primary document (e.g., a web page) on a display device (e.g., a computer monitor) before receiving the search request from a user. The primary document includes one or more document links, each document link referencing one of the secondary documents. After identifying the secondary document, which may be another web page or the like, the computer displays at least a portion of the identified secondary document to the user. The displayed portion of the secondary document preferably includes one or more search keywords in the search request.

[00351] In some embodiments, the computer locates within the identified secondary document one or more chunks that satisfy the search request using the approaches as described above and displays one or more of the identified chunks to the user.

[00352] In some embodiments, the primary document includes many document links pointing to a large number of secondary documents, many of which may have nothing to do with the user's search interest. For example, many web pages include links to boilerplate-type secondary documents such as "About Us," "Contact Us," "Sitemap," "Disclaimer," etc. Searching out these secondary documents rarely returns any useful search results. Thus, in some embodiments, rather than searching all the secondary documents referenced by the primary document, the user is allowed to select a subset of secondary documents to be searched by identifying document links associated with the user-selected secondary document.

[00353] For example, each of the subset of secondary documents can be selected by a respective mouse click of the corresponding document link in the primary document. Alternatively, the computer defines a region in the primary document using an input device and then identifies document links within the defined region as the user-selected document links. For example, the computer presses down a mouse's button at a first location and drags the mouse from one location to another location until releasing the mouse's button at a second location. By doing so, the user-selected region is a rectangle area defined by the first

location and the second location and all the document links falling into this rectangle area are document links to secondary documents to be further searched in response to a search request.

[00354] In some embodiments, the computer searches both the primary and secondary documents for chunks that satisfy the search request, and as a result, identifies at least one chunk in the primary document and at least one chunk in one of the secondary documents, both chunks satisfying the search request. The chunks associated with the primary and secondary documents are visually separated by a bar such that it is intuitive for a user to distinguish chunks identified within the primary document and chunks identified within the secondary documents.

[00355] In some embodiments, the search of secondary documents is a recursive process. In response to a user request to search a secondary document, the computer recursively retrieves the secondary document and documents referenced by this secondary document. Thus, the search results may not only include chunks identified within the primary document but also chunks within a secondary document that is indirectly referenced by the primary document.

[00356] Figures 22B through 22D are screenshots of a graphical user interface on a computer display illustrative of features associated with the process as shown in Figure 22A in accordance with some embodiments.

[00357] Figure 22B is a screenshot of a web browser window rendering a web page identified by the URL <http://www.rainingdata.com/products/index.html>. There are two user-specified search keywords “tigerlogic xdms” in the search box. The screenshot depicts at least chunks 2217-A through 2217-D that match the two search keywords. The web page includes many document links. Some of the document links (e.g., links 2219) are likely to be related to the search keywords 2213 and others (e.g., links 2220) probably have nothing to do with the search keywords 2213. In this example, the user avoids searching secondary documents associated with the links 2220 by either mouse-clicking the links or defining a rectangle region covering the links.

[00358] After a user mouse-click of the “Chunk Page Links” icon 2215, the computer generates a plurality of chunks identified within the primary document and the secondary documents identified by the links 2219 as shown in the screenshot of Figure 22C. Note that

the search results 2221 associated with the primary document (including the chunks 2217-A through 2217-C) are separated from the search results 2225 and 2229, which are associated with the two secondary documents identified by the two links 2219, each including a respective set of matching chunks 2227's and 2231's. Figure 22D is another screenshot that only depicts the search results from the secondary documents, nothing from the primary document.

[00359] Figure 23 is a block diagram of an exemplary document search server 2300 computer in accordance with some embodiments.

[00360] The exemplary document search server 2300 typically includes one or more processing units (CPU's) 2302, one or more network or other communications interfaces 2310, memory 2312, and one or more communication buses 2314 for interconnecting these components. The communication buses 2314 may include circuitry (sometimes called a chipset) that interconnects and controls communications between system components. The document search server 2300 may optionally include a user interface, for instance a display and a keyboard. Memory 2312 may include high speed random access memory and may also include non-volatile memory, such as one or more magnetic disk storage devices. Memory 2312 may include mass storage that is remotely located from the CPU's 2302. In some embodiments, memory 2312 stores the following programs, modules and data structures, or a subset or superset thereof:

- an operating system 2316 that includes procedures for handling various basic system services and for performing hardware dependent tasks;
- a network communication module 2318 that is used for connecting the document search server 2300 to other servers or computers via one or more communication networks (wired or wireless), such as the Internet, other wide area networks, local area networks, metropolitan area networks, and so on;
- a system initialization module 2320 that initializes other modules and data structures stored in memory 2312 required for the appropriate operation of the document search server 2300;
- a query engine 2322 for processing a user-driven search query and preparing relevant chunks in response to the search query;

- a cache engine 2324 for identifying candidate documents in response to the search query;
- a stream engine 2326 for retrieving candidate documents and identifying candidate chunks therein; and
- 5 • an index database 2328 for storing index information of a number of candidate documents 2330 accessible to the document search server 2300.

[00361] Figure 24 is a block diagram of an exemplary client computer 2400 in accordance with some embodiments.

[00362] The exemplary client computer 2400 typically includes one or more
10 processing units (CPU's) 2402, one or more network or other communications interfaces 2410, memory 2412, and one or more communication buses 2414 for interconnecting these components. The communication buses 2414 may include circuitry (sometimes called a chipset) that interconnects and controls communications between system components. The client computer 2400 may include a user input device 2410, for instance a display and a
15 keyboard. Memory 2412 may include high speed random access memory and may also include non-volatile memory, such as one or more magnetic disk storage devices. Memory 2412 may include mass storage that is remotely located from the CPU's 2402. In some embodiments, memory 2412 stores the following programs, modules and data structures, or a subset or superset thereof:

- 20 • an operating system 2416 that includes procedures for handling various basic system services and for performing hardware dependent tasks;
- a network communication module 2418 that is used for connecting the client computer 2400 to the document search server 2300 or other computers via one or more communication networks (wired or wireless), such as the Internet, other wide
25 area networks, local area networks, metropolitan area networks, and so on;
- a system initialization module 2419 that initializes other modules and data structures stored in memory 2412 required for the appropriate operation of the client computer 2400;
- a web browser 2420 for retrieving and displaying candidate documents including web
30 pages from remote web servers;

- a search toolbar 2425 attached to the web browser 2420 for identifying relevant chunks within the retrieved candidate document and displaying the relevant chunks;
- one or more applications 2430 such as Microsoft Office Word application 2431, Microsoft Office PowerPoint application 2433, Microsoft Office Excel application 2435, etc.; and
- an add-in application 2437 attached to the Microsoft Office applications for displaying relevant chunks associated with user-specified search keywords and re-using the relevant chunks based on user instructions.

[00363] The foregoing description, for purpose of explanation, has been described with reference to specific embodiments. However, the illustrative discussions above are not intended to be exhaustive or to limit the invention to the precise forms disclosed. Many modifications and variations are possible in view of the above teachings. For example, the aforementioned processes of identifying a relevant chunk within a document are by no means limited to a particular language such as English. Actually, the same processes are equally applicable to documents written in other languages and/or multi-lingual documents. The embodiments were chosen and described in order to best explain the principles of the invention and its practical applications, to thereby enable others skilled in the art to best utilize the invention and various embodiments with various modifications as are suited to the particular use contemplated.

What is claimed is:

1. A computer-implemented method, comprising:

at a client computer,

submitting one or more user-specified search keywords to a server computer;

5 receiving a set of search results from the server computer, wherein each search result identifies a document located at a respective document source that satisfies the search keywords in accordance with a first set of predefined criteria;

for each identified document,

retrieving the document from the corresponding document source;

10 identifying a chunk within the document that satisfies the search query in accordance with a second set of predefined criteria; and

displaying the identified chunk and a link to the identified chunk within the document.

2. A computer-implemented method, comprising:

15 identifying multiple resource identifiers, each resource identifier corresponding to a document at a respective data source;

for at least one of the resource identifiers,

retrieving the corresponding document from the respective document source;

20 identifying within the retrieved document a chunk that satisfies one or more user-specified search keywords; and

displaying the identified chunk and a link to the identified chunk within the document to the user.

3. The method of claim 2, wherein terms in the identified chunk that match the search keywords are either ordered differently from the user-specified search keywords or separated
25 from one another by at least one term not matching any of the search keywords.

4. The method of claim 2, wherein the identified chunk includes an identical instance of the search keywords appearing as a phrase.

5. The method of claim 2, further comprising:

30 highlighting different terms matching different search keywords in the identified chunk in different manners.

6. The method of claim 2, further comprising:
highlighting different terms matching different search keywords in the identified chunk in different colors.
7. The method of claim 2, further comprising:
5 identifying one or more sponsored links associated with at least one of the search keywords; and
displaying the sponsored links adjacent the identified chunk.
8. The method of claim 2, further comprising:
submitting the user-specified search keywords to a search engine; and
10 receiving a set of search results from the search engine, wherein each search result includes an abbreviated document segment identified by the search engine as satisfying the search keywords and the abbreviated document segment is different from the identified chunk.
9. The method of claim 2, further comprising:
15 identifying multiple chunks within the retrieved document, wherein each chunk satisfies the search keywords; and
displaying the multiple chunks and a link to each of the multiple chunks within the document.
10. The method of claim 9, further comprising:
20 displaying the multiple chunks in an order consistent with their relative relevancy to the search keywords.
11. The method of claim 9, further comprising:
displaying the multiple chunks in an order consistent with their relative locations within the document.
- 25 12. The method of claim 2, further comprising:
for at least one of the resource identifiers,
retrieving the corresponding document from the respective document source;
identifying within the retrieved document no chunk that satisfies the search keywords; and

displaying a link to search for chunks that satisfy any of the search keywords within the document.

13. The method of claim 12, further comprising:

in response to a user selection of the link to search for chunks that satisfy any of the search keywords within the document,

re-processing the document retrieved from the respective document source;

identifying within the retrieved document one or more chunks that satisfy at least one of the search keywords; and

displaying the identified chunks and a link to each of the identified chunks within the document.

14. The method of claim 2, wherein the identified document is a web page and the respective document source is a web server hosting the web page.

15. The method of claim 2, further comprising:

in response to a user selection of the link to the identified chunk,

displaying at least a portion of the identified document in a document view window, wherein the displayed portion includes, at least partially, the identified chunk;

highlighting the identified chunk in the document view window; and

highlighting terms that match the search keywords in the identified chunk such that they are visually distinguishable over the rest of the identified chunk.

16. The method of claim 15, wherein the identified chunk is defined by a pair of unique chunk virtual delimiters within the identified document.

17. The method of claim 16, wherein the pair of unique chunk virtual delimiters is invisible when the identified document is rendered by an application.

18. The method of claim 2, wherein the identification of a link within the retrieved document further includes:

generating a search query using the search keywords; and

applying the search query to the retrieved document to identify the chunk.

19. The method of claim 2, further comprising:

for the at least one of the resource identifiers, displaying a chunk page-link icon for searching chunks within documents that are referenced by the corresponding document; and

in response to a user selection of the chunk page-link icon, displaying one or more referenced document links, wherein each referenced document link has one or more associated chunks identified within a corresponding referenced document as satisfying the user-specified search keywords.

5 20. The method of claim 2, further comprising:

for the at least one of the resource identifiers, displaying a hide-chunk icon; and

in response to a user selection of the hide-chunk icon, causing the disappearance of the identified chunk.

21. A computer system, comprising:

10 memory;

one or more processors;

one or more programs stored in the memory and configured for execution by the one or more processors, the one or more programs including:

instructions for identifying multiple resource identifiers, each resource

15 identifier corresponding to a document at a respective data source;

instructions for retrieving the corresponding document from the respective document source for at least one of the resource identifiers;

instructions for identifying within the retrieved document a chunk that satisfies one or more user-specified search keywords; and

20 instructions for displaying the identified chunk and a link to the identified chunk within the document to the user.

22. The computer system of claim 21, further comprising:

instructions for ordering terms in the identified chunk that match the search keywords either differently from the user-specified search keywords or separating the terms from one
25 another by at least one term not matching any of the search keywords.

23. The computer system of claim 21, wherein the identified chunk includes an identical instance of the search keywords appearing as a phrase.

24. The computer system of claim 21, further comprising:

instructions for highlighting different terms matching different search keywords in the
30 identified chunk in different manners.

25. The computer system of claim 21, further comprising:
instructions for highlighting different terms matching different search keywords in the identified chunk in different colors.

26. The computer system of claim 21, further comprising:

5 instructions for identifying one or more sponsored links associated with at least one of the search keywords; and
instructions for displaying the sponsored links adjacent the identified chunk.

27. The computer system of claim 21, further comprising:

instructions for submitting the user-specified search keywords to a search engine; and
10 instructions for receiving a set of search results from the search engine, wherein each search result includes an abbreviated document segment identified by the search engine as satisfying the search keywords and the abbreviated document segment is different from the identified chunk.

28. The computer system of claim 21, further comprising:

15 instructions for identifying multiple chunks within the retrieved document, wherein each chunk satisfies the search keywords; and
instructions for displaying the multiple chunks and a link to each of the multiple chunks within the document.

29. The computer system of claim 28, further comprising:

20 instructions for displaying the multiple chunks in an order consistent with their relative relevancy to the search keywords.

30. The computer system of claim 28, further comprising:

instructions for displaying the multiple chunks in an order consistent with their relative locations within the document.

25 31. The computer system of claim 21, further comprising:

instructions for retrieving the corresponding document from the respective document source for at least one of the resource identifiers;

instructions for identifying within the retrieved document no chunk that satisfies the search keywords; and

instructions for displaying a link to search for chunks that satisfy any of the search keywords within the document.

32. The computer system of claim 31, further comprising:

instructions for re-processing the document retrieved from the respective document source in response to a user selection of the link to search for chunks that satisfy any of the search keywords within the document;

instructions for identifying within the retrieved document one or more chunks that satisfy at least one of the search keywords; and

instructions for displaying the identified chunks and a link to each of the identified chunks within the document.

33. The computer system of claim 21, wherein the identified document is a web page and the respective document source is a web server hosting the web page.

34. The computer system of claim 21, further comprising:

instructions for displaying at least a portion of the identified document in a document view window in response to a user selection of the link to the identified chunk, wherein the displayed portion includes, at least partially, the identified chunk;

instructions for highlighting the identified chunk in the document view window; and
instructions for highlighting terms that match the search keywords in the identified chunk such that they are visually distinguishable over the rest of the identified chunk.

35. The computer system of claim 34, wherein the identified chunk is defined by a pair of unique chunk virtual delimiters within the identified document.

36. The computer system of claim 35, wherein the pair of unique chunk virtual delimiters is invisible when the identified document is rendered by an application.

37. The computer system of claim 21, wherein the instructions for identifying a link within the retrieved document further include:

instructions for generating a search query using the search keywords; and
instructions for applying the search query to the retrieved document to identify the chunk.

38. The computer system of claim 21, further comprising:

instructions for displaying a chunk page-link icon for searching chunks within documents that are referenced by the corresponding document for the at least one of the resource identifiers; and

instructions for displaying one or more referenced document links in response to a user selection of the chunk page-link icon, wherein each referenced document link has one or more associated chunks identified within a corresponding referenced document as satisfying the user-specified search keywords.

39. The computer system of claim 21, further comprising:

instructions for displaying a hide-chunk icon for the at least one of the resource identifiers; and

instructions for causing the disappearance of the identified chunk in response to a user selection of the hide-chunk icon.

40. A computer readable storage medium having stored therein instructions, which when executed by a computer system cause the computer system to:

identify multiple resource identifiers, each resource identifier corresponding to a document at a respective data source;

retrieve the corresponding document from the respective document source for at least one of the resource identifiers;

identify within the retrieved document a chunk that satisfies one or more user-specified search keywords; and

display the identified chunk and a link to the identified chunk within the document to the user.

