



(51) **International Patent Classification:**
Not classified

(21) **International Application Number:**
PCT/US2019/056992

(22) **International Filing Date:**
18 October 2019 (18.10.2019)

(25) **Filing Language:** English

(26) **Publication Language:** English

(30) **Priority Data:**
201841039495 18 October 2018 (18.10.2018) IN
16/581,138 24 September 2019 (24.09.2019) US

(71) **Applicant: ORACLE INTERNATIONAL CORPORATION** [US/US]; 500 Oracle Parkway M/S 5OP7, Redwood Shores, California 94065 (US).

(72) **Inventors: GHOSHAL, Sandip;** Flat No- 414, Building- Uber Heights, Street- Uber, Area- KMR Estate, Serilingalpalli, Kondapur, Hyderabad / Telangana 500084 (IN). **PATTANAYAK, Nalini Kanta;** Flat-E1005, Aparna Cyber Commune, Nallagandla, Hyderabad / Telangana 500019 (IN). **PETER, Vivek;** 902, Block B, Aparna Towers, Kondapur, Hyderabad / Telangana 500084 (IN). **KADLA-BALU, Hareesh;** 91 Victor Street, Plainview, New York 11803 (US).

(74) **Agent: BRISNEHAN, Brian et al.;** 1100 Peachtree Street NE, Suite 2800, Mailstop: IP Docketing - 22, Atlanta, Georgia 30309 (US).

(81) **Designated States** (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN,

(54) **Title:** TECHNIQUES FOR RANKING CONTENT ITEM RECOMMENDATIONS

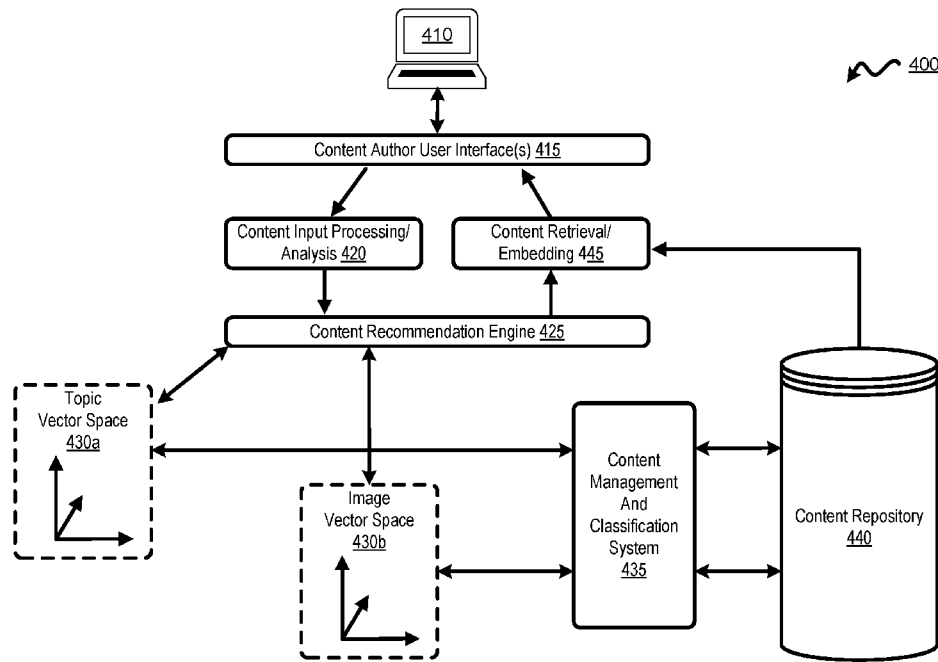


FIG. 4

(57) **Abstract:** Techniques describes herein include using software tools and feature vector comparisons to analyze and recommend images, text content, and other relevant media content from a content repository. A digital content recommendation tool may communicate with a number of back-end services and content repositories to analyze text and/or visual input, extract keywords or topics from the input, classify and tag the input content, and store the classified/tagged content in one or more content repositories. Input text and/or input images may be converted into vectors within a multi-dimensional vector space, and compared to a plurality of feature vectors within a vector space to identify relevant content items within a content repository. Such comparisons may include exhaustive deep searches and/or efficient tag-based filtered searches. Relevant content items (e.g., images, audio and/or video clips, links to related articles, etc.), may be retrieved and presented to a content author and embedded within original authored content.



HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:

— *without international search report and to be republished upon receipt of that report (Rule 48.2(g))*

TECHNIQUES FOR RANKING CONTENT ITEM RECOMMENDATIONS

PRIORITY CLAIM

[0001] This application:

(1) Claims priority to and is a continuation-in-part of U.S. Patent Application No. 16/581,138, filed September 24, 2019, and entitled “SMART CONTENT RECOMMENDATIONS FOR CONTENT AUTHORS,” which in turn claims the benefit of and priority to India Provisional Patent Application No. 201841039495, filed October 18, 2018, and entitled “SMART CONTENT RECOMMENDATIONS FOR CONTENT AUTHORS”; and

(2) Claims the benefit of and priority to India Provisional Patent Application No. 201841039495, filed October 18, 2018, and entitled “SMART CONTENT RECOMMENDATIONS FOR CONTENT AUTHORS.”

The entire contents of both U.S. Patent Application No. 16/581,138 and India Provisional Patent Application No. 201841039495 are incorporated herein by reference for all purposes.

BACKGROUND

[0002] Generators and authors of original content for online publication and/or transmission may use a variety of different software-based tools and techniques for generating, editing, and storing the newly generated content. For example, new original content may include on-line articles, newsletters, emails, blog entries, and the like, and such content may be generated via software-based word processor tools, email client applications, web-development tools, etc.

During the generation of new authored content, it may be useful for the author to identify and incorporate additional relevant items within the content, such as topical images and links to related articles, etc.

[0003] However, even when using search engines, keyword searching, and the like, locating and incorporating relevant images, links, and other related content may prove difficult and manually time-consuming. For example, when a generator or author of new content would like to incorporate related items into their newly authored text, the author must locate desired and relevant items, assure that the items are safe to insert into the authored content (e.g., free of

viruses and malware), confirm that the author and/or organization are authorized to use the items, and then embed the items within the newly authored content. Each of these processes may be technically challenging, manually intensive, and inefficient.