41. The computer readable storage medium of claim 40, further comprising:

instructions for ordering terms in the identified chunk that match the search keywords either differently from the user-specified search keywords or separating the terms from one another by at least one term not matching any of the search keywords.

42. The computer readable storage medium of claim 40, wherein the identified chunk includes an identical instance of the search keywords appearing as a phrase.

43. The computer readable storage medium of claim 40, further comprising:

instructions for highlighting different terms matching different search keywords in the identified chunk in different manners.

44. The computer readable storage medium of claim 40, further comprising:
instructions for highlighting different terms matching different search keywords in the
identified chunk in different colors.

45. The computer readable storage medium of claim 40, further comprising:
5 instructions for identifying one or more sponsored links associated with at least one of
the search keywords; and
instructions for displaying the sponsored links adjacent the identified chunk.

46. The computer readable storage medium of claim 40, further comprising:
instructions for submitting the user-specified search keywords to a search engine; and
10 instructions for receiving a set of search results from the search engine, wherein each
search result includes an abbreviated document segment identified by the search engine as
satisfying the search keywords and the abbreviated document segment is different from the
identified chunk.

47. The computer readable storage medium of claim 40, further comprising:
15 instructions for identifying multiple chunks within the retrieved document, wherein
each chunk satisfies the search keywords; and
instructions for displaying the multiple chunks and a link to each of the multiple
chunks within the document.

48. The computer readable storage medium of claim 47, further comprising:
20 instructions for displaying the multiple chunks in an order consistent with their
relative relevancy to the search keywords.

49. The computer readable storage medium of claim 47, further comprising:
instructions for displaying the multiple chunks in an order consistent with their
relative locations within the document.

25 50. The computer readable storage medium of claim 40, further comprising:
instructions for retrieving the corresponding document from the respective document
source for at least one of the resource identifiers;
instructions for identifying within the retrieved document no chunk that satisfies the
search keywords; and

instructions for displaying a link to search for chunks that satisfy any of the search keywords within the document.

51. The computer readable storage medium of claim 50, further comprising:

instructions for re-processing the document retrieved from the respective document source in response to a user selection of the link to search for chunks that satisfy any of the search keywords within the document;

instructions for identifying within the retrieved document one or more chunks that satisfy at least one of the search keywords; and

instructions for displaying the identified chunks and a link to each of the identified chunks within the document.

52. The computer readable storage medium of claim 40, wherein the identified document is a web page and the respective document source is a web server hosting the web page.

53. The computer readable storage medium of claim 40, further comprising:

instructions for displaying at least a portion of the identified document in a document view window in response to a user selection of the link to the identified chunk, wherein the displayed portion includes, at least partially, the identified chunk;

instructions for highlighting the identified chunk in the document view window; and

instructions for highlighting terms that match the search keywords in the identified chunk such that they are visually distinguishable over the rest of the identified chunk.

54. The computer readable storage medium of claim 53, wherein the identified chunk is defined by a pair of unique chunk virtual delimiters within the identified document.

55. The computer readable storage medium of claim 54, wherein the pair of unique chunk virtual delimiters is invisible when the identified document is rendered by an application.

56. The computer readable storage medium of claim 40, wherein the instructions for identifying a link within the retrieved document further include:

instructions for generating a search query using the search keywords; and

instructions for applying the search query to the retrieved document to identify the chunk.

57. The computer readable storage medium of claim 40, further comprising:

instructions for displaying a chunk page-link icon for searching chunks within documents that are referenced by the corresponding document for the at least one of the resource identifiers; and

5 instructions for displaying one or more referenced document links in response to a user selection of the chunk page-link icon, wherein each referenced document link has one or more associated chunks identified within a corresponding referenced document as satisfying the user-specified search keywords.

58. The computer readable storage medium of claim 40, further comprising:

10 instructions for displaying a hide-chunk icon for the at least one of the resource identifiers; and

instructions for causing the disappearance of the identified chunk in response to a user selection of the hide-chunk icon.

59. A graphical user interface on a computer display, comprising:

15 one or more document links, wherein each document link has one or more associated chunks identified within the corresponding document as satisfying one or more user-specified search keywords, wherein each chunk has an associated chunk link and includes terms matching each of the user-specified search keywords, wherein the terms are highlighted in the chunk in a visually distinguishable manner, wherein:

20 in response to a user selection of a chunk's chunk link, the corresponding document is displayed in a window on the computer display, wherein at least a portion of the chunk is highlighted in the window.

60. The graphical user interface of claim 59, wherein each document link has an associated chunk page-link icon for searching chunks within documents that are referenced by the corresponding document, wherein:

25 in response to a user selection of a document link's associated chunk page-link icon, one or more referenced document links are displayed on the computer display, wherein each referenced document link has one or more associated chunks identified within a corresponding referenced document as satisfying the user-specified search keywords.

30 61. The graphical user interface of claim 59, wherein each document link has an associated hide-chunk icon, wherein:

in response to a user selection of a document link's associated hide-chunk icon, the one or more chunks associated with the document link and their associated chunk links disappear from the computer display.

62. The graphical user interface of claim 59, wherein one or more chunks associated with a respective document link are displayed in an order consistent with their relative relevancy to the user-specified search keywords.

63. The graphical user interface of claim 59, wherein one or more chunks associated with a respective document link are displayed in an order consistent with their relative locations within the corresponding document.

64. A computer-implemented method, comprising:
while a user is browsing a document,
receiving multiple user-specified search keywords, wherein the search keywords have a first sequence;
identifying within the document at least one chunk that satisfies the search keywords; and
displaying a portion of the document including the identified chunk, wherein the search keywords are highlighted in the identified chunk and have a second sequence that is different from the first sequence.

65. A computer-implemented method, comprising:
displaying a portion of a document to a user;
upon receiving a user-specified text string that includes multiple search keywords, identifying a chunk within the document that satisfies the search keywords; and
displaying the identified chunk to the user, wherein terms in the identified chunk that match the search keywords are either ordered differently from the search keywords in the user-specified text string or separated from one another by at least one term not matching any of the search keywords.

66. The method of claim 65, further comprising:
highlighting different terms matching different search keywords in the identified chunk in different manners.

67. The method of claim 65, further comprising:

highlighting different terms matching different search keywords in the identified chunk in different colors.

68. The method of claim 65, further comprising:

identifying multiple chunks within the document, wherein each chunk satisfies the search keywords and appears at a respective location in the document; and displaying, at least partially, the chunk that appears above the other chunks in the document and its associated context.

69. The method of claim 65, further comprising:

identifying multiple chunks within the document, wherein each chunk satisfies the search keywords and has a ranking metric of its relevancy to the search keywords; and displaying, at least partially, the chunk that has the highest ranking metric than the other chunks in the document and its associated context.

70. The method of claim 65, wherein the identified document is a web page.

71. The method of claim 65, further comprising:

if there is no chunk within the document that satisfies each of the search keywords, identifying within the document one or more chunks that satisfy any of the search keywords; and displaying the identified chunks to the user.

72. The method of claim 65, wherein the identifying further includes:

generating a search query using the search keywords; and applying the search query to the document to identify the chunk.

73. The method of claim 65, wherein the identified chunk is not within the initially displayed portion of the document.

74. A computer system, comprising:

memory;

one or more processors;

one or more programs stored in the memory and configured for execution by the one or more processors, the one or more programs including:

instructions for displaying a portion of a document to a user;

instructions for identifying a chunk within the document that satisfies the search keywords upon receiving a user-specified text string that includes multiple search keywords; and

instructions for displaying the identified chunk to the user, wherein terms in the identified chunk that match the search keywords are either ordered differently from the search keywords in the user-specified text string or separated from one another by at least one term not matching any of the search keywords.

75. The computer system of claim 74, further comprising:

instructions for highlighting different terms matching different search keywords in the identified chunk in different manners.

76. The computer system of claim 74, further comprising:

instructions for highlighting different terms matching different search keywords in the identified chunk in different colors.

77. The computer system of claim 74, further comprising:

instructions for identifying multiple chunks within the document, wherein each chunk satisfies the search keywords and appears at a respective location in the document; and instructions for displaying, at least partially, the chunk that appears above the other chunks in the document and its associated context.

78. The computer system of claim 74, further comprising:

instructions for identifying multiple chunks within the document, wherein each chunk satisfies the search keywords and has a ranking metric of its relevancy to the search keywords; and

instructions for displaying, at least partially, the chunk that has the highest ranking metric than the other chunks in the document and its associated context.

79. The computer system of claim 74, wherein the identified document is a web page.

80. The computer system of claim 74, further comprising:

instructions for identifying within the document one or more chunks that satisfy any of the search keywords if there is no chunk within the document that satisfies each of the search keywords; and

instructions for displaying the identified chunks to the user.

81. The computer system of claim 74, wherein the instructions for identifying a chunk within the document further include:

instructions for generating a search query using the search keywords; and
instructions for applying the search query to the document to identify the chunk.

5 82. The computer system of claim 74, wherein the identified chunk is not within the initially displayed portion of the document.

83. A computer readable storage medium having stored therein instructions, which when executed by a computer system cause the computer system to:

display a portion of a document to a user;

10 identify a chunk within the document that satisfies the search keywords upon receiving a user-specified text string that includes multiple search keywords; and

display the identified chunk to the user, wherein terms in the identified chunk that match the search keywords are either ordered differently from the search keywords in the user-specified text string or separated from one another by at least one term not matching any
15 of the search keywords.

84. The computer readable storage medium of claim 83, further comprising:

instructions for highlighting different terms matching different search keywords in the identified chunk in different manners.

85. The computer readable storage medium of claim 83, further comprising:

20 instructions for highlighting different terms matching different search keywords in the identified chunk in different colors.

86. The computer readable storage medium of claim 83, further comprising:

instructions for identifying multiple chunks within the document, wherein each chunk satisfies the search keywords and appears at a respective location in the document; and

25 instructions for displaying, at least partially, the chunk that appears above the other chunks in the document and its associated context.

87. The computer readable storage medium of claim 83, further comprising:

instructions for identifying multiple chunks within the document, wherein each chunk satisfies the search keywords and has a ranking metric of its relevancy to the search

30 keywords; and

instructions for displaying, at least partially, the chunk that has the highest ranking metric than the other chunks in the document and its associated context.

88. The computer readable storage medium of claim 83, wherein the identified document is a web page.

5 89. The computer readable storage medium of claim 83, further comprising:
instructions for identifying within the document one or more chunks that satisfy any of the search keywords if there is no chunk within the document that satisfies each of the search keywords; and
instructions for displaying the identified chunks to the user.

10 90. The computer readable storage medium of claim 83, wherein the instructions for identifying a chunk within the document further include:
instructions for generating a search query using the search keywords; and
instructions for applying the search query to the document to identify the chunk.

15 91. The computer readable storage medium of claim 83, wherein the identified chunk is not within the initially displayed portion of the document.

92. A computer-implemented method, comprising:
identifying a document in response to a search request from a user, wherein the document includes content data and metadata and the search request includes one or more search keywords;
20 generating a hierarchical semantic model of the content data of the document by applying heuristics to the metadata of the document;
identifying a chunk within the document by scanning the hierarchical semantic model, wherein the identified chunk includes a subset of the content data that satisfies the search keywords and the corresponding metadata; and
25 returning the identified chunk to the requesting user.

93. The method of claim 92, wherein the document is a semi-structured document.

94. The method of claim 92, wherein generating the hierarchical semantic model further includes identifying one or more candidate chunks in the document, each candidate chunk corresponding to a respective subset of the document.

95. The method of claim 94, wherein applying the heuristics further includes identifying a subset of the document as a candidate chunk if the subset of the document has at least one instance of predefined metadata.

96. The method of claim 94, wherein applying the heuristics further includes identifying a subset of the document as a candidate chunk if the subset of the document has at least two instances of predefined metadata.

97. The method of claim 94, wherein a first subset of the document associated with a first candidate chunk encompasses a second subset of the document associated with a second candidate chunk.

98. The method of claim 97, wherein the search keywords include a first search keyword and a second search keyword and scanning the hierarchical semantic model further includes:
identifying content data in the first candidate chunk and preceding the second candidate chunk that satisfies the first search keyword;
identifying content data in the second candidate chunk that satisfies the second search keyword; and
choosing the first candidate chunk as the identified chunk.

99. The method of claim 98, further comprising:
identifying content data in the second candidate chunk that satisfies the first search keyword and the second keyword; and
choosing the second candidate chunk as the identified chunk.

100. The method of claim 94, wherein a first subset of the document associated with a first candidate chunk encompasses a second subset of the document associated with a second candidate chunk and a third subset of the document associated with a third candidate chunk and wherein there is no overlapping between the second candidate chunk and the third candidate chunk.

101. The method of claim 100, wherein the search keywords include a first search keyword and a second search keyword and scanning the hierarchical semantic model further includes:
identifying content data satisfying the first search keyword in the second candidate chunk;

identifying content data satisfying the second search keyword in the third candidate chunk; and

choosing the first candidate chunk as the identified chunk.

102. The method of claim 100, wherein there is no intermediate candidate chunk within the first candidate chunk that encompasses the second candidate chunk and the third candidate chunk.

103. The method of claim 92, further comprising:

identifying the chunk within the document by sequentially scanning the hierarchical semantic model.

104. A computer system, comprising:

memory;

one or more processors;

one or more programs stored in the memory and configured for execution by the one or more processors, the one or more programs including:

instructions for identifying a document in response to a search request from a user, wherein the document includes content data and metadata and the search request includes one or more search keywords;

instructions for generating a hierarchical semantic model of the content data of the document by applying heuristics to the metadata of the document;

instructions for identifying a chunk within the document by scanning the hierarchical semantic model, wherein the identified chunk includes a subset of the content data that satisfies the search keywords and the corresponding metadata; and

instructions for returning the identified chunk to the requesting user.

105. The computer system of claim 104, wherein the document is a semi-structured document.

106. The computer system of claim 104, wherein the instructions for generating the hierarchical semantic model further include:

instructions for identifying one or more candidate chunks in the document, each candidate chunk corresponding to a respective subset of the document.

107. The computer system of claim 106, wherein the instructions for applying the heuristics further include:

instructions for identifying a subset of the document as a candidate chunk if the subset of the document has at least one instance of predefined metadata.

5 108. The computer system of claim 106, wherein the instructions for applying the heuristics further include:

instructions for identifying a subset of the document as a candidate chunk if the subset of the document has at least two instances of predefined metadata.

10 109. The computer system of claim 106, wherein a first subset of the document associated with a first candidate chunk encompasses a second subset of the document associated with a second candidate chunk.

110. The computer system of claim 109, wherein the search keywords include a first search keyword and a second search keyword and the instructions for scanning the hierarchical semantic model further include:

15 instructions for identifying content data in the first candidate chunk and preceding the second candidate chunk that satisfies the first search keyword;

instructions for identifying content data in the second candidate chunk that satisfies the second search keyword; and

instructions for choosing the first candidate chunk as the identified chunk.

20 111. The computer system of claim 109, further comprising:

instructions for identifying content data in the second candidate chunk that satisfies the first search keyword and the second keyword; and

instructions for choosing the second candidate chunk as the identified chunk.

25 112. The computer system of claim 106, wherein a first subset of the document associated with a first candidate chunk encompasses a second subset of the document associated with a second candidate chunk and a third subset of the document associated with a third candidate chunk and wherein there is no overlapping between the second candidate chunk and the third candidate chunk.

113. The computer system of claim 112, wherein the search keywords include a first search keyword and a second search keyword and the instructions for scanning the hierarchical semantic model further include:

instructions for identifying content data satisfying the first search keyword in the
5 second candidate chunk;

instructions for identifying content data satisfying the second search keyword in the
third candidate chunk; and

instructions for choosing the first candidate chunk as the identified chunk.