BRIEF SUMMARY

5 [0004] Aspects of the present disclosure relate to an artificial intelligence (AI) driven tool configured to function as a smart digital assistant to recommend images, text content, and other relevant media content from a content repository. Certain embodiments, may include front-end software tools having graphical user interfaces to supplement the content authoring interfaces being used to author the original media content (e.g., blog entries, online articles, etc.). In some 10 cases, the additional GUI screens and features may be incorporated into an existing content authoring software tool, for example, as software plug-ins. The smart digital content recommendation tool may communicate with a number of back-end services and content repositories, for example, to analyze text and/or visual input, extract keywords or topics from the input, classify and tag the input content, and store the classified/tagged content in one or more 15 content repositories.

[0005] Additional techniques performed in various embodiments of the smart digital content recommendation tool (e.g., either directly by the software tool, and/or indirectly by invoking back-end services), may include converting input text and/or images into vectors within a multi-dimensional vector space, and comparing input content to a plurality of repository content to find 20 a number of relevant content options within the content repository. Such comparisons may include thorough and exhaustive deep searches and/or more efficient tag-based filtered searches. Finally, relevant content items (e.g., images, audio and/or video clips, links to related articles, etc.), may be retrieved and presented to the content author to be reviewed and embedded within the original authored content.

25 BRIEF DESCRIPTION OF THE DRAWINGS

[0006] A further understanding of the nature and advantages of embodiments according to the present disclosure may be realized by reference to the remaining portions of the specification in conjunction with the following appended figures.

[0007] FIG. 1 is a diagram of an example computer system architecture including a data integration cloud platform in which certain embodiments of the present disclosure may be implemented.

5 [0008] FIG. 2 is an example screen of a customized dashboard in a user interface used to configure, monitor, and control a service instance, in accordance with certain embodiments of the present disclosure.

[0009] FIG. 3 is an architectural diagram of a data integration cloud platform, in accordance with certain embodiments of the present disclosure.

10 [0010] FIG. 4 is a diagram of an example computing environment configured to perform content classification and recommendation, in accordance with certain embodiments of the present disclosure.

[0011] FIG. 5 is another diagram of an example computing environment configured to perform content classification and recommendation, in accordance with certain embodiments of the present disclosure.

15 [0012] FIG. 6 is a flowchart illustrating a process for generating feature vectors based on content resources within a content repository, in accordance with certain embodiments of the present disclosure.

[0013] FIG. 7 is an example image identifying a plurality of image features, in accordance with certain embodiments of the present disclosure.

20 [0014] FIG. 8 is an example of a text document illustrating a keyword extraction process, in accordance with certain embodiments of the present disclosure.

[0015] FIGS. 9-11 are diagrams illustrating a process of generating and storing image tags, in accordance with certain embodiments of the present disclosure.

25 [0016] FIG. 12 is a flowchart illustrating another process for comparing feature vectors and identifying related content within a content repository, in accordance with certain embodiments of the present disclosure.

[0017] FIG. 13 is a diagram illustrating a technique of transforming an image filed into a feature vector, in accordance with certain embodiments of the present disclosure.

[0018] FIG. 14 is an illustrative vector space populated with feature vectors, in accordance with certain embodiments of the present disclosure.

5 [0019] FIG. 15 is a diagram illustrating a deep feature space vector comparison, in accordance with certain embodiments of the present disclosure.

[0020] FIGS. 16-17 are diagrams illustrating a filtered feature space vector comparison, in accordance with certain embodiments of the present disclosure.

10 [0021] FIG. 18 is a diagram representing a process of receiving and processing text input to identify related images or articles, in accordance with certain embodiments of the present disclosure.

[0022] FIG. 19 is an example diagram illustrating a comparison of extracted keywords to image tags, in accordance with certain embodiments of the present disclosure.

15 [0023] FIG. 20 is an example of a keyword analysis within a 3D word vector space, in accordance with certain embodiments of the present disclosure.

[0024] FIG. 21 is a diagram illustrating a keyword-to-tag vector space analysis, in accordance with certain embodiments of the present disclosure.

[0025] FIG. 22 is a diagram illustrating an example of homonym image tags, in accordance with certain embodiments of the present disclosure.

20 [0026] FIGS. 23-24 are diagrams illustrating an example disambiguation process, in accordance with certain embodiments of the present disclosure.

[0027] FIGS. 25-28 are diagrams illustrating a process of comparing feature vectors and identifying related content within a content repository, in accordance with certain embodiments of the present disclosure.

25 [0028] FIGS. 29-30 are example diagrams of a text document illustrating a topic extraction process, in accordance with certain embodiments of the present disclosure.

[0029] FIGS. 31-35 are diagrams illustrating a process of identifying related articles based on input text data, in accordance with certain embodiments of the present disclosure.

[0030] FIG. 36 is a diagram of an example semantic text analyzer system, in accordance with certain embodiments of the present disclosure.

5 [0031] FIGS. 37-38 are example user interface screens showing image recommendations provided to a user during the creation of content, in accordance with certain embodiments of the present disclosure.

[0032] FIG. 39 depicts a simplified diagram of a distributed system for implementing certain embodiments in accordance with present disclosure.

10 [0033] FIG. 40 is a simplified block diagram of one or more components of a system environment by which services provided by one or more components of a system may be offered as cloud services, in accordance with certain embodiments of the present disclosure.

[0034] FIG. 41 illustrates an exemplary computer system, in which various embodiments may be implemented.

15 [0035] FIG. 42 is a diagram of an example computing environment configured to evaluate and rank content items from a content repository in response to input content received from user or client systems, in accordance with certain embodiments of the present disclosure.

[0036] FIG. 43 is a flowchart illustrating a process for identifying and ranking content items that are relevant to user content, in accordance with certain embodiments of the present
20 disclosure.

[0037] FIG. 44 is an example screen of a content authoring user interface in accordance with certain embodiments of the present disclosure.

[0038] FIG. 45 shows an example table of a matching set of content items identified by the content recommendation system, in accordance with certain embodiments of the present
25 disclosure.

[0039] FIG. 46 shows another example table of a matching set of content items including ranking scores, in accordance with certain embodiments of the present disclosure.

[0040] FIG. 47 is another example screen of a content authoring user interface in accordance with certain embodiments of the present disclosure.

[0041] In the appended figures, similar components and/or features may have the same reference label. Further, various components of the same type may be distinguished by following the reference label by a dash and a second label that distinguishes among the similar components. If only the first reference label is used in the specification, the description is applicable to any one of the similar components having the same first reference label irrespective of the second reference label.