114. The computer system of claim 112, wherein there is no intermediate candidate chunk
10 within the first candidate chunk that encompasses the second candidate chunk and the third candidate chunk.

115. The computer system of claim 104, further comprising:

instructions for identifying the chunk within the document by sequentially scanning
the hierarchical semantic model.

116. A computer readable storage medium having stored therein instructions, which when
15 executed by a computer system cause the computer system to:

identify a document in response to a search request from a user, wherein the
document includes content data and metadata and the search request includes one or more
search keywords;

20 generate a hierarchical semantic model of the content data of the document by
applying heuristics to the metadata of the document;

identify a chunk within the document by scanning the hierarchical semantic model,
wherein the identified chunk includes a subset of the content data that satisfies the search
keywords and the corresponding metadata; and

25 return the identified chunk to the requesting user.

117. The computer readable storage medium of claim 116, wherein the document is a
semi-structured document.

118. The computer readable storage medium of claim 116, wherein the instructions for
generating the hierarchical semantic model further include:

30 instructions for identifying one or more candidate chunks in the document, each
candidate chunk corresponding to a respective subset of the document.

119. The computer readable storage medium of claim 118, wherein the instructions for applying the heuristics further include:

instructions for identifying a subset of the document as a candidate chunk if the subset of the document has at least one instance of predefined metadata.

5 120. The computer readable storage medium of claim 118, wherein the instructions for applying the heuristics further include:

instructions for identifying a subset of the document as a candidate chunk if the subset of the document has at least two instances of predefined metadata.

10 121. The computer readable storage medium of claim 118, wherein a first subset of the document associated with a first candidate chunk encompasses a second subset of the document associated with a second candidate chunk.

122. The computer readable storage medium of claim 121, wherein the search keywords include a first search keyword and a second search keyword and the instructions for scanning the hierarchical semantic model further include:

15 instructions for identifying content data in the first candidate chunk and preceding the second candidate chunk that satisfies the first search keyword;

instructions for identifying content data in the second candidate chunk that satisfies the second search keyword; and

instructions for choosing the first candidate chunk as the identified chunk.

20 123. The computer readable storage medium of claim 121, further comprising:

instructions for identifying content data in the second candidate chunk that satisfies the first search keyword and the second keyword; and

instructions for choosing the second candidate chunk as the identified chunk.

25 124. The computer readable storage medium of claim 118, wherein a first subset of the document associated with a first candidate chunk encompasses a second subset of the document associated with a second candidate chunk and a third subset of the document associated with a third candidate chunk and wherein there is no overlapping between the second candidate chunk and the third candidate chunk.

125. The computer readable storage medium of claim 124, wherein the search keywords include a first search keyword and a second search keyword and the instructions for scanning the hierarchical semantic model further include:

instructions for identifying content data satisfying the first search keyword in the
5 second candidate chunk;

instructions for identifying content data satisfying the second search keyword in the
third candidate chunk; and

instructions for choosing the first candidate chunk as the identified chunk.

126. The computer readable storage medium of claim 124, wherein there is no intermediate
10 candidate chunk within the first candidate chunk that encompasses the second candidate chunk and the third candidate chunk.

127. The computer readable storage medium of claim 116, further comprising:

instructions for identifying the chunk within the document by sequentially scanning
the hierarchical semantic model.

128. A computer-implemented method, comprising:

receiving a search keyword provided by a user;

selecting an archetype for the search keyword;

identifying one or more search results in accordance with the archetype; and

returning at least one of the search results to the user.

129. The method of claim 128, wherein the archetype has an enumerable set of instances
20 and the search keyword is associated with one of the instances.

130. The method of claim 128, further comprising:

after selecting the archetype,

identifying at least one query operator for the selected archetype;

constructing a search query using the query operator; and

executing the search query against one or more data sources.

131. The method of claim 130, wherein the query operator has a scope and the search
query is constructed to limit search results within the scope.

132. The method of claim 131, wherein the scope is from 0% to 100%.

133. The method of claim 130, wherein the query operator has a pattern and the search query is constructed to limit search results including the pattern.

134. The method of claim 133, wherein the pattern is a date pattern.

135. The method of claim 128, further comprising:

5 after selecting the archetype and before identifying the search results,
 soliciting user instructions in connection with the archetype;
 constructing the search query in accordance with the user instructions; and
 executing the search query against the data sources.

136. The method of claim 135, further comprising:

10 after selecting the archetype and before identifying the search results,
 generating feedback to the user instructions;
 receiving more user instructions in connection with the archetype and the
feedback; and
 repeating said generating feedback and said receiving more user instructions
15 until receiving a search query execution request from the user.

137. A computer system, comprising:

 memory;
 one or more processors;
 one or more programs stored in the memory and configured for execution by the one
20 or more processors, the one or more programs including:
 instructions for receiving a search keyword provided by a user;
 instructions for selecting an archetype for the search keyword;
 instructions for identifying one or more search results in accordance with the
archetype; and
25 instructions for returning at least one of the search results to the user.

138. The computer system of claim 137, wherein the archetype has an enumerable set of instances and the search keyword is associated with one of the instances.

139. The computer system of claim 137, further comprising:

 instructions for identifying at least one query operator for the selected archetype;
30 instructions for constructing a search query using the query operator; and

instructions for executing the search query against one or more data sources.

140. The computer system of claim 139, wherein the query operator has a scope and the search query is constructed to limit search results within the scope.

141. The computer system of claim 140, wherein the scope is from 0% to 100%.

5 142. The computer system of claim 139, wherein the query operator has a pattern and the search query is constructed to limit search results including the pattern.

143. The computer system of claim 142, wherein the pattern is a date pattern.

144. The computer system of claim 137, further comprising:

instructions for soliciting user instructions in connection with the selected archetype;

10 instructions for constructing the search query in accordance with the user instructions;

and

instructions for executing the search query against the data sources.

145. The computer system of claim 144, further comprising:

instructions for generating feedback to the user instructions;

15 instructions for receiving more user instructions in connection with the archetype and the feedback; and

instructions for repeating said generating feedback and said receiving more user instructions until receiving a search query execution request from the user.

146. A computer readable storage medium having stored therein instructions, which when

20 executed by a computer system cause the computer system to:

receive a search keyword provided by a user;

select an archetype for the search keyword;

identify one or more search results in accordance with the archetype; and

return at least one of the search results to the user.

25 147. The computer readable storage medium of claim 146, wherein the archetype has an enumerable set of instances and the search keyword is associated with one of the instances.

148. The computer readable storage medium of claim 146, further comprising:

instructions for identifying at least one query operator for the selected archetype;

instructions for constructing a search query using the query operator; and
instructions for executing the search query against one or more data sources.

149. The computer readable storage medium of claim 148, wherein the query operator has a scope and the search query is constructed to limit search results within the scope.

5 150. The computer readable storage medium of claim 149, wherein the scope is from 0% to 100%.

151. The computer system of claim 139, wherein the query operator has a pattern and the search query is constructed to limit search results including the pattern.

10 152. The computer readable storage medium of claim 151, wherein the pattern is a date pattern.

153. The computer readable storage medium of claim 146, further comprising:
instructions for soliciting user instructions in connection with the selected archetype;
instructions for constructing the search query in accordance with the user instructions;
and
15 instructions for executing the search query against the data sources.

154. The computer readable storage medium of claim 153, further comprising:
instructions for generating feedback to the user instructions;
instructions for receiving more user instructions in connection with the archetype and
the feedback; and
20 instructions for repeating said generating feedback and said receiving more user
instructions until receiving a search query execution request from the user.

155. A computer-implemented method, comprising:
displaying an application user interface, the application user interface including a
document authoring window and a search results window;
25 displaying in the search results window a set of search results associated with one or
more user-specified search keywords in a text-only display format, wherein each search result
includes a chunk within a respective document that satisfies the search keywords;
in response to a user request to view a chunk, launching a document display window
in the application user interface and displaying therein a portion of the corresponding
30 document that includes the chunk in its native display format; and

in response to a user request to duplicate a segment of the corresponding document in the document authoring window, generating therein an instance of the segment of the corresponding document in its native display format.

156. The method of claim 155, wherein the portion of the corresponding document in the document display window includes more information about the search keywords than the chunk in the search results window.

157. The method of claim 156, wherein the more information includes a search keyword's location in the corresponding document.

158. The method of claim 156, wherein the more information includes textual contents adjacent the search keywords in the corresponding document.

159. The method of claim 155, wherein the set of search results includes a first chunk within a first document having a first content type and a second chunk within a second document having a second content type, wherein the first content type is different from the second content type.

160. The method of claim 159, further comprising:
in response to a user request to duplicate the first chunk from the search results window into the document authoring window, generating therein an instance of a first segment of the first document in its native display format, wherein the first segment includes the first chunk; and

in response to a user request to duplicate the second chunk from the search results window into the document authoring window, generating therein an instance of a second segment of the second document in its native display format, wherein the second segment includes the second chunk;

wherein the first document and the second document have different native display formats.

161. The method of claim 159, further comprising:

in response to a first user selection of the first chunk, launching a first document display window in the application user interface and displaying therein a first document that includes the first chunk in its native display format; and

in response to a second user selection of the second chunk, launching a second document display window in the application user interface and displaying therein a second document that includes the second chunk in its native display format;

wherein the first document and the second document have different native display formats.

162. The method of claim 161, further comprising:

in response to a user request to duplicate a segment of the first document from the first document display window into the document authoring window, generating therein an instance of the segment of the first document in its native display format; and

in response to a user request to duplicate a segment of the second document the second document display window into the document authoring window, generating therein an instance of the segment of the second document in its native display format.

163. The method of claim 161, further comprising:

closing the first document display window before launching the second document display window in response to the second user selection of the second chunk.

164. The method of claim 155, further comprising:

in response to the user request to view the chunk,
generating an empty region in the application user interface by shrinking the document authoring window; and

occupying the empty region with the document display window in the application user interface.

165. The method of claim 155, wherein the document display window is a preview-only window of the corresponding document.

166. The method of claim 155, wherein the document display window is a second document authoring window that is different from the original document authoring window in the application user interface.

167. The method of claim 155, wherein different search keywords are highlighted in different manners in the search results window.

168. The method of claim 155, wherein the portion of the corresponding document in the document display window includes at least one highlighted search keyword.

169. A computer system, comprising:

memory;

one or more processors;

one or more programs stored in the memory and configured for execution by the one

5 or more processors, the one or more programs including:

instructions for displaying an application user interface, the application user interface including a document authoring window and a search results window;

instructions for displaying in the search results window a set of search results associated with one or more user-specified search keywords in a text-only display format,

10 wherein each search result includes a chunk within a respective document that satisfies the search keywords;

instructions for, in response to a user request to view a chunk, launching a document display window in the application user interface and displaying therein a portion of the corresponding document that includes the chunk in its native display format; and

15 instructions for, in response to a user request to duplicate a segment of the corresponding document in the document authoring window, generating therein an instance of the segment of the corresponding document in its native display format.

170. The computer system of claim 169, wherein the portion of the corresponding document in the document display window includes more information about the search
20 keywords than the chunk in the search results window.

171. The computer system of claim 170, wherein the more information includes a search keyword's location in the corresponding document.

172. The computer system of claim 170, wherein the more information includes textual contents adjacent the search keywords in the corresponding document.

25 173. The computer system of claim 169, wherein the set of search results includes a first chunk within a first document having a first content type and a second chunk within a second document having a second content type, wherein the first content type is different from the second content type.

174. The computer system of claim 173, further comprising:

30 instructions for, in response to a user request to duplicate the first chunk from the search results window into the document authoring window, generating therein an instance of

a first segment of the first document in its native display format, wherein the first segment includes the first chunk; and

instructions for, in response to a user request to duplicate the second chunk from the search results window into the document authoring window, generating therein an instance of a second segment of the second document in its native display format, wherein the second segment includes the second chunk;

wherein the first document and the second document have different native display formats.

175. The computer system of claim 173, further comprising:

instructions for, in response to a first user selection of the first chunk, launching a first document display window in the application user interface and displaying therein a first document that includes the first chunk in its native display format; and

instructions for, in response to a second user selection of the second chunk, launching a second document display window in the application user interface and displaying therein a second document that includes the second chunk in its native display format;

wherein the first document and the second document have different native display formats.

176. The computer system of claim 175, further comprising:

instructions for, in response to a user request to duplicate a segment of the first document from the first document display window into the document authoring window, generating therein an instance of the segment of the first document in its native display format; and

instructions for, in response to a user request to duplicate a segment of the second document the second document display window into the document authoring window, generating therein an instance of the segment of the second document in its native display format.

177. The computer system of claim 175, further comprising:

instructions for closing the first document display window before launching the second document display window in response to the second user selection of the second chunk.

178. The computer system of claim 169, further comprising:

instructions for generating an empty region in the application user interface by shrinking the document authoring window in response to the user request to view the chunk; and

instructions for occupying the empty region with the document display window in the application user interface.

179. The computer system of claim 169, wherein the document display window is a preview-only window of the corresponding document.

180. The computer system of claim 169, wherein the document display window is a second document authoring window that is different from the original document authoring window in the application user interface.

181. The computer system of claim 169, wherein different search keywords are highlighted in different manners in the search results window.

182. The computer system of claim 169, wherein the portion of the corresponding document in the document display window includes at least one highlighted search keyword.

183. A computer readable storage medium having stored therein instructions, which when executed by a computer system cause the computer system to:

display an application user interface, the application user interface including a document authoring window and a search results window;

display in the search results window a set of search results associated with one or more user-specified search keywords in a text-only display format, wherein each search result includes a chunk within a respective document that satisfies the search keywords;

in response to a user request to view a chunk, launch a document display window in the application user interface and display therein a portion of the corresponding document that includes the chunk in its native display format; and

in response to a user request to duplicate a segment of the corresponding document in the document authoring window, generate therein an instance of the segment of the corresponding document in its native display format.

184. The computer readable storage medium of claim 183, wherein the portion of the corresponding document in the document display window includes more information about the search keywords than the chunk in the search results window.

185. The computer readable storage medium of claim 184, wherein the more information includes a search keyword's location in the corresponding document.

186. The computer readable storage medium of claim 184, wherein the more information includes textual contents adjacent the search keywords in the corresponding document.

5 187. The computer readable storage medium of claim 183, wherein the set of search results includes a first chunk within a first document having a first content type and a second chunk within a second document having a second content type, wherein the first content type is different from the second content type.

188. The computer readable storage medium of claim 187, further comprising:

10 instructions for, in response to a user request to duplicate the first chunk from the search results window into the document authoring window, generating therein an instance of a first segment of the first document in its native display format, wherein the first segment includes the first chunk; and

15 instructions for, in response to a user request to duplicate the second chunk from the search results window into the document authoring window, generating therein an instance of a second segment of the second document in its native display format, wherein the second segment includes the second chunk;

wherein the first document and the second document have different native display formats.

20 189. The computer readable storage medium of claim 187, further comprising:

instructions for, in response to a first user selection of the first chunk, launching a first document display window in the application user interface and displaying therein a first document that includes the first chunk in its native display format; and

25 instructions for, in response to a second user selection of the second chunk, launching a second document display window in the application user interface and displaying therein a second document that includes the second chunk in its native display format;

wherein the first document and the second document have different native display formats.

190. The computer readable storage medium of claim 189, further comprising:

30 instructions for, in response to a user request to duplicate a segment of the first document from the first document display window into the document authoring window,

generating therein an instance of the segment of the first document in its native display format; and

instructions for, in response to a user request to duplicate a segment of the second document the second document display window into the document authoring window,
5 generating therein an instance of the segment of the second document in its native display format.

191. The computer readable storage medium of claim 189, further comprising:
instructions for closing the first document display window before launching the
second document display window in response to the second user selection of the second
10 chunk.