DETAILED DESCRIPTION

[0042] In the following description, for the purposes of explanation, specific details are set forth in order to provide a thorough understanding of various implementations and examples. It will be apparent, however, that various implementations may be practiced without these specific details. For example, circuits, systems, algorithms, structures, techniques, networks, processes, and other components may be shown as components in block diagram form in order not to obscure the implementations in unnecessary detail. The figures and description are not intended to be restrictive.

[0043] Some examples, such as those disclosed with respect to the figures in this disclosure, may be described as a process which is depicted as a flowchart, a flow diagram, a data flow diagram, a structure diagram, a sequence diagram, or a block diagram. Although a sequence diagram or a flowchart may describe the operations as a sequential process, many of the operations may be performed in parallel or concurrently. In addition, the order of the operations may be re-arranged. A process is terminated when its operations are completed, but could have additional steps not included in a figure. A process may correspond to a method, a function, a procedure, a subroutine, a subprogram, etc. When a process corresponds to a function, its termination may correspond to a return of the function to the calling function or the main function.

[0044] The processes depicted herein, such as those described with reference to the figures in this disclosure, may be implemented in software (e.g., code, instructions, program) executed by one or more processing units (e.g., processors cores), hardware, or combinations thereof. The

software may be stored in a memory (e.g., on a memory device, on a non-transitory computer-readable storage medium). In some examples, the processes depicted in sequence diagrams and flowcharts herein can be implemented by any of the systems disclosed herein. The particular series of processing steps in this disclosure are not intended to be limiting. Other sequences of steps may also be performed according to alternative examples. For example, alternative examples of the present disclosure may perform the steps outlined above in a different order. Moreover, the individual steps illustrated in the figures may include multiple sub-steps that may be performed in various sequences as appropriate to the individual step. Furthermore, additional steps may be added or removed depending on the particular applications. One of ordinary skill in the art would recognize many variations, modifications, and alternatives.

[0045] In some examples, each process in the figures of this disclosure can be performed by one or more processing units. A processing unit may include one or more processors, including single core or multicore processors, one or more cores of processors, or combinations thereof. In some examples, a processing unit can include one or more special purpose co-processors such as graphics processors, Digital Signal Processors (DSPs), or the like. In some examples, some or all of the processing units can be implemented using customized circuits, such as Application Specific Integrated Circuits (ASICs), or Field programmable gate arrays (FPGAs).

[0046] Certain embodiments described herein may be implemented as part of a Data Integration Platform Cloud (DIPC). In general, data integration involves combining data residing in different data sources and providing users with unified access and a unified view of the data. This process often arises and becomes significant in many situations, such as merging commercial entities with existing legacy databases. Data integration is beginning to appear more frequently in Enterprise software systems as the volume of data continues to increase along with the ability to analyze the data to provide useful results (“big data”). For example, consider a web application where users can query a variety of types of travel information (e.g., weather, hotels, airlines, demographics, crime statistics, etc.). Instead of requiring all of these various data types to be stored in a single database with a single schema, an Enterprise application can instead use the unified views and virtual schemas in the DIPC to combine many heterogeneous data sources such that they can be presented in a unified view to the user.

[0047] The DIPC is a cloud-based platform for data transformation, integration, replication, and governance. It provides batch and real-time data movement between cloud and on-premises data sources while maintaining data consistency with default-tolerance and resiliency. The DIPC may be used to connect to various data sources and prepare, transform, replicate, govern, and/or
5 monitored data from these various sources as they are combined into one or more data warehouses. The DIPC can work with any type of data source and support any type of data in any format. The DIPC can use a Platform as a Service (PaaS) or Infrastructure as a Service (IaaS) architecture to provide cloud-based data integration for an Enterprise.

[0048] The DIPC may provide a number of different utilities, including transferring entire data
10 sources to new cloud-based deployments and allowing cloud databases to be easily accessed from the cloud platform. Data can be streamed in real-time to new data sources up-to-date and to keep any number of distributed data sources synchronized. Loads may be divided amongst synchronized data sources such that they remain highly available to end-users. An underlying data management system can be used to reduce the amount of data moved over the network for
15 deployments into a database cloud, a big data cloud, third-party clouds, and so forth. A drag-and-drop user interface can be used to execute reusable Extract, Load, and Transform, (ELT) functions and templates. Real-time test environments can be created to perform reporting and data analytics in the cloud on replicated data sources such that the data can remain highly available to end-users. Data migrations can be executed with zero downtime using duplicated,
20 synchronized data sources. The synchronized data sources can also be used for a seamless disaster recovery that maintains availability.

[0049] FIG. 1 illustrates a computer system architecture that utilizes the DIPC to integrate data from various existing platforms, according to some embodiments. A first data source 102 may include a cloud-based storage repository. A second data source 104 may include an on-premises
25 data center. In order to provide uniform access and views to the first data source 102 and the second data source 104, the DIPC 108 can use an existing library of high-performance ELT functions 106 to copy data from the first data source 102 and the second data source 104. The DIPC 108 can also extract, enrich, and transform the data as it is stored in a new cloud platform. The DIPC 108 can then provide access to any big data utilities that are resident in or accessible
30 by the cloud platform. In some embodiments, the original data sources 102 and 104 may

continue to provide access to customers, while replicated data sources in the cloud platform can be used for testing, monitoring, governance, and big data analytics. In some embodiments, data governance may be provided to profile, cleanse, and govern data sources within an existing set of customized dashboards in a user interface.

5 [0050] FIG. 2 illustrates one of the customized dashboards in a user interface that can be used to configure, monitor, and control a service instance in the DIPC 108. A summary dashboard 202 can provide a control 204 that allows users to create a service instance. Next, a series of progressive web forms can be presented to walk the user through the types of information used to create a service instance. In a first step, the user will be asked to provide a service name and
10 description with an email address and a service edition type. The user may also be asked for a cluster size, which specifies a number of virtual machines used in the service. The service edition type determines which applications are installed on the virtual machines. In a second step and corresponding web form, the user may provide a running cloud database deployment to store the schemas of the DIPC server. The same database may later be used to store data entities
15 and perform integration tasks. Additionally, a storage cloud may be specified and/or provisioned as a backup utility. The user may also provide credentials that can be used to access the existing data sources used in the data integration. In a third step, the provisioning information can be confirmed and a service instance can be created. The new service instance may then be displayed in the summary area 206 of the summary dashboard 202. From there, the user can
20 access any information for any running data integration service instance.

[0051] FIG. 3 illustrates an architectural diagram of the DIPC, according to some embodiments. Requests may be received through a browser client 302, which may be implemented using a JavaScript Extension Toolkit (JET) set of components. Alternatively or
25 additionally, the system may receive requests through a DIPC agent 304 that operates at a customer's on-premises data center 306. The DIPC agent 304 may include a data integrator agent 308 and an agent 310 for a replication service, such as Oracle's GoldenGate® service. Each of these agents 308, 310 may retrieve information from the on-premises data center 306 during normal operations and transmit data using a connectivity service 312 back to the DIPC.

[0052] Incoming requests can be passed through a sign-in service 314, which may include
30 load-balancing or other utilities for routing requests through the DIPC. The sign-in service 314

may use an identity management service, such as an identity cloud service 316 to provide security and identity management for the cloud platform as part of an integrated enterprise security fabric. The identity cloud service 316 can manage user identities for both the cloud deployments and the on-premises applications described in this embodiment. In addition to the
5 identity cloud service 316, the DIPC may also use a PaaS Service Manager (PSM) tool 318 to provide an interface to manage the lifecycle of the platform service in the cloud deployment. For example, the PSM tool 318 can be used to create and manage instances of the data integration service in the cloud platform.

[0053] The DIPC can be implemented on a Web logic server 320 for building and deploying
10 enterprise applications in the cloud environment. The DIPC may include a local repository 322 that stores data policies, design information, metadata, and audit data for information passing through the DIPC. It may also include a monitoring service 324 to populate the local repository 322. A catalog service 326 may include a collection of machine-readable open APIs to provide access to many of the SaaS and PaaS applications in the cloud deployment. The catalog service
15 326 may also be available for a search application 338 that uses a distributed indexing service, such as the Apache Solr®. A connectivity service 328 and a mediator service 330 can manage connections and provide transformation, validation, and routing logic for information that passes through the DIPC. Information within the DIPC may be passed using an Event Driven Architecture (EDA) and a corresponding message bus 332.

[0054] The DIPC may also include an orchestration service 334. The orchestration service 334
20 may enable automation tasks by calling REST endpoints, scripts, third-party automation frameworks, and so forth. These tasks can then be executed by the orchestration service 334 to provide the DIPC functionality. The orchestration service 334 can use run-time services to import, transform, and store data. For example, an ELT runtime service 334 can execute the
25 library of ELT functions described above, while the replication runtime service 342 can copy data from various data sources into a cloud-deployed DIPC repository 316. Additionally, the DIPC may include a code generation service 336 that provides automatic code generation for both ELT functions and replication functions.

I. SMART CONTENT – SMART CONTENT RECOMMENDATIONS

[0055] As noted above, when a user is creating/authoring original media content (e.g., articles, newsletters, emails, blog posts, etc.), it is often useful to enhance the authored content with relevant additional content such as related images, audio/video clips, links to related articles or other content. However, searching for such additional content, as well as embedding that additional content within the user's original authored content, may be difficult in several respects. The initial difficulties may involve finding safe/secure additional content from a trusted source, and ensuring that user/author is authorized to incorporate that content within their work. Additionally, from any such secure and authorized content repositories, it may be a manually intensive and inefficient process for the user/author to locate and incorporate/embed any relevant content within their original authored content.

[0056] Therefore, certain aspects described herein relate to a smart digital content recommendation tool. In certain embodiments, the smart digital content recommendation tool may be an artificial intelligence (AI) driven tool configured to process and analyze input content (e.g., text, images) from a content author in real-time, and to recommend relevant images, additional text content, and/or other relevant media content (e.g., audio or video clips, graphics, social media posts, etc.), from one or more trusted content repositories. A smart digital content recommendation tool may communicate with a number of back-end services and content repositories, for example, to analyze text and/or visual input, extract keywords or topics from the input, classify and tag the input content, and store the classified/tagged content in one or more content repositories.

[0057] Additional aspects described herein, each of which may be performed directly via a smart digital content recommendation tool executing on a client operated by the content author, and/or indirectly by invoking various back-end services, may include (a) receiving original content as input, in the form of text and/or images, (b) extracting keywords and/or topics from the original content, (c) determining and storing associated keyword and/or topic tags for the original content, (d) converting the original content (e.g., input text and/or images) into vectors within a multi-dimensional vector space, (e) comparing such vectors to a plurality of other content vectors, each of which represents additional content in a content repository, in order to find and identify various potentially-relevant additional content related to the original content

input authored by the user/author, and finally (f) retrieve and present the identified additional content to the author via the smart digital content recommendation tool. In some embodiments, each additional content item (e.g., images, links to related articles or web pages, audio or video files, graphics, social media posts, etc.) may be displayed and/or thumbnailed by the smart digital content recommendation tool, in a GUI-based tool that allows the user to drag-and-drop or otherwise place the additional content within the user's original authored content, including content positioning, formatting, re-sizing, etc.

[0058] With reference now to FIG. 4, a block diagram is shown illustrating various components of a system 400 for smart content classification and recommendation, including a client device 410, content input processing and analysis service 420, content recommendation engine 425, a content management and classification system 435, and a content retrieval and embedding service 445. Additionally, system 400 includes one or more content repositories 440 storing content files/resources, and one or more vector spaces 430. As described in more detail below, a vector space may refer to a multi-dimensional data structure configured to store one or more feature vectors. In some embodiments, the recommendation engine 425, associated software components and services 420 and 445, content management and classification system 435, and content repositories 440 (which may store one or more data stores or other data structures) may be implemented and stored as back-end server systems remote from the front-end client device 410. Thus, interactions between client device 410 and content recommendation engine 425 may be Internet-based web browsing sessions, or client-server application sessions, during which users access may input original authored content via the client device 410, and receive content recommendations from content recommendation engine 425 in the form of additional content that is retrieved from the content repository 440 and linked or embedded into the content authoring user interface at the client device 410. Additionally or alternatively, the content recommendation engine 425 and/or content repository 440 and related services may be implemented as specialized software components executing on the client device 410.