192. The computer readable storage medium of claim 183, further comprising:
instructions for generating an empty region in the application user interface by
shrinking the document authoring window in response to the user request to view the chunk;
and

15 instructions for occupying the empty region with the document display window in the application user interface.

193. The computer readable storage medium of claim 183, wherein the document display window is a preview-only window of the corresponding document.

194. The computer readable storage medium of claim 183, wherein the document display
20 window is a second document authoring window that is different from the original document authoring window in the application user interface.

195. The computer readable storage medium of claim 183, wherein different search keywords are highlighted in different manners in the search results window.

196. The computer readable storage medium of claim 183, wherein the portion of the
25 corresponding document in the document display window includes at least one highlighted search keyword.

197. A graphical user interface on a computer display, comprising:
an application user interface, the application user interface including a document
authoring window and a search results window, wherein a set of search results associated
30 with one or more user-specified search keywords is displayed in the search results window in

a text-only display format, wherein each search result includes one or more chunks identified within a respective document as satisfying the user-specified search keywords, wherein:

in response to a user request to duplicate a chunk within a document in the document authoring window, an instance of the chunk is displayed in the document authoring window
5 in the document's native display format.

198. The graphical user interface of claim 197, wherein two chunks identified within two different documents have different native display formats.

199. The graphical user interface of claim 197, wherein each chunk has an associated chunk link, wherein:

10 in response to a user selection of a respective chunk link, a document display window is displayed in the application user interface and a portion of the corresponding document that includes the corresponding chunk is displayed in the document display window in the document's native display format.

200. The graphical user interface of claim 197, wherein each chunk includes terms that
15 match the user-specified search keywords an associated chunk link, and different terms matching different search keywords are highlighted in the search results window in a visually distinguishable manner.

201. The graphical user interface of claim 197, wherein the one or more chunks identified within a document are displayed in the search results window in an order consistent with their
20 relative relevancy to the user-specified search keywords.

202. The graphical user interface of claim 197, wherein the one or more chunks identified within a document are displayed in the search results window in an order consistent with their relative locations within the corresponding document.

203. A computer-implemented method, comprising:

25 receiving a search request that includes one or more user-specified search keywords;
identifying a first document and a second document, each document having at least one chunk that satisfies the search keywords, wherein a first text string in the first document has a first content type and the first text string in the second document has a second content type that is different from the first content type;

30 receiving a user request to replace the first text string with a second text string; and

substituting the second text string for the first text string in the first document and the second document, wherein the replacing second text string in the first document has the first content type and the replacing second text string in the second document has the second content type.

5 204. The method of claim 203, further comprising:

 after identifying the first document and the second document,

 displaying a first chunk from the first document and a second chunk from the second document, each chunk including at least one instance of the first text string, wherein the instances of the first text string within the first and second chunks are displayed in a text-only display format; and

10

 after substituting the second text string for the first text string,

 replacing the displayed instances of the first text string within the first and second chunks with respective instances of the second text string, wherein the instances of the replacing second text string within the first and second chunks are displayed in the text-only display format.

15

205. The method of claim 203, wherein the first document has a first document type and the second document has a second document type that is different from the first document type.

206. The method of claim 203, wherein the first text string in the first document has a first appearance when the first document is rendered by its native application and the first text string in the second document has a second appearance that is different from the first appearance when the second document is rendered by its native application.

20

207. The method of claim 203, wherein the first document includes an original second text string that has a content type different from the replacing second text string.

25 208. The method of claim 207, wherein, when the first document is rendered by its native application, the original second text string and the replacing second text string have different appearances.

209. The method of claim 203, wherein the user request stipulates that the first text string at one or more user-specified locations in the first document and the second document be replaced by the second text string, further comprising:

30

substituting the second text string for the first text string at the user-specified locations in the first document and the second document, respectively.

210. The method of claim 209, wherein the user-specified locations include one or more selected from the group consisting of title, paragraph, table, header, footer, slide, spreadsheet, and all.

211. A computer-implemented method, comprising:

receiving a user request to replace a first text string with a second text string in a first document and a second document, respectively, wherein the first text string in the first document has a first content type and the first text string in the second document has a second content type that is different from the first content type; and

substituting the second text string for the first text string in the first document and the second document, wherein the replacing second text string in the first document has the first content type and the replacing second text string in the second document has the second content type.

212. The method of claim 211, wherein the user request stipulates that the first text string at user-specified locations in the first document and the second document be replaced by the second text string.

213. The method of claim 212, wherein the user-specified locations include one or more selected from the group consisting of title, paragraph, table, header, footer, slide, spreadsheet, and all.

214. The method of claim 212, further comprising:

substituting the second text string for the first text string at the user-specified locations in the first document and the second document, respectively.

215. The method of claim 211, wherein the first document and the second document are identified in response to a search request for the first text string within a set of documents.

216. The method of claim 215, further comprising:

after identifying the first document and the second document,
displaying a first chunk from the first document and a second chunk from the second document, each chunk including at least one instance of the first text string, wherein

the instances of the first text string within the first and second chunks are displayed in a text-only display format; and

after substituting the second text string for the first text string,

replacing the displayed instances of the first text string within the first and

5 second chunks with respective instances of the second text string, wherein the instances of the replacing second text string within the first and second chunks are displayed in the text-only display format.

217. The method of claim 211, wherein the first document has a first document type and the second document has a second document type that is different from the first document
10 type.

218. The method of claim 211, wherein the first text string in the first document has a first appearance when the first document is rendered by its native application and the first text string in the second document has a second appearance that is different from the first appearance when the second document is rendered by its native application.

15 219. The method of claim 211, wherein the first document includes an original second text string that has a content type different from the replacing second text string.

220. The method of claim 219, wherein the original second text string and the replacing second text string have different appearances when the first document is rendered by its native application.

20 221. A computer system, comprising:

memory;

one or more processors;

one or more programs stored in the memory and configured for execution by the one or more processors, the one or more programs including:

25 instructions for receiving a user request to replace a first text string with a second text string in a first document and a second document, respectively, wherein the first text string in the first document has a first content type and the first text string in the second document has a second content type that is different from the first content type; and

30 instructions for substituting the second text string for the first text string in the first document and the second document, wherein the replacing second text string in the first

document has the first content type and the replacing second text string in the second document has the second content type.

222. The computer system of claim 221, wherein the user request stipulates that the first text string at user-specified locations in the first document and the second document be replaced by the second text string.

223. The computer system of claim 222, wherein the user-specified locations include one or more selected from the group consisting of title, paragraph, table, header, footer, slide, spreadsheet, and all.

224. The computer system of claim 222, further comprising:

instructions for substituting the second text string for the first text string at the user-specified locations in the first document and the second document, respectively.

225. The computer system of claim 221, wherein the first document and the second document are identified in response to a search request for the first text string within a set of documents.

226. The computer system of claim 225, further comprising:

instructions for displaying a first chunk from the first document and a second chunk from the second document, each chunk including at least one instance of the first text string, wherein the instances of the first text string within the first and second chunks are displayed in a text-only display format; and

instructions for replacing the displayed instances of the first text string within the first and second chunks with respective instances of the second text string, wherein the instances of the replacing second text string within the first and second chunks are displayed in the text-only display format.

227. The computer system of claim 221, wherein the first document has a first document type and the second document has a second document type that is different from the first document type.

228. The computer system of claim 221, wherein the first text string in the first document has a first appearance when the first document is rendered by its native application and the first text string in the second document has a second appearance that is different from the first appearance when the second document is rendered by its native application.

229. The computer system of claim 221, wherein the first document includes an original second text string that has a content type different from the replacing second text string.

230. The computer system of claim 229, wherein the original second text string and the replacing second text string have different appearances when the first document is rendered
5 by its native application.

231. A computer readable storage medium having stored therein instructions, which when executed by a computer system cause the computer system to:

receive a user request to replace a first text string with a second text string in a first document and a second document, respectively, wherein the first text string in the first
10 document has a first content type and the first text string in the second document has a second content type that is different from the first content type; and

substitute the second text string for the first text string in the first document and the second document, wherein the replacing second text string in the first document has the first content type and the replacing second text string in the second document has the second
15 content type.

232. The computer readable storage medium of claim 231, wherein the user request stipulates that the first text string at user-specified locations in the first document and the second document be replaced by the second text string.

233. The computer readable storage medium of claim 232, wherein the user-specified
20 locations include one or more selected from the group consisting of title, paragraph, table, header, footer, slide, spreadsheet, and all.

234. The computer readable storage medium of claim 232, further comprising:
instructions for substituting the second text string for the first text string at the user-specified locations in the first document and the second document, respectively.

235. The computer readable storage medium of claim 231, wherein the first document and the second document are identified in response to a search request for the first text string within a set of documents.

236. The computer readable storage medium of claim 235, further comprising:
instructions for displaying a first chunk from the first document and a second chunk
30 from the second document, each chunk including at least one instance of the first text string,

wherein the instances of the first text string within the first and second chunks are displayed in a text-only display format; and

instructions for replacing the displayed instances of the first text string within the first and second chunks with respective instances of the second text string, wherein the instances of the replacing second text string within the first and second chunks are displayed in the text-only display format.

237. The computer readable storage medium of claim 231, wherein the first document has a first document type and the second document has a second document type that is different from the first document type.

238. The computer readable storage medium of claim 231, wherein the first text string in the first document has a first appearance when the first document is rendered by its native application and the first text string in the second document has a second appearance that is different from the first appearance when the second document is rendered by its native application.

239. The computer readable storage medium of claim 231, wherein the first document includes an original second text string that has a content type different from the replacing second text string.

240. The computer readable storage medium of claim 239, wherein the original second text string and the replacing second text string have different appearances when the first document is rendered by its native application.

241. A computer-implemented method, comprising:

receiving a first user request including a first set of search keywords;

identifying a first set of chunks within multiple documents, wherein each chunk includes terms matching the first set of search keywords;

displaying at least a portion of the first set of chunks, including highlighting the terms matching the first set of search keywords in the displayed portion in a first manner;

receiving a second user request to search among the documents for documents that satisfy a second set of search keywords;

identifying a second set of chunks within the documents, wherein each chunk includes terms matching the second set of search keywords; and

displaying at least a portion of the second set of chunks, including highlighting the terms matching the second set of search keywords in the displayed portion in a second manner that is different from the first manner.

242. The method of claim 241, wherein the second set of chunks includes at least one chunk that is included in the first set of chunks.

243. The method of claim 242, further comprising:

displaying the at least one chunk in an order consistent with its relevancy to the first set of search keywords in the first set of chunks; and

displaying the at least one chunk in an order consistent with its relevancy to the second set of search keywords in the second set of chunks.

244. The method of claim 241, wherein the second set of chunks includes at least one chunk that is not included in the first set of chunks.

245. The method of claim 241, wherein the second set of search keywords is different from the first set of search keywords.

246. A computer-implemented method, comprising:

receiving a first user request including a first set of search keywords;

identifying multiple documents, wherein each document includes at least one chunk that satisfies the first set of search keywords;

receiving a second user request to search among the chunks in the identified documents for chunks that satisfy a second set of search keywords; and

identifying a subset of the chunks, wherein each chunk in the subset satisfies the second set of search keywords.

247. The method of claim 246, further comprising:

in response to the first user request, displaying the chunks in the identified documents in accordance with their relevancy to the first set of search keywords; and

in response to the second user request, displaying the subset of chunks in accordance with their relevancy to the second set of search keywords.

248. The method of claim 247, further comprising:

highlighting terms in the subset of chunks matching the first set of search keywords, if any, in a first manner; and

highlighting terms in the subset of chunks matching the second set of search keywords in a second manner that is different from the first manner.

249. The method of claim 246, wherein the subset of chunks includes a first chunk and a second chunk, further comprising:

- 5 in response to the first user request, displaying the first chunk ahead of the second chunk; and
 in response to the second user request, displaying the second chunk ahead of the first chunk.

10 250. The method of claim 246, wherein the second set of search keywords is different from the first set of search keywords.

251. A computer system, comprising:

- memory;
one or more processors;
one or more programs stored in the memory and configured for execution by the one
15 or more processors, the one or more programs including:
 instructions for receiving a first user request including a first set of search keywords;
 instructions for identifying a first set of chunks within multiple documents, wherein each chunk includes terms matching the first set of search keywords;
20 instructions for displaying at least a portion of the first set of chunks, including highlighting the terms matching the first set of search keywords in the displayed portion in a first manner;
 instructions for receiving a second user request to search among the documents for documents that satisfy a second set of search keywords;
25 instructions for identifying a second set of chunks within the documents, wherein each chunk includes terms matching the second set of search keywords; and
 instructions for displaying at least a portion of the second set of chunks, including highlighting the terms matching the second set of search keywords in the displayed portion in a second manner that is different from the first manner.

30 252. The computer system of claim 251, wherein the second set of chunks includes at least one chunk that is included in the first set of chunks.

253. The computer system of claim 252, further comprising:

instructions for displaying the at least one chunk in an order consistent with its relevancy to the first set of search keywords in the first set of chunks; and

instructions for displaying the at least one chunk in an order consistent with its relevancy to the second set of search keywords in the second set of chunks.

254. The computer system of claim 251, further comprising:

instructions for highlighting terms in the second set of chunks matching the first set of search keywords, if any, in a first manner; and

instructions for highlighting terms in the second set of chunks matching the second set of search keywords in a second manner that is different from the first manner.

255. The computer system of claim 251, wherein both the first set of chunks and the second set of chunks include a first chunk and a second chunk, further comprising:

instructions for displaying the first chunk ahead of the second chunk in response to the first user request; and

instructions for displaying the second chunk ahead of the first chunk in response to the second user request.

256. The computer system of claim 251, wherein the second set of chunks includes at least one chunk that is not included in the first set of chunks.

257. The computer system of claim 251, wherein the second set of search keywords is different from the first set of search keywords.

258. A computer readable storage medium having stored therein instructions, which when executed by a computer system cause the computer system to:

receive a first user request including a first set of search keywords;

identify a first set of chunks within multiple documents, wherein each chunk includes terms matching the first set of search keywords;

display at least a portion of the first set of chunks, including highlighting the terms matching the first set of search keywords in the displayed portion in a first manner;

receive a second user request to search among the documents for documents that satisfy a second set of search keywords;

identify a second set of chunks within the documents, wherein each chunk includes terms matching the second set of search keywords; and

display at least a portion of the second set of chunks, including highlighting the terms matching the second set of search keywords in the displayed portion in a second manner that is different from the first manner.

259. The computer readable storage medium of claim 258, wherein the second set of
5 chunks includes at least one chunk that is included in the first set of chunks.

260. The computer readable storage medium of claim 259, further comprising:
instructions for displaying the at least one chunk in an order consistent with its
relevancy to the first set of search keywords in the first set of chunks; and
instructions for displaying the at least one chunk in an order consistent with its
10 relevancy to the second set of search keywords in the second set of chunks.

261. The computer readable storage medium of claim 258, further comprising:
instructions for highlighting terms in the second set of chunks matching the first set of
search keywords, if any, in a first manner; and
instructions for highlighting terms in the second set of chunks matching the second set
15 of search keywords in a second manner that is different from the first manner.

262. The computer readable storage medium of claim 258, wherein both the first set of
chunks and the second set of chunks include a first chunk and a second chunk, further
comprising:
instructions for displaying the first chunk ahead of the second chunk in response to
20 the first user request; and
instructions for displaying the second chunk ahead of the first chunk in response to
the second user request.

263. The computer readable storage medium of claim 258, wherein the second set of
chunks includes at least one chunk that is not included in the first set of chunks.

25 264. The computer readable storage medium of claim 258, wherein the second set of
search keywords is different from the first set of search keywords.