[0059] The various computing infrastructure elements shown in this example (e.g., content recommendation engine 425, software components/services 420, 435, and 445, and content repositories 440) may correspond to a high-level computer architecture created and maintained by an enterprise or organization that provides Internet-based services and/or content to various

client devices 410. The content described herein (which also may be referred to as content resources and/or content files, content links, etc.) may be stored in one or more content repositories, retrieved and classified by the content recommendation engine 425, and provided to content authors at client devices 410. In various embodiments, various different media types or file types of content may be input as original content by a content author at client device 410, and similarly various different media types or file types of content may be stored in the content repositories 440 and recommended for/ embedded into the front-end user interfaces at the client device 410. These various different media types authored by or recommended to the content author may include text (e.g., authoring a letter, article, or blog), images (selected by or for the author), audio or video content resources, graphics, social media content, (e.g., posts, messages, or tweets on Facebook, Twitter, Instagram, etc.).

[0060] In some embodiments, system 400 shown in FIG. 4 may be implemented as a cloud-based multi-tier system, in which upper-tier user devices 410 may request and receive access to the network-based resources and services via the content processing/analysis component 420, and wherein the application servers may be deployed and executed upon an underlying set of resources (e.g., cloud-based, SaaS, IaaS, PaaS, etc.) including hardware and/or software resources. Additionally, although cloud-based systems may be used in some embodiments, systems 400 in other examples may use on-premise datacenters, server farms, distributed computing systems, and other various non-cloud computing architectures. Some or all of the functionality described herein for the content processing/analysis component 420, content recommendation engine 425, content management and classification system 435, content retrieval and embedding component 445, and the generation and storage of the vector spaces 430 may be performed by Representational State Transfer (REST) services and/or web services including Simple Object Access protocol (SOAP) web services or APIs, and/or web content exposed via Hypertext Transfer Protocol (HTTP) or HTTP Secure protocols. Thus, although not shown in FIG. 4 so as not to obscure the components shown with additional details, the computing environment 400 may include additional client devices 410, one or more computer network(s), one or more firewalls 435, proxy servers, and/or other intermediary network devices, facilitating the interactions between the client devices 410, content recommendation engine 425, and back-end content repositories 440. Another embodiment of similar system 500 is shown in more detail in FIG. 5.

[0061] Referring briefly to FIG. 5, another example diagram is shown of a computing environment 500 showing a data flow/data transformation diagram for performing content classification and recommendation. Thus, the computing environment 500 shown in this example may correspond to one possible implementation of the computing environment 400 described above in FIG. 4. In FIG. 5, several of the diagram blocks shown represent specific data states or data transformations, rather than the structural hardware and/or software components described above in FIG. 4. Thus, block 505 may represent input content data received via a user interface. Block 510 represents a set of keywords determined by the system 400, based on the input content 505. As discussed above, the keywords 510 may be determined using one or more keyword extraction and/or topic modeling processes, by the input processing/analysis component 420, and a text feature vector 515 may be generated based on the determined keywords 510.

[0062] Continuing with the example shown in FIG. 5, a number of additional feature vectors 520 may be retrieved from the content repository 440. In this example, the additional feature vectors 520 may be selected from the content repository 440 by executing one or more neural network trained image models, and providing the determined keywords 510 into the trained models. The resulting feature vectors 520 may further be narrowed to exclude those having a feature vector probability of less than $z\%$, based on the output of the trained model, resulting in a subset of the retrieved feature vectors 525. A feature space comparison 530 then may be performed between the test feature vector 515 and the subset of retrieved feature vectors 525. In some embodiments, and as shown in this example closest Euclidean distance calculation may be used to identify the closest retrieved feature vectors 525 to the test feature 515. Based on the feature space comparison 530, one or more recommendations 530 may be determined, each recommendation 530, based on the associated feature vector 525 having a threshold closeness to the test feature vector 515, and each recommendation 530 corresponding to an image within the content repository 440.

[0063] The components shown in system 400 for providing AI-based and feature vector analysis-based content recommendations and services to client devices 410 may be implemented in hardware, software, or a combination of hardware and software. For example, web services may be generated, deployed, and executed within datacenters 440, using the underlying system

hardware or software components such as data storage devices, network resources, computing resources (e.g., servers), and various software components. In some embodiments, web services may correspond to different software components executing on the same underlying computer server(s), networks, data stores, and/or within the same virtual machines. Some web-based content, computing infrastructure instances, and/or web services provided within the content recommendation engine 425 may use dedicated hardware and/or software resources, while others may share underlying resources (e.g., shared cloud). In either case, certain higher-level services (e.g., user applications), as well as users at client devices, need not ever be aware of the underlying resources being used to support the services.

10 **[0064]** In such implementations, the various application servers, database servers and/or cloud storage systems, and other infrastructure components such as web caches, network components, etc. (not shown in this example), may include various hardware and/or software components (e.g., application programming interfaces (APIs), cloud resource managers, etc.) to provide and monitor the classification and vectorization of content resources, as well as to manage the underlying storage/server/network resources. The underlying resources of the content repositories 440 may be stored within content databases and/or cloud storage systems which may include, for example, a set of non-volatile computer memory devices implemented as databases, file-based storage, etc., a set of network hardware and software components (e.g., routers, firewalls, gateways, load balancers, etc.), a set of host servers, and various software resources such as, store software images, installations, builds, templates, configuration files, and the like, corresponding to different versions of various platform, server, middleware, and application software. The datacenters housing the application servers of the recommendation engine 425, vector spaces 430, and related services/components also may include additional resources such as hypervisors, host operating systems, resource managers and other cloud-based applications, along with the hardware and software infrastructure to support various Internet-based services, such as Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and Software as a Service (SaaS). Additionally, the underlying hardware of the datacenters may be configured to support a number of internal shared services, which may include, for example, security and identity services, integration services, repository services, enterprise management services, virus scanning services, backup and recovery services, notification services, file transfer services, and the like.

[0065] As noted above, many different types of computer architectures (cloud-based, web-based, hosting, multi-tier computing environments, distributed computing environments, etc.) may be used to provide the web-based content recommendations from the content recommendation engine 425 (which may be implemented via one or more content recommendation application servers) to the client devices 410 in accordance with various embodiments described herein. However, in certain implementations, cloud computing platforms may be used to provide certain advantageous features for generation and management of web-based content. For example, cloud computing platforms may provide elasticity and scalability to quickly provision, configure, and deploy many different types of computing infrastructure instances, in contrast to non-cloud based implementations having fixed architectures and limited hardware resources. Moreover, public cloud, private cloud, and public-private hybrid cloud platforms may be used in various embodiments to leverage the features and advantages of each distinct architecture.

[0066] Additionally, as shown in this example, system 400 also includes a content management and classification system 435. In some embodiments, the content management and classification system 435 may include a distributed storage processing system, one or more machine learning-based classification algorithms (and/or non-machine learning-based algorithms), and/or a storage architecture. As discussed in more detail below, in some embodiments the content management and classification system 435 may access content resources (e.g., web-based articles, images, audio files, video files, graphics, social media content, etc.) via one or more content repositories 440 (e.g., network-based document stores, web-based content providers, etc.). For example, within system 400, specialized JavaScript or other software components may be installed and operate on one or more application servers, database servers and/or cloud systems storing content objects or network-based content. These software components may be configured to retrieve and transmit content resources (e.g., articles, images, web pages, documents, etc.), to the content management and classification system 435 for analysis and classification. For example, each time a user within the operating organization of the system 400 imports or creates new content such as an image or article, a software component may retrieve the content back to the content management and classification system 435, for various processing and analysis described below (e.g., image processing, keyword extraction, topic analysis, etc.). Additionally, although in this example, the content management

and classification system 435 is depicted as being implemented separately from the content recommendation engine 425 and content repositories 440, in other examples, content management and classification systems 435 may be implemented locally with either the content recommendation engine 425 and/or the storage devices storing content repositories 440, and thus
5 need not receive separate transmissions of content from those device, but rather may analyze and classify those content resources that are stored or provided by their respective systems.

[0067] One or more vector spaces 430 also may be generated and used to store feature vectors corresponding to the different content items within the content repositories 440, and to compare the feature vectors for original authored content (e.g., received from a client device 410) to the
10 feature vectors of the additional content items in the content repositories 440. In some embodiments, a plurality of multi-dimensional feature spaces 430 may be implemented within the system 400, such as first feature space 430a for topics of text input/articles, and a second feature space 430b for images. In other embodiments, additional separate multi-dimensional feature spaces 430 may be generated for different types of content media (e.g., a feature space
15 for audio data/files, a feature space for video data/files, a feature space for graphics, a feature space for social media content, etc.). As discussed below, comparison algorithms may be used to determine distances between the vectors within a feature space. Thus, in a feature space of image feature vectors, the algorithms may be used to identify the closest images to a received input image, in a feature space of text feature vectors, the algorithms may be used to identify the
20 closest texts (e.g., articles) to a received input text block, and so on. Additionally or alternatively, the comparison algorithm may use keywords/tags of the vector space to determine similarity between various media types.

[0068] In various implementations, system 400 may be implemented using one or more computing systems and/or networks. These computing systems can include one or more
25 computers and/or servers, which may be general purpose computers, specialized server computers (such as desktop servers, UNIX servers, mid-range servers, mainframe computers, rack-mounted servers, etc.), server farms, server clusters, distributed servers, or any other appropriate arrangement and/or combination of computing hardware. The content recommendation engine 425 may run an operating system and/or a variety of additional server
30 applications and/or mid-tier applications, including HyperText Transport Protocol (HTTP)

servers, File Transport Service (FTP) servers, Common Gateway Interface (CGI) servers, Java servers, database servers, and other computing systems. The content repository 440 may include database servers, for example, those commercially available from Oracle, Microsoft, etc. Each component in system 400 may be implemented using hardware, firmware, software, or combinations of hardware, firmware, and software.

[0069] In various implementations, each component within system 400 may include at least one memory, one or more processing units (e.g., processor(s)), and/or storage. The processing unit(s) can be implemented as appropriate in hardware (e.g., integrated circuits), computer-executable instructions, firmware, or combinations of hardware and instructions. In some examples, the various components of system 400 may include several subsystems and/or modules. The subsystems and/or modules in the content recommendation engine 425 may be implemented in hardware, software (e.g., program code or instructions executable by a processor) executing on hardware, or combinations thereof. In some examples, the software may be stored in a memory (e.g., a non-transitory computer-readable medium), on a memory device, or some other physical memory, and may be executed by one or more processing units (e.g., one or more processors, one or more processor cores, one or more Graphics Process Units (GPUs), etc.). Computer-executable instructions or firmware implementations of the processing unit(s) can include computer-executable or machine-executable instructions written in any suitable programming language, which can perform the various operations, functions, methods, and/or processes described herein. The memory may store program instructions that are loadable and executable on the processing unit(s), as well as data generated during the execution of these programs. The memory may be volatile (such as random access memory (RAM)) and/or non-volatile (such as read-only memory (ROM), flash memory, etc.). The memory may be implemented using any type of persistent storage device, such as computer readable storage media. In some examples, computer-readable storage media may be configured to protect a computer from an electronic communication containing malicious code. .

[0070] Referring now to FIG. 6, a flowchart is shown illustrating a process for generating feature vectors based on content resources within a content repository 440, and storing the feature vectors within a feature space 430. As described below, the steps in this process may be performed by one or more components in the computing environment 400, such as the content

management and classification system 435, and the various subsystems and subcomponents implemented therein.

[0071] In step 602, a content resource may be retrieved from a content repository 440 or other data store. As discussed above, an individual content resource (which also may be referred to as content or a content item) may correspond to a data object of any various content type: a text item (e.g., text file, article, email, blog entry, etc.), an image, an audio file, a video file, a 2D or 3D graphic object, a social media data item, etc. In some embodiments, content items may be retrieved from particular content repositories 440, such as proprietary data stores owned and operated by a particular trusted organization. Although content repositories 440 may be external data sources, such as Internet web servers or other remote data stores, systems 400 that retrieve and vectorize content from local and/or privately-controlled content repositories 440 may realize several technical advantages in the operation of system 400, including assurances that the content from repository 440 will be preserved and accessible when needed, and that the user/author is authorized to use and reproduce the content from repository 440. In some cases, the retrieval in step 602 (and subsequent steps 604-608) may be triggered in response to a new content item being stored in the content repository 440, and/or to a modification to an item in the content repository 440.