265. A graphical user interface on a computer display, comprising:
a first set of search results displayed in a text-only display format, wherein each
search result includes one or more chunks identified within a respective document as
30 satisfying a first set of search keywords, wherein:

in response to a user request to search among the identified chunks for chunks that satisfy a second set of search keywords, the first set of search results is replaced by a second set of search results, wherein each search result in the second set includes one or more chunks identified within a respective document as satisfying both the first set of search keywords and the second set of search keywords.

266. The graphical user interface of claim 265, wherein two chunks identified within two different documents have different native display formats.

267. The graphical user interface of claim 265, wherein the second set of search keywords is different from the first set of search keywords.

268. The graphical user interface of claim 265, wherein terms matching the first set of search keywords and terms matching the second set of search keywords within a respective chunk are highlighted in a visually distinguishable manner.

269. The graphical user interface of claim 265, wherein:

the one or more chunks identified within a respective document as satisfying the first set of search keywords are displayed in an order consistent with their relative relevancy to the first set of search keywords; and

the one or more chunks identified within a respective document as satisfying both the first set of search keywords and the second set of search keywords are displayed in an order consistent with their relative relevancy to the second set of search keywords.

270. The graphical user interface of claim 265, wherein the one or more chunks identified within a respective document as satisfying any of the first and second sets of search keywords are displayed in an order consistent with their relative locations within the corresponding document.

271. A computer-implemented method, comprising:

identifying a first candidate document at a first data source and a second candidate document at a second data source in response to a request from a user, wherein the request includes one or more keywords;

generating a first node stream for the first candidate document and a second node stream for the second candidate document using data packets received from the respective first and second data sources; and

alternatively processing the first node stream and the second node stream until a candidate chunk matching at least one of the keywords is identified therein, wherein the candidate chunk includes a set of nodes within a respective data source.

272. The method of claim 271, further comprising:

5 returning the candidate chunk as a relevant chunk to the user if the candidate chunk satisfies the keywords.

273. The method of claim 271, further comprising:

submitting a document-retrieval request to the first data source; and

10 receiving a response from the first data source, the response including multiple data packets corresponding to the first candidate document.

274. The method of claim 273, wherein the document-retrieval request and the response are based on a communication protocol selected from the group consisting of TCP/IP and FTP.

275. The method of claim 273, further comprising:

15 receiving one of the data packets from the first data source;
extracting one or more nodes from the data packet; and
inserting the one or more nodes into the first node stream.

276. The method of claim 275, further comprising:

20 extracting a node fragment from the data packet;
forming a node by combining the node fragment with another node fragment; and
inserting the formed node into the first node stream.

277. The method of claim 271, wherein the alternative processing of the first node stream and the second node stream further includes:

25 after processing nodes currently in the first node stream,
processing nodes currently in the second node stream if no new node appears in the first node stream for a first amount of time; and
identifying the candidate chunk in the second node stream.

278. The method of claim 277, wherein the alternative processing of the first node stream and the second node stream further includes:

30 after processing nodes currently in the second node stream,

processing nodes currently in the first node stream if no new node appears in the second node stream for the first amount of time; and
identifying the candidate chunk in the first node stream.

279. The method of claim 277, wherein the alternative processing of the first node stream and the second node stream further includes:

discarding processing results associated with the first node stream if no new node appears in the first node stream for a second amount of time, wherein the second amount of time is no less than the first amount of time.

280. The method of claim 271, wherein the first data source is different from the second data source.

281. The method of claim 271, wherein the first data source and the second data source are two different web servers.

282. The method of claim 271, wherein the first candidate document is an HTML web page.

283. A computer system, comprising:

memory;

one or more processors;

one or more programs stored in the memory and configured for execution by the one or more processors, the one or more programs including:

instructions for identifying a first candidate document at a first data source and a second candidate document at a second data source in response to a request from a user, wherein the request includes one or more keywords;

instructions for generating a first node stream for the first candidate document and a second node stream for the second candidate document using data packets received from the respective first and second data sources; and

instructions for alternatively processing the first node stream and the second node stream until a candidate chunk matching at least one of the keywords is identified therein, wherein the candidate chunk includes a set of nodes within a respective data source.

284. The computer system of claim 283, further comprising:

instructions for returning the candidate chunk as a relevant chunk to the user if the candidate chunk satisfies the keywords.

285. The computer system of claim 283, further comprising:

instructions for submitting a document-retrieval request to the first data source; and

instructions for receiving a response from the first data source, the response including multiple data packets corresponding to the first candidate document.

286. The computer system of claim 285, wherein the document-retrieval request and the response are based on a communication protocol selected from the group consisting of TCP/IP and FTP.

287. The computer system of claim 285, further comprising:

instructions for receiving one of the data packets from the first data source;

instructions for extracting one or more nodes from the data packet; and

instructions for inserting the one or more nodes into the first node stream.

288. The computer system of claim 287, further comprising:

instructions for extracting a node fragment from the data packet;

instructions for forming a node by combining the node fragment with another node fragment; and

instructions for inserting the formed node into the first node stream.

289. The computer system of claim 283, wherein the instructions for alternatively processing the first node stream and the second node stream further include:

instructions for processing nodes currently in the second node stream if no new node appears in the first node stream for a first amount of time; and

instructions for identifying the candidate chunk in the second node stream.

290. The computer system of claim 289, wherein the instructions for alternatively processing the first node stream and the second node stream further include:

instructions for discarding processing results associated with the first node stream if no new node appears in the first node stream for a second amount of time, wherein the second amount of time is no less than the first amount of time.

291. The computer system of claim 283, wherein the first data source is different from the second data source.

292. The computer system of claim 283, wherein the first data source and the second data source are two different web servers.

293. The computer system of claim 283, wherein the first candidate document is an HTML web page.

5 294. A computer readable storage medium having stored therein instructions, which when executed by a computer system cause the computer system to:

identify a first candidate document at a first data source and a second candidate document at a second data source in response to a request from a user, wherein the request includes one or more keywords;

10 generate a first node stream for the first candidate document and a second node stream for the second candidate document using data packets received from the respective first and second data sources; and

alternatively process the first node stream and the second node stream until a candidate chunk matching at least one of the keywords is identified therein, wherein the
15 candidate chunk includes a set of nodes within a respective data source.

295. The computer readable storage medium of claim 294, further comprising:
instructions for returning the candidate chunk as a relevant chunk to the user if the candidate chunk satisfies the keywords.

296. The computer readable storage medium of claim 294, further comprising:
20 instructions for submitting a document-retrieval request to the first data source; and
instructions for receiving a response from the first data source, the response including multiple data packets corresponding to the first candidate document.

297. The computer readable storage medium of claim 296, wherein the document-retrieval request and the response are based on a communication protocol selected from the group
25 consisting of TCP/IP and FTP.

298. The computer readable storage medium of claim 296, further comprising:
instructions for receiving one of the data packets from the first data source;
instructions for extracting one or more nodes from the data packet; and
instructions for inserting the one or more nodes into the first node stream.

30 299. The computer readable storage medium of claim 298, further comprising:

instructions for extracting a node fragment from the data packet;
instructions for forming a node by combining the node fragment with another node fragment; and
instructions for inserting the formed node into the first node stream.

- 5 300. The computer readable storage medium of claim 294, wherein the instructions for alternatively processing the first node stream and the second node stream further include:
 instructions for processing nodes currently in the second node stream if no new node appears in the first node stream for a first amount of time; and
 instructions for identifying the candidate chunk in the second node stream.

- 10 301. The computer readable storage medium of claim 300, wherein the instructions for alternatively processing the first node stream and the second node stream further include:
 instructions for discarding processing results associated with the first node stream if no new node appears in the first node stream for a second amount of time, wherein the second amount of time is no less than the first amount of time.

- 15 302. The computer readable storage medium of claim 294, wherein the first data source is different from the second data source.

303. The computer readable storage medium of claim 294, wherein the first data source and the second data source are two different web servers.

- 20 304. The computer readable storage medium of claim 294, wherein the first candidate document is an HTML web page.

305. A computer-implemented method, comprising:
 retrieving a document from a data source, wherein the document has a structure type;
 generating a customized data model for the document in accordance with its structure type;

- 25 identifying one or more candidate chunks within the customized data model in accordance with a set of heuristic rules associated with the structure type.

306. The method of claim 305, further comprising:
 selecting at least one of the candidate chunks that satisfy one or more search keywords.

307. The method of claim 305, wherein the customized data model is a hierarchical data model.

308. The method of claim 305, wherein the structure type is one selected from the group consisting of structured, semi-structured, and unstructured.

5 309. The method of claim 305, wherein the data source is a web server.

310. The method of claim 305, wherein the document is an HTML web page.

311. The method of claim 310, wherein the HTML web page includes multiple pairs of HTML tags, further comprising:

10 identifying a first subset of the HTML web page between a first pair of HTML tags as a first candidate chunk if the first pair of HTML tags satisfies one of the set of heuristic rules.

312. The method of claim 311, further comprising:

recursively identifying a second subset of the HTML web page within the first subset of the HTML web page between a second pair of HTML tags as a second candidate chunk if the second pair of HTML tags satisfy one of the set of heuristic rules.

15 313. The method of claim 305, wherein the document is a plain-text document.

314. The method of claim 313, wherein generating the customized data model further includes inserting metadata into the data model that separates the plain-text document into multiple candidate chunks.

20 315. The method of claim 314, wherein the metadata in the data model is one or more XML tags and the text following at least one of the XML tags is identified as a candidate chunk.

316. A computer system, comprising:

memory;

one or more processors;

25 one or more programs stored in the memory and configured for execution by the one or more processors, the one or more programs including:

instructions for retrieving a document from a data source, wherein the document has a structure type;

instructions for generating a customized data model for the document in accordance with its structure type; and

instructions for identifying one or more candidate chunks within the customized data model in accordance with a set of heuristic rules associated with the structure type.

317. The computer system of claim 316, further comprising:

instructions for selecting at least one of the candidate chunks that satisfy one or more search keywords.

318. The computer system of claim 316, wherein the customized data model is a hierarchical data model.

319. The computer system of claim 316, wherein the structure type is one selected from the group consisting of structured, semi-structured, and unstructured.

320. The computer system of claim 316, wherein the data source is a web server.

321. The computer system of claim 316, wherein the document is an HTML web page.

322. The computer system of claim 321, wherein the HTML web page includes multiple pairs of HTML tags, further comprising:

instructions for identifying a first subset of the HTML web page between a first pair of HTML tags as a first candidate chunk if the first pair of HTML tags satisfy one of the set of heuristic rules.

323. The computer system of claim 322, further comprising:

instructions for recursively identifying a second subset of the HTML web page within the first subset of the HTML web page between a second pair of HTML tags as a second candidate chunk if the second pair of HTML tags satisfy one of the set of heuristic rules.

324. The computer system of claim 316, wherein the document is a plain-text document.

325. The computer system of claim 324, wherein the instructions for generating the customized data model further include instructions for inserting metadata into the data model that separates the plain-text document into multiple candidate chunks.

326. The computer system of claim 325, wherein the metadata in the data model is one or more XML tags and the text following at least one of the XML tags is identified as a candidate chunk.

327. A computer readable storage medium having stored therein instructions, which when executed by a computer system cause the computer system to:

retrieve a document from a data source, wherein the document has a structure type; generate a customized data model for the document in accordance with its structure type; and

identify one or more candidate chunks within the customized data model in accordance with a set of heuristic rules associated with the structure type.

328. The computer readable storage medium of claim 327, further comprising:

instructions for selecting at least one of the candidate chunks that satisfy one or more search keywords.

329. The computer readable storage medium of claim 327, wherein the customized data model is a hierarchical data model.

330. The computer readable storage medium of claim 327, wherein the structure type is one selected from the group consisting of structured, semi-structured, and unstructured.

331. The computer readable storage medium of claim 327, wherein the data source is a web server.

332. The computer readable storage medium of claim 327, wherein the document is an HTML web page.

333. The computer readable storage medium of claim 332, wherein the HTML web page includes multiple pairs of HTML tags, further comprising:

instructions for identifying a first subset of the HTML web page between a first pair of HTML tags as a first candidate chunk if the first pair of HTML tags satisfy one of the set of heuristic rules.

334. The computer readable storage medium of claim 333, further comprising:

instructions for recursively identifying a second subset of the HTML web page within the first subset of the HTML web page between a second pair of HTML tags as a second candidate chunk if the second pair of HTML tags satisfy one of the set of heuristic rules.

335. The computer readable storage medium of claim 327, wherein the document is a plain-text document.

336. The computer readable storage medium of claim 335, wherein the instructions for generating the customized data model further include instructions for inserting metadata into the data model that separates the plain-text document into multiple candidate chunks.

337. The computer readable storage medium of claim 336, wherein the metadata in the data model is one or more XML tags and the text following at least one of the XML tags is identified as a candidate chunk.

338. A computer-implemented method, comprising:

identifying within a document multiple matching chunks in response to a search request from a user, wherein the search request includes one or more search keywords and each of the multiple matching chunks matches at least one of the search keywords;

partitioning the matching chunks into multiple groups, wherein matching chunks within a respective group have an associated matching level to the search request; and

returning one or more groups of the matching chunks to the user in an order consistent with their respective matching levels to the search request.

339. The method of claim 338, wherein each of the search keywords has an associated weight indicative of its relevance to the user's search interest.

340. The method of claim 339, wherein one of the search keywords has an associated weight of zero.

341. The method of claim 338, wherein the matching level of a respective group of matching chunks is, at least partially, determined by summing the weights of unique search keywords within one of the matching chunks.

342. The method of claim 338, wherein the matching level of a respective group of matching chunks is, at least partially, determined by the number of unique search keywords within one of the matching chunks.

343. The method of claim 338, wherein partitioning the matching chunks into multiple groups further includes:

selecting one of the matching chunks;

determining the selected matching chunk's matching level and length;

5 invalidating the selected matching chunk if its matching level is less than a minimum matching level or if its length is outside a predefined range; and

inserting the selected matching chunk into one of the groups of matching chunks if its matching level is not less than the minimum matching level and its length is not outside the predefined range.

10 344. The method of claim 343, wherein the length of the matching chunk is the total word count of the textual content of the matching chunk.

345. The method of claim 343, wherein the length of the matching chunk is the total character count of the textual content of the matching chunk after white-space normalization.

15 346. The method of claim 343, wherein inserting the selected matching chunk into one of the groups of matching chunks further includes:

identifying a respective group of matching chunks whose matching levels are the same or similar to the selected matching chunk's matching level; and

adding the selected matching chunk to the group of matching chunks.

347. The method of claim 346, further comprising:

20 identifying within a group of matching chunks one or more matching chunks that are descendants of the selected matching chunk in a hierarchical data model of the document; and

invalidating the identified descendant matching chunks from the group of matching chunks.

348. The method of claim 343, further comprising:

25 after inserting the selected matching chunk into one of the groups of matching chunks,

determining a total count of matching chunks whose matching levels are no less than the minimum matching level;

updating the minimum matching level if the total count of matching chunks is greater than a predefined count threshold; and

30 invalidating at least a subset of one of the groups of matching chunks whose matching levels are less than the updated minimum matching level.

349. The method of claim 338, wherein returning one or more groups of the matching chunks to the user further includes:

selecting among the groups of matching chunks a group of matching chunks that has a highest matching level;

5 returning the selected group of matching chunks to the user; and

repeating said selecting and returning of a group of matching chunks having a next highest matching level until the total count of the returned matching chunks is not less than a predefined count threshold.

350. The method of claim 338, wherein returning one or more groups of the matching
10 chunks to the user further includes:

displaying a respective relevancy indicator adjacent each of the returned matching chunks.