[0072] In step 604, the content item retrieved in step 602 may be parsed/analyzed/etc. in order to extract a set of item features or characteristics. The type of parsing, processing, feature extraction, and/or analysis performed in step 604 may depend on the type of the content item. For image content items, artificial intelligence-based image classification tools may be used to identify particular image features and/or generate image tags. As shown in the example image of FIG. 7, the image analyses may identify a plurality of image features (e.g., smile, waitress, counter, sale machine, coffee cup, cake, hand, food, human, café, etc.), and the image may be tagged with each of these identified features. For text-based content items, such as blog posts, letters, emails, articles, etc., the analysis performed in step 604 may include keyword extraction and processing tools (e.g., stemming, synonym retrieval, etc.) as shown in FIG. 8. One or both types of analyses (i.e., feature extraction from images as shown in FIG. 9, and keyword/topic extraction from text content as shown in FIG. 8), may be performed via REST-based services or other web services, using analytics, machine-learning algorithms, and/or artificial intelligence

(AI), such as an AI-based cognitive image analysis service, or a similar AI/REST cognitive text service to be used for the text content of FIG. 8. Similar techniques may be used in step 604 for other types of content items, such as video files, audio files, graphics, or social media posts, where a specialized web-service of the system 400 is used to extract and analyze specific
5 features (e.g., words, objects within images/video, facial expressions, etc.) depending on the media type of the content item.

[0073] In step 606, after extracting/determining the specific content features (e.g., visual objects, keywords, topics, etc.) from the content item, a feature vector may be generated based on the extracted/determined features. Using various transformation techniques, each set of
10 features associated with a content item may be transformed into a vector that can be input into a common vector space 430. The transformation algorithm may output a predetermined vector format (e.g., 1 x 4096 dimension vector). Then, in step 608, the feature vector may be stored in one or more of the vector spaces 430 (e.g., topic vector space 430a for text content, image vector space 430b for image content, and/or a combined vector space for multiple content types). The
15 vector spaces and feature vectors stored therein may be generated and maintained by the content management and classification system 435, the content recommendation engine 425, and/or other components of the system 400.

[0074] In some embodiments, a subset of the extracted/determined content features also may be saved as tags associated with the content item. Example processes for generating and storing
20 image tags based on images, and conversely retrieving images based on image tags, are shown in FIGS. 9-11. Although these examples relate to image content items, similar tagging processes and/or keyword or topic extraction may be performed for text content items, audio/video content items, etc. As shown in FIG. 9, in step 901 an image may be created and/or uploaded, for example, into a content repository 440. In step 902, the image may be transmitted from the
25 content repository 440 to an artificial intelligence (AI) based REST service configured to analyze the image and extract topics, themes, specific visual feature, and the like. The AI REST service may determine one or more specific image tags based on the identified image features, and in step 903 may transmit the image tags back to the content repository, to be stored either within or associated with the image in step 904. In FIG. 10, an identical process to that described in FIG.
30 9 is shown for generating and storing image tags based on images. Additionally, FIG. 10 shows

several illustrative features 1001 that may be implemented within the AI REST service in certain embodiments, including an image tag determination/retrieval component, an Apache MxNet component, and a cognitive image service. After determining one or more tags for a content item, the tags may be stored either back in the content repository 440, or in a separate storage location. For instance, referring to the example image of FIG. 7, a dozen or more potential image features may be extracted from this single image, all of which may be incorporated into the feature vector. However, AI REST service and/or the content repository 440 may determine that it is optimal for content match to tag the image only with a few of the most prevalent themes of the image (e.g., coffee, retail).

10 **[0075]** Referring briefly now to FIG. 11, another example process is shown relate to the described in FIGS. 9 and 10 for generating and storing image tags based on images. In FIG. 11, a converse process is illustrated, in which image tags are used to retrieve matching images from a content repository 440. In step 1101, a content authoring user interface 415 or other front end interface may determine one or more content tags based on the input received via the interface. 15 In this example, a single content tag (“waitress”) is determine from the received user input, and in step 1102 the content tag is transmitted to a search API associated with the content repository 440. A search API may be implemented within one or more separate layers of the computing system, including within a content input processing/analysis component 420, the content recommendation engine 425, and/or the content management and classification system 435. In 20 step 1103, data identifying the matching images determined by the search API may be transmitted back to the content authoring user interface 415 to be integrated within the interface or otherwise presented to the user.

[0076] Thus, upon the completion of the steps 602-608 for a plurality of content resources in a content repository 440, one or more vector spaces 430 may be populated with vectors wherein 25 each vector corresponds to a content item in the repository 440. Additionally, in some embodiments, a separate set of metadata tags may be generated for some or all of the content items, and stored as separate objects from the vectors in vector spaces 430. Such tags may be stored in any data storage or component shown in FIG. 4, or a separate data store, and each tag may be associated with content items in a repository 440, vectors in a vector space 430, or both.

[0077] Referring now to FIG. 12, another flowchart is shown illustrating a second process for receiving original authored content from a user via the client device 410, extracting features and/or tags from the content in real-time (or near real-time) during the user's authoring session, vectorizing the authored content (also in real-time or near real-time), and comparing the vector
5 of the original authored content to one or more existing vector spaces 430 in order to identify and retrieve related/associated content from one or more available content repositories 440. The steps in this process also may be performed by one or more components within the computing environment 400, such as the content recommendation engine 425 in collaboration with the client device 410, input processing/analysis component 420 and retrieval/embedding component
10 445, and the various subsystems and subcomponents implemented therein.

[0078] In step 1202, original authored content may be received from a user via the client device 410. As discussed above, the original authored content may correspond to text typed by the user, a new image created or imported by the user, new audio or video input recorded or imported by the user, a new graphic created by the user, etc. Thus, step 1202 may be similar to
15 step 602, discussed above. However, whereas the content in step 602 may be previously authored/stored content that is retrieved from a repository 440, in step 1202 the content may be newly authored content received via a user interface, such as a web-based text input control, image importer control, image creation control, audio/video creation control, etc.

[0079] In step 1204, the content received in step 1202 (e.g., original authored content) may be
20 processed, for example, by the input processing/analysis component 420. Step 1204 may be similar or identical to step 604, discussed above, with respect to the parsing steps, processing steps, keyword/data feature extraction steps, etc. For example, for text input received in step 1202 (e.g., blog posts, letters, emails, articles, etc.), the processing in step 1204 may include text parsing, identification of keywords, stemming, synonym analysis/retrieval, and the like. In other
25 examples, when an image is received in step 1202 (which may be uploaded by the user, imported from another system, and/or created or modified manually by the user via the content author user interface 415), then step 1204 may include using AI-based image classification tools as described above to identify particular image features and/or generate image tags. (See FIG. 7). These analyses in step 1204 may be performed via REST-based services or other web services, using
30 analytics, machine-learning algorithms, and/or AI, such as an AI-based cognitive image analysis

service, and/or AI/REST cognitive text services. Similar techniques/services may be used in step 1204 for other types of content items, such as video files, audio files, graphics, or social media posts, where a specialized web-service of the system 400 is used to extract and analyze specific features (e.g., words, objects within images/video, facial expressions, etc.) depending on the media type of the content item. Step 1204 may also include any of the tagging processes described herein for tagging of text blocks, images, audio/video data, and/or other content with tags corresponding to any identified content topics, categories, or features.