351. The method of claim 350, wherein the relevancy indicator comprises image, text, number or a combination thereof.

15 352. A computer system, comprising:

memory;

one or more processors;

one or more programs stored in the memory and configured for execution by the one or more processors, the one or more programs including:

20 instructions for identifying within a document multiple matching chunks in response to a search request from a user, wherein the search request includes one or more search keywords and each of the multiple matching chunks matches at least one of the search keywords;

25 instructions for partitioning the matching chunks into multiple groups, wherein matching chunks within a respective group have an associated matching level to the search request; and

instructions for returning one or more groups of the matching chunks to the user in an order consistent with their respective matching levels to the search request.

30 353. The computer system of claim 352, wherein each of the search keywords has an associated weight indicative of its relevance to the user's search interest.

354. The computer system of claim 353, wherein one of the search keywords has an associated weight of zero.

355. The computer system of claim 352, wherein the matching level of a respective group of matching chunks is, at least partially, determined by summing the weights of unique search
5 keywords within one of the matching chunks.

356. The computer system of claim 352, wherein the matching level of a respective group of matching chunks is, at least partially, determined by the number of unique search keywords within one of the matching chunks.

357. The computer system of claim 352, wherein the instructions for partitioning the
10 matching chunks into multiple groups further include:

instructions for selecting one of the matching chunks;

instructions for determining the selected matching chunk's matching level and length;

instructions for invalidating the selected matching chunk if its matching level is less than a minimum matching level or if its length is outside a predefined range; and

15 instructions for inserting the selected matching chunk into one of the groups of matching chunks if its matching level is not less than the minimum matching level and its length is not outside the predefined range.

358. The computer system of claim 357, wherein the length of the matching chunk is the total word count of the textual content of the matching chunk.

20 359. The computer system of claim 357, wherein the length of the matching chunk is the total character count of the textual content of the matching chunk after white-space normalization.

360. The computer system of claim 357, wherein the instructions for inserting the selected matching chunk into one of the groups of matching chunks further include:

25 instructions for identifying a respective group of matching chunks whose matching levels are the same or similar to the selected matching chunk's matching level; and

instructions for adding the selected matching chunk to the group of matching chunks.

361. The computer system of claim 360, further comprising:

instructions for identifying within a group of matching chunks one or more matching chunks that are descendants of the selected matching chunk in a hierarchical data model of the document; and

5 instructions for invalidating the identified descendant matching chunks from the group of matching chunks.

362. The computer system of claim 357, further comprising:

instructions for determining a total count of matching chunks whose matching levels are no less than the minimum matching level;

10 instructions for updating the minimum matching level if the total count of matching chunks is greater than a predefined count threshold; and

instructions for invalidating at least a subset of one of the groups of matching chunks whose matching levels are less than the updated minimum matching level.

363. The computer system of claim 352, wherein the instructions for returning one or more groups of the matching chunks to the user further include:

15 instructions for selecting among the groups of matching chunks a group of matching chunks that has a highest matching level;

instructions for returning the selected group of matching chunks to the user; and

20 instructions for repeating said selecting and returning of a group of matching chunks having a next highest matching level until the total count of the returned matching chunks is not less than a predefined count threshold.

364. The computer system of claim 352, wherein the instructions for returning one or more groups of the matching chunks to the user further includes:

displaying a respective relevancy indicator adjacent each of the returned matching chunks.

25 365. The computer system of claim 364, wherein the relevancy indicator comprises image, text, number or a combination thereof.

366. A computer readable storage medium having stored therein instructions, which when executed by a computer system cause the computer system to:

30 identify within a document multiple matching chunks in response to a search request from a user, wherein the search request includes one or more search keywords and each of the multiple matching chunks matches at least one of the search keywords;

partition the matching chunks into multiple groups, wherein matching chunks within a respective group have an associated matching level to the search request; and

return one or more groups of the matching chunks to the user in an order consistent with their respective matching levels to the search request.

5 367. The computer readable storage medium of claim 366, wherein each of the search keywords has an associated weight indicative of its relevance to the user's search interest.

368. The computer readable storage medium of claim 367, wherein one of the search keywords has an associated weight of zero.

10 369. The computer readable storage medium of claim 366, wherein the matching level of a respective group of matching chunks is, at least partially, determined by summing the weights of unique search keywords within one of the matching chunks.

370. The computer readable storage medium of claim 366, wherein the matching level of a respective group of matching chunks is, at least partially, determined by the number of unique search keywords within one of the matching chunks.

15 371. The computer readable storage medium of claim 366, wherein the instructions for partitioning the matching chunks into multiple groups further include:
instructions for selecting one of the matching chunks;
instructions for determining the selected matching chunk's matching level and length;
instructions for invalidating the selected matching chunk if its matching level is less
20 than a minimum matching level or if its length is outside a predefined range; and
instructions for inserting the selected matching chunk into one of the groups of matching chunks if its matching level is not less than the minimum matching level and its length is not outside the predefined range.

25 372. The computer readable storage medium of claim 371, wherein the length of the matching chunk is the total word count of the textual content of the matching chunk.

373. The computer readable storage medium of claim 371, wherein the length of the matching chunk is the total character count of the textual content of the matching chunk after white-space normalization.

374. The computer readable storage medium of claim 371, wherein the instructions for inserting the selected matching chunk into one of the groups of matching chunks further include:

instructions for identifying a respective group of matching chunks whose matching
5 levels are the same or similar to the selected matching chunk's matching level; and
instructions for adding the selected matching chunk to the group of matching chunks.

375. The computer readable storage medium of claim 374, further comprising:

instructions for identifying within a group of matching chunks one or more matching
chunks that are descendants of the selected matching chunk in a hierarchical data model of
10 the document; and

instructions for invalidating the identified descendant matching chunks from the
group of matching chunks.

376. The computer readable storage medium of claim 371, further comprising:

instructions for determining a total count of matching chunks whose matching levels
15 are no less than the minimum matching level;

instructions for updating the minimum matching level if the total count of matching
chunks is greater than a predefined count threshold; and

instructions for invalidating at least a subset of one of the groups of matching chunks
whose matching levels are less than the updated minimum matching level.

20 377. The computer readable storage medium of claim 366, wherein the instructions for returning one or more groups of the matching chunks to the user further include:

instructions for selecting among the groups of matching chunks a group of matching
chunks that has a highest matching level;

instructions for returning the selected group of matching chunks to the user; and

25 instructions for repeating said selecting and returning of a group of matching chunks
having a next highest matching level until the total count of the returned matching chunks is
not less than a predefined count threshold.

378. The computer readable storage medium of claim 366, wherein the instructions for returning one or more groups of the matching chunks to the user further includes:

30 displaying a respective relevancy indicator adjacent each of the returned matching
chunks.

379. The computer readable storage medium of claim 378, wherein the relevancy indicator comprises image, text, number or a combination thereof.

380. A graphical user interface on a computer display, comprising:

a set of search keywords; and

5 a document link followed by a first chunk and a second chunk identified within the corresponding document,

wherein the first chunk and the second chunk are identified in response to a user search request associated with the set of search keywords,

10 wherein the first chunk satisfies a first subset of the set of search keywords and the second chunk satisfies a second subset of the set of search keywords that is not identical to the first subset of search keywords.

381. The graphical user interface of claim 380, wherein the first chunk and the second chunk are displayed in an order consistent with their relevancy to their respective subsets of search keywords.

15 382. The graphical user interface of claim 380, wherein there is at least one search keyword in both the first subset of search keywords and the second subset of search keywords.

383. The graphical user interface of claim 380, wherein there is a respective relevancy indicator adjacent each of the first and second chunks.

20 384. The graphical user interface of claim 383, wherein the relevancy indicator comprises image, text, number or a combination thereof.

385. A computer-implemented method, comprising:

receiving a request to search one or more secondary documents, wherein at least one of the secondary documents is associated with a primary document;

25 searching at least a subset of the secondary documents for documents that satisfy the search request; and

identifying at least one secondary document that satisfies the search request.

386. The method of claim 385, further comprising:

displaying the primary document on a display device before receiving the search request from a user, wherein the primary document includes one or more document links, each document link referencing one of the secondary documents; and

displaying at least a portion of the identified secondary document to the user.

5 387. The method of claim 386, further comprising:

locating within the identified secondary document one or more chunks that satisfy the search request; and

displaying one or more of the identified chunks to the user.

10 388. The method of claim 385, wherein the subset of the secondary documents is identified in response to a user selection of one or more document links in the primary document, each user-selected document link referencing a respective secondary document in the subset.

389. The method of claim 388, wherein each of the one or more document links in the primary document is selected by a respective mouse click of the document link.

15 390. The method of claim 388, wherein the one or more document links in the primary document are selected by:

defining a region in the primary document using an input device; and

identifying document links within the defined region as the selected document links.

391. The method of claim 390, wherein the input device is a mouse and defining a region in the primary document further includes:

20 pressing down the mouse's button at a first location;

dragging the mouse to a second location; and

releasing the mouse's button at the second location, wherein the region is a rectangle area defined by the first location and the second location.

25 392. The method of claim 385, wherein the primary document is a web page and the web page includes one or more document links, each document link referencing a respective secondary document.

393. The method of claim 385, further comprising:

searching the primary and secondary documents for chunks that satisfy the search request; and

identifying at least a first chunk in the primary document and at least a second chunk in one of the secondary documents, wherein both the first and second chunks satisfy the search request.

394. The method of claim 393, further comprising:

5 displaying the first chunk and the second chunk in a visually distinguishable manner.

395. The method of claim 385, wherein the search request includes one or more search keywords.

396. The method of claim 385, wherein the secondary documents include a first document and a second document, wherein the second document is indirectly referenced by the primary
10 document through at least the first document, further comprising:

searching the primary document and the secondary documents for chunks that satisfy the search request; and

identifying at least a first chunk in the primary document and at least a second chunk in the second document, wherein both the first and second chunks satisfy the search request.

15 397. A computer system, comprising:

memory;

one or more processors;

one or more programs stored in the memory and configured for execution by the one or more processors, the one or more programs including:

20 instructions for receiving a request to search one or more secondary documents, wherein at least one of the secondary documents is associated with a primary document;

instructions for searching at least a subset of the secondary documents for documents that satisfy the search request; and

25 instructions for identifying at least one secondary document that satisfies the search request.

398. The computer system of claim 397, further comprising:

instructions for displaying the primary document on a display device before receiving the search request from a user, wherein the primary document includes one or more

30 document links, each document link referencing one of the secondary documents; and

instructions for displaying at least a portion of the identified secondary document to the user.

399. The computer system of claim 398, further comprising:

instructions for locating within the identified secondary document one or more chunks
5 that satisfy the search request; and
instructions for displaying one or more of the identified chunks to the user.

400. The computer system of claim 397, wherein the subset of the secondary documents is identified in response to a user selection of one or more document links in the primary document, each user-selected document link referencing a respective secondary document in
10 the subset.

401. The computer system of claim 400, wherein each of the one or more document links in the primary document is selected by a respective mouse click of the document link.

402. The computer system of claim 400, wherein the one or more document links in the primary document are selected by execution of:

15 instructions for defining a region in the primary document using an input device; and
instructions for identifying document links within the defined region as the selected document links.

403. The computer system of claim 402, wherein the input device is a mouse and the instructions for defining a region in the primary document further include:

20 instructions for pressing down the mouse's button at a first location;
instructions for dragging the mouse to a second location; and
instructions for releasing the mouse's button at the second location, wherein the region is a rectangle area defined by the first location and the second location.

404. The computer system of claim 397, wherein the primary document is a web page and
25 the web page includes one or more document links, each document link referencing a respective secondary document.

405. The computer system of claim 397, further comprising:

instructions for searching the primary and secondary documents for chunks that satisfy the search request; and

instructions for identifying at least a first chunk in the primary document and at least a second chunk in one of the secondary documents, wherein both the first and second chunks satisfy the search request.

406. The computer system of claim 405, further comprising:

5 instructions for displaying the first chunk and the second chunk in a visually distinguishable manner.

407. The computer system of claim 397, wherein the search request includes one or more search keywords.

408. The computer system of claim 397, wherein the secondary documents include a first
10 document and a second document, wherein the second document is indirectly referenced by the primary document through at least the first document, further comprising:

instructions for searching the primary document and the secondary documents for chunks that satisfy the search request; and

15 instructions for identifying at least a first chunk in the primary document and at least a second chunk in the second document, wherein both the first and second chunks satisfy the search request.

409. A computer readable storage medium having stored therein instructions, which when executed by a computer system cause the computer system to:

20 receive a request to search one or more secondary documents, wherein at least one of the secondary documents is associated with a primary document;

search at least a subset of the secondary documents for documents that satisfy the search request; and

identify at least one secondary document that satisfies the search request.

410. The computer readable storage medium of claim 409, further comprising:

25 instructions for displaying the primary document on a display device before receiving the search request from a user, wherein the primary document includes one or more document links, each document link referencing one of the secondary documents; and

instructions for displaying at least a portion of the identified secondary document to the user.

30 411. The computer readable storage medium of claim 410, further comprising:

instructions for locating within the identified secondary document one or more chunks that satisfy the search request; and

instructions for displaying one or more of the identified chunks to the user.

412. The computer readable storage medium of claim 409, wherein the subset of the secondary documents is identified in response to a user selection of one or more document links in the primary document, each user-selected document link referencing a respective secondary document in the subset.

413. The computer readable storage medium of claim 412, wherein each of the one or more document links in the primary document is selected by a respective mouse click of the document link.

414. The computer readable storage medium of claim 412, wherein the one or more document links in the primary document are selected by execution of:

instructions for defining a region in the primary document using an input device; and

instructions for identifying document links within the defined region as the selected document links.

415. The computer readable storage medium of claim 414, wherein the input device is a mouse and the instructions for defining a region in the primary document further include:

instructions for pressing down the mouse's button at a first location;

instructions for dragging the mouse to a second location; and

instructions for releasing the mouse's button at the second location, wherein the region is a rectangle area defined by the first location and the second location.

416. The computer readable storage medium of claim 409, wherein the primary document is a web page and the web page includes one or more document links, each document link referencing a respective secondary document.

417. The computer readable storage medium of claim 409, further comprising:

instructions for searching the primary and secondary documents for chunks that satisfy the search request; and

instructions for identifying at least a first chunk in the primary document and at least a second chunk in one of the secondary documents, wherein both the first and second chunks satisfy the search request.

418. The computer readable storage medium of claim 417, further comprising:
instructions for displaying the first chunk and the second chunk in a visually
distinguishable manner.

419. The computer readable storage medium of claim 409, wherein the search request
includes one or more search keywords.

420. The computer readable storage medium of claim 409, wherein the secondary
documents include a first document and a second document, wherein the second document is
indirectly referenced by the primary document through at least the first document, further
comprising:

instructions for searching the primary document and the secondary documents for
chunks that satisfy the search request; and

instructions for identifying at least a first chunk in the primary document and at least a
second chunk in the second document, wherein both the first and second chunks satisfy the
search request.

421. A graphical user interface on a computer display, comprising:

a set of search keywords;

a primary document; and

one or more document links, each document link referencing a respective secondary
document, wherein:

in response to a user selection of at least one of the document links, the user-
selected document links are highlighted; and

in response to a user request to search the primary and secondary documents
associated with the user-selected document links, one or more chunks identified within the
primary and secondary documents are displayed, wherein each chunk satisfies at least one of
the set of search keywords.

422. The graphical user interface of claim 421, wherein each of the document links in the
primary document is selected by a respective mouse click of the document link.

423. The graphical user interface of claim 421, wherein the document links in the primary
document are selected by:

defining a region in the primary document using an input device; and

identifying document links within the defined region as the selected document links.

424. The graphical user interface of claim 423, wherein the input device is a mouse and the region in the primary document is defined by:

pressing down the mouse's button at a first location;

dragging the mouse to a second location; and

5 releasing the mouse's button at the second location, wherein the region is a rectangle area defined by the first location and the second location.

425. The graphical user interface of claim 421, wherein the chunks identified within the primary document and the chunks identified within the secondary documents are displayed in a visually distinguishable manner.




Figure 1




Figure 2




Figure 3




Figure 4




Figure 5




Figure 6




Figure 7




Figure 8A




Figure 8B




Figure 8C




Figure 9A

```

<html>
  <body>
    <div align="center">
960 <center>
      <table>
962 <tr>
964 <td>
966 <div align="right">
      <table>
968 <tr>
970 <td>
      <table border="0" cellpadding="0">
        <tbody>
972 <tr>
974 <td>
          <p align="left">One of the best sites to search for CAD jobs
            is<a shape="rect" href="http://www.cadtalent.com/">CadTalent</a>.
            CadTalent specializes in CAD jobs and only CAD jobs. You can
976 search CAD jobs by platform and location.
          </p>
978 <p align="left"><span lang="en-us">W</span>ondering where you
            can find the most job listings?
            <a>Monster.com</a>, true to it's name, leads the pack with
980 a million job listings and almost 100,000 internship listings.
          </p>
982 </td>
        </tr>
        </tbody>
      </table>
    </td>
  </tr>
</table>
</div>
</body>
</html>

```

Figure 9B illustrates a nested HTML structure for a web page layout. The structure is as follows:

- Outer Table (942):** A single-row table with one cell (944).
 - Cell 944:** Contains a centered **div** (946) with `align="center"`.
 - Inner Table (948):** A single-row table with one cell (950).
 - Cell 950:** Contains a **table** (952) with `border="0" cellpadding="0"`.
 - Table Body (954):** Contains one row (956) with one cell (958).
 - Cell 958:** Contains two paragraphs of text.
 - Paragraph 1 (976): "One of the best sites to search for CAD jobs is [CadTalent](http://www.cadtalent.com/). CadTalent specializes in CAD jobs and only CAD jobs. You can search CAD jobs by platform and location."
 - Paragraph 2 (980): "Wondering where you can find the most job listings? [Monster.com](#), true to it's name, leads the pack with a million job listings and almost 100,000 internship listings."

Figure 9B




Figure 10A




Figure 10B




Figure 11A

1110

```

<bib>
{
  for $b in doc('bib.xml')//book
  where $b/publisher = "Addison-Wesley"
  and   $b/@year > 1991
  return
    <book year="{ $b/@year }">
      { $b/title }
    </book>
}
</bib>

```

Figure 11B

1115

Input Sequence	Path Filter	Mode	Parent
Node Sub-Stream (0)	bib/book	Context	N/A
Node Sub-Stream (1)	bib/book/publisher	Content	Node Sub-Stream (0)
Node Sub-Stream (2)	bib/book/@year	Context	Node Sub-Stream (0)
Node Sub-Stream (3)	bib/book/@year	Context	Node Sub-Stream (0)
Node Sub-Stream (4)	bib/book/title	All	Node Sub-Stream (0)

Figure 11C




Figure 11D




Figure 11E




Figure 11F




Figure 11G




Figure 12A




Figure 12B

1220

1221

TigerLogic [einstein bohr debate]  **Chunkle**

☐ Best Match ☒ Match All ☐ "Exact" Match ☐ Match Any

1226-A 1226-B 1226-C 1226-D

1 2 3 4 5 6 7 8 9 10 

Results: 1 - 10 of 63,700

Bohr-Einstein debates - Wikipedia, the free encyclopedia 1225

The second phase of *Einstein's debate* with *Bohr* and the orthodox interpretation is characterized by an acceptance of the fact that it is, as a practical ...
<http://en.wikipedia.org/Chunkle/Reggie/Links/Hide/Chunkle>

1231-A The *Bohr-Einstein* debate 1223 1224 that was actually a series of epistemological challenges presented by Albert *Einstein* against what has come to be called the standard or Copenhagen interpretation of quantum mechanics. Since *Einstein's* closest friend and primary interlocutor in the "school" of Copenhagen was the physicist Niels *Bohr*, and since it was *Bohr* who provided answers to most of the challenges presented by *Einstein*, what was actually a friendly and fruitful series of exchanges of ideas has taken on the label of a "debate".

1231-B *Einstein's* natural reference point, as mentioned above, was always Niels *Bohr*, as the person who, more than other members of the School of Copenhagen, was animated by a particular interest for the philosophical and epistemological aspects of the theory and drew inspiration from the surprising aspects of the microscopical world in order to present daring hypotheses about reality and about knowledge, such as his idea of complementarity. These two giants of scientific thought nurtured a profound respect for each other and they were both extremely attentive to the acute and penetrating observations of the other. The *debate* is not only of historical interest as we will see *Einstein's* attacks often provoked reactions on the part of *Bohr* which called into question the current elements of the formalization of QM and its interpretation. This articulated process, in which many other important scientists, from Ehrenfest and Heisenberg to Born and from Schrödinger to John von Neumann, took part, brought more and more detailed attention on certain particularly problematic points of the theory.

1231-C The second phase of *Einstein's debate* with *Bohr* and the orthodox interpretation is characterized by an acceptance of the fact that it is, as a practical matter, impossible to simultaneously determine the values of certain incompatible quantities, but the rejection that this implies that these quantities do not actually have precise values. He rejects the probabilistic interpretation of Born and insists that quantum probabilities are epistemic and

1227

Sponsored Links

Unlock the Power of Office 2007
 TigerLogic Office Search
 Download Free Trial
www.TigerLogic.com

1229-A 1229-B 1229-C

Figure 12C

1220





☒ Best Match
 ☐ Match All
 ☐ "Exact" Match
 ☐ Match Any

1226-A

[Back to Results](#)
[Chunked: http://en.wikipedia.org/wiki/Bohr-Einstein_debates](http://en.wikipedia.org/wiki/Bohr-Einstein_debates)

http://en.wikipedia.org/wiki/Bohr-Einstein_debates
[Chunk Page Links](#)
[Show Chunks](#)

<http://en.wikipedia.org>

Page Links:

1253-A

<http://en.wikipedia.org>
[Chunk Page Links](#)
[Show Chunks](#)

1254-A


 Niels **Bohr** and Albert **Einstein**. The picture was taken at Ehrenfest's home in Leiden, the occasion was most likely the 50th anniversary of Hendrik Lorentz doctorate (December 11, 1925).


 Niels **Bohr** and Alberet **Einstein** Foto by Ehrenfest {{GFDL}} category: Albert **Einstein** category: Neils **Bohr**

1253-B

<http://en.wikipedia.org>
[Chunk Page Links](#)
[Show Chunks](#)

1254-B

Niels Bohr


 Niels **Bohr** and Albert **Einstein** debating quantum theory at Paul Ehrenfest's home in Leiden (December 1925).


Bohr also conceived the principle of complementarity : that items could be separately analyzed as having several contradictory properties. For example, physicists currently conclude that light is both a wave and a stream of particles – two apparently mutually exclusive properties – on the basis of this principle. **Bohr** also found philosophical applications for this daringly original principle. Albert **Einstein** much preferred the determinism of classical physics over the probabilistic new physics of **Bohr** (to which Max Planck and **Einstein** himself had contributed). He and **Bohr** had good-natured arguments over the truth of this principle throughout their lives (see **Bohr Einstein debate**). One of **Bohr** most famous students was Werner Heisenberg, a crucial figure in the development of quantum mechanics, who was also head of the German atomic bomb project.

Figure 12D



☐ Best Match
 ☒ Match All
 ☐ "Exact" Match
 ☐ Match Any



☐ 1 ☐ 2 ☐ 3 ☐ 4 ☐ 5 ☐ 6 ☐ 7 ☐ 8 ☐ 9 ☐ 10 ☐ 21

Bohr-Einstein debates - Wikipedia, the free encyclopedia

The second phase of *Einstein's* "debate" with *Bohr* and the orthodox interpretation is characterized by an acceptance of the fact that it is, as a practical ...
<http://en.wikipedia.org> [Chunk Rage Links](#) [Hide Chunks](#)

Bohr v Einstein

1257-A

The *debate* between *Bohr* and *Einstein* over the interpretation of quantum theory

The *debate* between *Einstein* and *Bohr* was conducted with the two talking
<http://homepages.paradise.net.nz> [Chunk Rage Links](#) [Hide Chunks](#) 1257-B

Einstein/Bohr Debate and Quantum Computing

[PDF]
<http://www.nanohub.org> [Chunk Rage Links](#) [Hide Chunks](#)

 Location skipped: content can not be chunked.

nanoHUB - Einstein?Bohr Debate and Quantum Computing

nanoHUB.org: online simulation and more... Our mission: To be the place where experiment, theory and simulation meet and move nanoscience to nanotechnology ..
<http://www.nanohub.org> [Chunk Rage Links](#) [Hide Chunks](#)

 No relevant content found at this location. Re-Chunk with match any keyword.

Einstein Exhibit -- The Quantum and The Cosmo

At the Solvay Conferences of 1927 and 1930 the *debate* between *Bohr* and *Einstein* went on day and night, neither man conceding defeat. Werner Heisenberg ...
<http://www.aiaa.org> [Chunk Rage Links](#) [Hide Chunks](#)

Figure 12E



☐ Best Match
 ☐ Match All
 ☒ "Exact" Match
 ☐ Match Any



1226-C

1233-A

1233-B

1

2

3

4

5

6

7

8

9

10

21

Bohr-Einstein debates - EPR paradox


A selection of articles related to [Bohr-Einstein debates](#) - EPR paradox
<http://www.experiencefestival.com> Chunk Rage Links Hide Chunks


 A Wisdom Archive on [Bohr-Einstein debates](#) - EPR paradox



 A selection of articles related to [Bohr-Einstein debates](#) - EPR paradox


 More material related to [Bohr-einstein Debates](#) can be found here:



 Main Page. for. [Bohr-einstein Debates](#)


 Index of Articles. related to. [Bohr-einstein Debates](#)


 Index of Articles. related to. [Bohr-Einstein debates](#) - E...


[Bohr-Einstein debates](#), [Bohr-Einstein debates](#) - Diffraction the single slit experiment, [Bohr-Einstein debates](#) - EPR paradox, [Bohr-Einstein debates](#) - Interference the double slit experiment, [Bohr-Einstein debates](#) - Photon in a box, Afshar's experiment, Complementarity, Copenhagen interpretation, Quantum eraser, Schrödinger's cat, Uncertainty principle, Wheeler's delayed choice experiment

ARTICLES RELATED TO [Bohr-Einstein debates](#) - EPR paradox





[Bohr-Einstein debates](#) - Diffraction the single slit experiment

Figure 12F

26/70





☐ Best Match
 ☐ Match All
 ☒ "Exact" Match
 ☐ Match Any

☐ 1
 ☐ 2
 ☐ 3
 ☐ 4
 ☐ 5
 ☐ 6
 ☐ 7
 ☐ 8
 ☐ 9
 ☐ 10
 ☒ 20

Bohr-Einstein debates - EPR paradox

A selection of articles related to [Bohr-Einstein debates](#) - EPR paradox
<http://www.experiencefestival.com> Chunk Rage Links Hide Chunks

The [Bohr-Einstein debates](#) on foundational aspects on quantum mechanics happened during the Solvay conferences. They consisted of analyses of thought experiments. Put simply, they were an attempt by Einstein to explain away the aspects of Bohr's interpretation of Quantum Mechanics that he disliked. Bohr attempted (and, most scholars agree, largely succeeded) to rebuild these challenges. The [Bohr-Einstein debates](#) remain among the most important in the history of the philosophy of physics, and are certainly ...

A Wisdom Archive on [Bohr-Einstein debates](#) - EPR paradox

A selection of articles related to [Bohr-Einstein debates](#) - EPR paradox

More material related to [Bohr-einstein Debates](#) can be found here:

Main Page for [Bohr-einstein Debates](#)

Index of Articles. related to. [Bohr-einstein Debates](#)

Index of Articles. related to. [Bohr-Einstein debates](#) - E...

[Bohr-Einstein debates](#), [Bohr-Einstein debates](#) - Diffraction the single slit experiment, [Bohr-Einstein debates](#) - EPR paradox, [Bohr-Einstein debates](#) - Interference the double slit experiment, [Bohr-Einstein debates](#) - Photon in a box, Aishar's experiment, Complementarity, Copenhagen interpretation, Quantum eraser, Schrödinger's cat, Uncertainty principle, Wheeler's delayed choice experiment

ARTICLES RELATED TO [Bohr-Einstein debates](#) - EPR paradox

[Bohr-Einstein debates](#) - Diffraction the single slit experiment

Figure 12G

Figure 12H





☐ Best Match
 ☒ Match All
 ☐ "Exact" Match
 ☐ Match Any

1

2

3

4

5

6

7

8

9

10

21

Cosmology: Big Bang Theory

The Big bang Model is a broadly accepted theory for the origin and evolution of ... Alber .Einstein at the chalkboard. The first key idea dates to 1916 when ...
<http://map.gsfc.nasa.gov> [Chunk Page Links](#) [Hide Chunks](#)

After the introduction of General Relativity a number of scientists, including **Einstein** tried to apply the new gravitational dynamics to the universe as a whole. At the time this required as assumption about how the matter in the universe was distributed. The simplest assumption to make is that if you viewed the contents of the universe with sufficiently poor vision, it would appear roughly the same everywhere and in every direction. That is, the matter in the universe is homogenous and isotropic when averaged over very large scales. This is called the Cosmological Principle. This assumption is being tested continuously as we actually observe the distribution of galaxies on ever larger scales. The accompanying picture shows how uniform the distribution of measured galaxies is over a 30° swath of sky. In addition the cosmic microwave background radiation, the remnant heat from the **Big Bang**, has a temperature which is highly uniform over the entire sky. This fact strongly supports the notion that the gas which emitted this radiation long ago was very uniformly distributed.

Stephen King, The Big Bang, and God

Einstein considered this to be the greatest blunder of his scientific career ... The .big bang theory is significant evidence against Hinduism ...
<http://www.eaderu.com> [Chunk Page Links](#) [Hide Chunks](#)

No relevant content found at this location. Re-Chunk with match any keyword.

1241

1243

The Big Bang Theory

So how does it all stack up? How much evidence do we actually have to support the Big Bang model of the universe? **Einstein's** Theory of Relativity. ...
<http://www.thechalkboard.co.uk> [Chunk Page Links](#) [Hide Chunks](#)

Figure 121

☐ Best Match
 ☐ Match All
 ☐ Exact Match
 ☒ Match Any

☒ 1
 ☐ 2
 ☐ 3
 ☐ 4
 ☐ 5
 ☐ 6
 ☐ 7
 ☐ 8
 ☐ 9
 ☐ 10
 ☒ 21
 1226-D

Stephen King, The Big Bang, and God

Einstein considered this to be the greatest blunder of his scientific career ... The big bang theory is significant evidence against Hinduism ... <http://www.leaderu.com> [Chunk Page Links](#) [Hide Chunks](#)

▲ The idea that the universe had a specific time of origin has been philosophically resisted by some very distinguished scientists. We could begin with Arthur Eddington, who experimentally confirmed **Einstein's** general theory of relativity in 1919. He stated a dozen years later: "Philosophically, the notion of a beginning to the present order is repugnant to me and I should like to find a genuine loophole." He later said, "We must allow evolution an infinite amount of time to get started."

▲ Albert **Einstein's** reaction to the consequences of his own general theory of relativity appear to acknowledge the threat of an encounter with God. Through the equations of general relativity, we can trace the origin of the universe backward in time to some sort of a beginning. However, before publishing his cosmological inferences **Einstein** introduced a cosmological constant, a "fudge factor," to yield a static model for the universe. **Einstein** later considered this to be the greatest blunder of his scientific career.

▲ **Einstein** ultimately gave grudging acceptance to what he called "the necessity for a beginning" and eventually to "the presence of a superior reasoning power." But he never did accept the reality of a personal God.

▲ In 1946, George Gamow, a Russian-born scientist, proposed that the primeval fireball, the **big bang**, was an intense concentration of pure energy. It was the source of all the matter that now exists in the universe. The theory predicts that all the galaxies in the universe should be rushing away from each other at high speeds as a result of that initial **big bang**. A dictionary definition of the **big bang** theory is "the entire physical universe, all the matter and energy and even the four dimensions of time and space, burst forth from a state of infinite or near infinite density, temperature, and pressure."

▲ The 1965 observation of the microwave background radiation by Arno Penzias and Robert Wilson from the Bell Telephone Laboratories convinced most scientists of the validity of the **Big Bang** theory. Further observations reported

Figure 12J




Figure 13A




Figure 13B

The screenshot shows a web browser window with the address bar displaying http://en.wikipedia.org/wiki/Bohr-Einstein_debates. The page title is "Bohr-Einstein debates" and the subtitle is "From Wikipedia, the free encyclopedia". The article text discusses the Bohr-Einstein debates, mentioning Niels Bohr, Albert Einstein, and the Copenhagen interpretation of quantum mechanics. A black and white photograph of Niels Bohr and Albert Einstein is included, with a caption: "Niels Bohr with Albert Einstein at Paul Ehrenfest's home in Leiden (December 1925)".

Bohr-Einstein debates
From Wikipedia, the free encyclopedia

The Bohr-Einstein debates is a popular name given to what actually a series of epistemological challenges presented by Albert Einstein against what has come to be called the *standard* or Copenhagen interpretation of quantum mechanics. Since Einstein's closest friend and primary interlocutor in the "school" of Copenhagen was the physicist Niels Bohr, and since it was Bohr who provided answers to most of the challenges presented by Einstein, what was actually a friendly and fruitful series of exchanges of ideas has taken on the label of a "debate".

Einstein's position with respect to quantum mechanics is significantly more subtle and open-minded than it has often been portrayed in technical manuals and popular science articles. Be that as it may, his constant and powerful criticisms of the quantum "orthodoxy" compelled the defenders of that orthodoxy to sharpen and refine their understanding of the philosophical and scientific implications of their own theory. Einstein's natural reference point as mentioned above, was always Niels Bohr as the person who more than other members of the School of Copenhagen, was animated by a particular interest for the philosophical and epistemological aspects of the theory and drew inspiration from the surprising aspects of the microscopic world in order to present daring hypothesis about reality and about knowledge such as his idea of complementarity. These two giants of scientific thought nurtured a profound aspect for each other and they were both extremely attentive to the acute and penetrating observations of the other. The debate is not only of historical interest as we will see Einstein's attacks often provoked reactions on the part pf Bohr which called into question the crucial elements of the formalization of QM and of its interpretation. This articulated process in which many other important scientists from Ehrenfest and Heisenberg to Born and from Schrödinger to John Neumann, took part, brought more and more detailed attention on certain particularly problematic points of the theory.

Navigation:

- Main Page
- Contents
- Featured content
- Current events
- Random article

Interaction:

- About Wikipedia
- Community portal
- Recent changes
- Contact Wikipedia
- Donate to Wikipedia
- Help

Search:

Go Search

toolbox

- What links here

Figure 13C

[W http://en.wikipedia.org/wiki/Bohr-Einstein_debates](http://en.wikipedia.org/wiki/Bohr-Einstein_debates)

Google TigerLogo

1324

Bookmarks • 1 blocked • Check • AutoLink • AutoFill • Send • Settings • Snagit

Search Options • Chunk Page • Chunk Page Links

☆ Bohr-Einstein debates - Wikipedia, the free encyclopedia

»


Home • Site Map • Page • Tools • »

wave of limited spatial extension has been created following the explanation given above. In order to challenge the indeterminacy relation between time and energy, it is necessary to find a way to determine with an adequate precision the energy that the photon has brought with it. At this point, Einstein turns to his celebrated relation between mass and energy of special relativity: $E = mc^2$. From this it follows that knowledge of the mass of an object provides a precise indication about its energy. The argument is therefore very simple: if one weighs the box before and after the opening of the shutter and if a certain amount of energy has escaped from the box, the box will be lighter. The variation in mass multiplied by c^2 will provide precise knowledge of the energy emitted. Moreover, the clock will indicate the precise time at which the event of the particle's emission took place. Since, in principle, the mass of the box can be determined to an arbitrary degree of accuracy, the energy emitted can be determined with a precision ΔE as accurate as one desires. Therefore, the product $\Delta E \Delta t$ can be rendered less than what is implied by the principle of indeterminacy.


The idea is particularly acute and the argument seemed unassailable. It's important to consider the impact of all of these exchanges on the people involved at the time. Leon Rosenfeld, a scientist who had participated in the Congress, described the event several years later.

It was a real shock for Bohr...who, at first, could not think of a solution. For the entire evening he was extremely agitated and he continued passing from one scientist to another, seeking to persuade them that it could not be the case, that it would have been the end of physics if Einstein were right, but he couldn't come up with any way to resolve the paradox. I will never forget the image of the two antagonists as they left the club. Einstein, with his tall and commanding figure, who walked tranquilly, with a mildly ironic smile, and Bohr who trotted along beside him, full of excitement... The morning after the triumph of Bohr.

The "triumph of Bohr" consisted in his demonstrating once again, that Einstein's subtle argument was not conclusive, but even more so in the way that he arrived at this conclusion by appealing precisely to one of the great ideas of Einstein: the principle of equivalence between gravitational mass and inertial mass. Bohr showed that in order for Einstein's experiment to function, the box would have to be suspended on a spring in the middle of a gravitational field. In order to obtain a measurement of weight, a pointer would have to be attached to the box which corresponded with the index on a scale. After the release of a photon, weights could be added to the box to restore it to its original position and this would allow us




Einstein's thought experiment of 1930 as designed by Bohr. Einstein's box was supposed to prove the violation of the indeterminacy relation between time and energy.



George Gamov's make-believe experiment apparatus for validating the thought experiment at the Niels Bohr Institute in Copenhagen.

Figure 13D






Figure 13F

Figure 13G




Figure 14




Figure 15




Figure 16A




Figure 16B




Figure 16C




Figure 16D




Figure 16E




Figure 16F

Document1 - Microsoft Word

Home Insert Page Layout References Mailings Review View Add-Ins TigerLogic

Keyword(s)
Einstein general relativity

Results Panel Search Within Results Find Search Assistant TigerLogic Search Previous Next

Contextual Search

File Name Last Mod. Type

Einstein- 1/15/2006 10:45 PM

Antigravity.doc:

PARA Starting in the early 1920's Albert *Einstein* and others began to speculate that *general relativity*, which describes gravitation and space-time, could be modified to include the laws of Maxwell that describe electromagnetism. In essence Albert *Einstein* sought to show that the laws of electricity and magnetism could be "unified" with the laws of gravitation. In other words such a theory would imply that all electrical and magnetic effects and all gravity effects are manifestations of an underlying "unified field". By the late 1920's Albert *Einstein's* papers on Unified Field theory began to be read by physicists interested in exploring experimental verification of his theories.

Einstein general relativity.pptx 1/15/2006 10:45 PM

SL 21 TITLE: Cosmology: Albert *Einstein*
"Cosmology Considerations on the *general Theory of relativity* - Albert *Einstein* - (1917) [see Cosmological Constants, eds Bernstein & Feinberg] - Copernican Principle gives way to Cosmological Principle - Model universe as spatially homogenous isotropic-Mach's

Matching Slides: Select einstein general relativity pptx

15 16 17 18 19 20 21 22 23 24 25 26 27 28

1657

Cosmology: Albert Einstein

"Cosmology Considerations on the General Theory of Relativity"
Albert Einstein (1917) [see Cosmological Constants, eds Bernstein & Feinberg]

Copernican Principle gives way to Cosmological Principle
Model the universe as spatially homogenous, isotropic

Mach's influence on relativity and inertia
No inertia relative to spacetime, but inertia of masses relative to one another

Observation and experiment
Account for the small kinetic motions of stars and nebulae

First mathematical model of the universe in general relativity

Click to add notes

Page: 1 of 1 Words: 0

Figure 16G

Document1 - Microsoft Word

Home Insert Page Layout References Mailings Review View Add-ins TigerLogic

Keyword(s)

Einstein general relativity

Results Panel Search Within Results Find Search Assistant TigerLogic Search

Contextual Search

Source A Source B Source C Web Sources Help TigerLogic Settings

Previous Next

Word PowerPoint Excel PDF Outlook Text File Types

1661

Cosmology: Albert Einstein

"Cosmology Considerations on the General Theory of Relativity" - Albert Einstein (1917) [see Cosmological Constant, eds. Bernstein & Feinberg]

Copernican Principle give way to Cosmological Principle

Model the universe as spatially homogeneous, isotropic

Mach's influence on relativity and inertia

No inertia relative to spacetime, but inertia of matter relative to one another

Observation and experiment

Account for the small kinetic motions of stars and nebulae

Final mathematical model of the universe in general relativity

1657

Einstein

Considerations on the General Theory of Relativity" [see Cosmological Constant, eds. Bernstein & Feinberg]

give way to Cosmological Principle

universe as spatially homogeneous, isotropic

inertia and inertia

due to spacetime, but inertia of masses another

experiment

a small kinetic motions of stars and nebulae

model of the universe in general relativity

SL 21

File Name Last Mod. Type

Einstein- 1/15/2006 10:45 PM

Antigravity.doc

PARA ☐ Starting in the early 1920's Albert Einstein and others began to speculate that general relativity, which describes gravitation and space-time, could be modified to include the laws of Maxwell that describe electromagnetism. In essence Albert Einstein sought to show that the laws of electricity and magnetism could be "unified" with the laws of gravitation. In other words such a theory would imply that all electrical and magnetic effects and all gravity effects are manifestations of an underlying "unified field". By the late 1920's Albert Einstein's papers on Unified Field theory began to be read by physicists interested in exploring experimental verification of his theories.

Einstein_general 1/15/2006 10:45 PM

relativity.pptx

SL 21 ☐ TITLE: Cosmology: Albert Einstein

"Cosmology Considerations on the general Theory of relativity - Albert Einstein - (1917) [see Cosmological Constant, eds Bernstein & Feinberg] - Copernican Principle gives way to Cosmological Principle - Model universe as spatially homogenous isotropic-Mach's

Page 1 of 1 Words: 0

Figure 16H




Figure 16I




Figure 16J




Figure 17A




Figure 17B




Figure 17C




Figure 17D




Figure 17E




Figure 18A




Figure 18B




Figure 18C




Figure 18D




Figure 19




Figure 20




Figure 21A




Figure 21B




Figure 21C



☒ Best Match
 ☐ Match All
 ☐ "Exact" Match
 ☐ Match Any

☒ 1
 ☐ 2
 ☐ 3
 ☐ 4
 ☐ 5
 ☐ 6
 ☐ 7
 ☐ 8
 ☐ 9
 ☐ 10
 ☒ 21

2154

Moon - Wikipedia the free encyclopedia

This results in a 3.8 cm yearly increase in the distance between the two bodies. [53] The .Moon will continue to move slowly away from the .Earth until the ...

<http://en.wikipedia.org> [Chunk Rage Links](#) [Hide Chunks](#)

Gravitational coupling between the Moon and the oceans affects the orbit of the Moon. From the Moon's point of view, the tidal bulges are carried ahead by the rotation of the Earth, so that they don't point directly toward the Moon. The gravitational coupling drains .kinetic energy and .angular momentum from the Earth's rotation. In turn, angular momentum is added to the Moon's orbit. Somewhat counterintuitively, this moves the Moon to a higher orbit with a longer period. This results in a 3.8 cm yearly increase in the distance between the two bodies. [53] The Moon will continue to move slowly away from the Earth until the tidal effects between the two are no longer of significance, whereupon the Moon's orbit will stabilize.

The Moon (Latin : Luna) is Earth's only natural satellite and is the .fifth largest one in the .Solar System. The average centre-to-centre distance from the Earth to the Moon is 384,403 km which is about thirty times the diameter of the Moon. The Moon has a diameter of 3,474 km.[6] - slightly more than a quarter that of the Earth, and about two-thirds of the average east-west distance across the United States. This means that the volume of the Moon is about 2 percent that of the Earth. The .gravitational pull at its surface is about 17 percent of the Earth's. The Moon makes a complete .orbit around the Earth every 27.3 days, and the periodic variations in the geometry of the Earth-Moon .Sun system are responsible for the .lunar phases that repeat every 29.5 days.

When the Moon is at apogee, it is 11% farther from the Earth than it is at perigee. This is far enough that it cannot entirely block the bright light, so eclipses which occur near apogee are not total.

Figure 21D




Figure 22A




Figure 22B







☒ Best Match
 ☐ Match All
 ☐ "Exact" Match
 ☐ Match Any

[Back To Results](#)
[History: C1> C2> C3> C4>](#)
[Chunked: http://www.rainingdata.com/products/index.html](#)

[TigerLogic XML Database, Insurance data Services, Dynamic Data Hub, Clinical Trials data Services, Geospatial data management, Multidimensional Databases, and Development Solutions,](#)
<http://www.rainingdata.com>
[Chunk Rage Links](#)
[Hide Chunks](#)

 ***TigerLogic XDMS***

 ***TigerLogic XDMS*** offers highly effective solutions across multiple use cases to meet the needs of evolving and complex data management such as dynamic data hub, geospatial server and SOA infrastructure enhancement to industry solutions such as clinical trials and insurance data services.

 A high performance advantage of ***TigerLogic XDMS*** is that it offers up to 22x faster XML query performance than other XML database products and relational databases.

Page Links:


[TigerLogic Dynamic Data Hub](#)
<http://www.rainingdata.com>
[Chunk Rage Links](#)
[Hide Chunks](#)

 ***TigerLogic*** Dynamic Data Hub ***TigerLogic*** Dynamic Data Hub supports data either physically stored in the ***TigerLogic*** database with data refresh policies or virtually available through federated queries of the data sources key feature of the ***TigerLogic*** Dynamic Data Hub id that the database structure evolves to handle changing data sources such as legacy systems, databases, data warehouses and Comma Separated Value (CSV) files. Specifically, the ***TigerLogic*** XML Data Management Server (***XDMS***) builds in dexes and stores data based on the structure of the incoming data. For example, if the incoming data structure changes from 10 fields to 12 or 8, then ***TigerLogic XDMS*** represents the data as XML and evolves to handle both the old structures as well as the new structure (no unload) rebuild, reload process is necessary). ***TigerLogic*** Dynamic Data Hub also triggers a process handle the new data structure (email to a data steward, etc.).

 The ***TigerLogic*** Dynamic Data Hub intrinsically handles temporal data or dynamic data values. The ***TigerLogic XDMS*** structure allows nesting so old data values and time stamps can be stored as sub-elements. This functionality can be easily turned on and off.

[TigerLogic Dynamic Data Hub](#)
<http://www.rainingdata.com>
[Chunk Rage Links](#)
[Hide Chunks](#)

 ***TigerLogic***® High Volume Geospatial Server . Geospatial Data Managemnt

 High volume geospatial data in XML messages are becoming more prevalent in government (Homeland Security, Department of Defense) and commercial applications (Wireless Service Providers). ***TigerLogic*** XML data Management Server ***XDMS*** provides the most efficient and low cost method to handle this data by using commodity hardware, keeping data in XML format and using the in-memory cache.

 ***TigerLogic XDMS*** is an ideal environment for rich query, conversion and management of geographical and spatial data sets represented in Geography Markup Language (GML). Geospatial support is seamlessly integrated into the ***TigerLogic XDMS*** by providing Geospatial XQuery (GSX) functions and operators over a variety of data sources, from fast access to the native data store, to Web Services and beyond.

Figure 22C




Figure 22D

Document Search Server 2300




Figure 23

Client Computer 2400




Figure 24