[0080] In step 1206, one or more vectors, compatible with one or more of the available vector spaces 430, may be generated based on the content received in step 1202. Step 1206 may be similar or identical to step 606, discussed above. As discussed above, a vector may be generated based on the particular features within the content (and/or tags) identified in step 1204. The vector generation process in step 1206 may use one or more data transformation techniques, whereby the set of features associated with the originally authored content item may be transformed into a vector compatible with one of the common vector spaces 430. For example, FIG. 13 illustrates a technique by which image input received in step 1202 may be transformed into a feature vector of a predetermined vector format (e.g., 1 x 4096 dimension vector) in step 1206. As shown in FIG. 13, an image may be provided as input to a model, which is configured to extract and learn features in the image (e.g., using convolution, pooling, and other functions) and output a feature vector representing the input image. For example, as described in the paper “Visualizing and Understanding Convolutional Networks,” by Matthew D. Zeiler and Rob Fergus, of New York University, 2014, which is incorporated by reference herein, within convolution neural networks, the initial layers of the neural network may detect simple features from the image like linear edge, in the later layers it detects more complex shapes and patterns. For instance, the first and/or second layers in a convolutional neural network may detect simple edges or patterns, while the later layers can detect the actual complex objects present in the image (a cup, a flower, a dog, etc.) As an example, when receiving and processing an image of a face using convolutional neural network, the first layer may detect the edges in various directions, the second layer may detects different parts (e.g., eyes, nose, etc.) of the given face, and the third layer may get the feature map of the entire face.

[0081] In step 1208, the feature vector generated in step 1206 may be compared to the compatible feature vector space 430 (or spaces 430a-n) that were populated during the processes 602-608, discussed above. For example, an illustrative vector space populated with feature vectors corresponding to images is shown in FIG. 14. In this example, each dot may represent a vectorized image, and the circles (and corresponding dot colors) in FIG. 14 may indicate one of three example tags associated with the images. The image tags in this case are “Coffee,” “Mountain,” and “Bird,” and it should be understood that these tags are not mutually-exclusive (i.e., an image may be tagged with one, two, or all three tags). Additionally, it should be understood that these tags and the layout of the multi-dimensional vector space in FIG. 14 are illustrative only. In various embodiments, there is no limit the number or type of tags that may be used, or to the number of dimensions of a vector space 430.

[0082] To perform the vector space comparison in step 1208, the content recommendation engine 425 may calculate the Euclidean distance between the feature vector generated in step 1206, and each of the other feature vectors stored in the vector space/spaces 430. Based on the distances calculated, the engine 425 may rank feature vectors in ascending order of feature space distance, so that the smaller the distance between the two feature vectors, the higher the rank. Such techniques may allow the content recommendation engine 425 to determine a set of highest ranking feature vectors within the vector space(s) 430, which are most similar in features/characteristics/etc. to the feature vector generated in step 1206 based on the input received in step 1202. In some cases, a predetermined number (N) of the highest-ranking feature vectors may be selected in step 1208 (e.g., the 5 most similar articles, 10 most similar images, etc.), while in other cases all of the feature vectors satisfying a particular closeness threshold (e.g., distance between vectors < threshold(T)) may be selected.

[0083] In some embodiments, the vector comparison in step 1208 may be a “deep feature space” comparison, illustrated in FIG. 15. In these embodiments, the feature vector generated in step 1206 may be compared without respect to any tags or other metadata. In other words, in a deep feature comparison, the feature vector generated in step 1206 may be compared to every other feature stored in the vector space 430. While a deep feature comparison may be guaranteed to find the closest vectors in the vector space 430, this sort of comparison may require additional processing resources and/or additional time to return the vector results. This is

especially true for large vector spaces, which may include thousands or even millions of features vectors, each of which representing a separate content object/resource stored in a repository 440. For instance, to compute the Euclidean distance between two image feature vectors of size 1×4096 requires approximately 10,000 addition and multiplication instructions to be performed by the system 400. Thus, if there are 10,000 images in the repository, 10,000,000 operations must be performed.

[0084] Accordingly, in other embodiments, the vector comparison in step 1208 may be a “filtered feature space” comparison, illustrated in FIGS. 16-17. In filtered feature space comparisons, the vector space may first be filtered based on tags (and/or other properties, such as resource media type, creation date, etc.), to identify subset of the feature vectors within the vector space(s) 430 having tags (and/or other properties) that match those of the feature vector generated in step 1206. Then, the feature vector generated in step 1206 may be compared to only those feature vectors in the subset having matching tags/properties. Accordingly, a filtered feature space comparison may be executed faster and more efficiently than a deep space comparison, although the possibility exists of missing close feature vectors that are filtered out and not compared.

[0085] As noted above, step 1208 may include a comparison of the feature vector generated in step 1206 to a single vector space, or to multiple vector spaces. In some embodiments, the feature vector generated in step 1206 may be compared to the corresponding type of vector space. For instance, when text input is received in step 1202, the resulting feature vector may be compared to the vectors within the topic vector space 430a, and when an image is received as input in step 1202, the resulting feature vector may be compared to the vectors within the image vector space 430b, and so on. In some embodiments, it may be possible to compare feature vectors corresponding to one type of input, to vector spaces contain different types of vectors (e.g., identifying the image resources most closely related to a text-based input, or vice versa). For instance, FIG. 18 represents a process of receiving text input in step 1402, and in step 1408 retrieving both similar images (e.g., the closest-to from an image vector space 430b) and similar articles (e.g., the closest-to from a topic vector space 430b).

[0086] For embodiments that involve retrieving and/or comparing tags associated with content resources, a problem may arise when the tags of one resource are related but do not exactly

match the corresponding tags/keywords/characteristics of another resource. An example of this potential problem is illustrated in FIG. 19, where the keywords extracted from an original authored text content resource are compared to a set of image tags stored for a set of image content resources. In this example, none of the extracted keywords (“Everest,” “Base Camp,” “Summit,” “Mountain,” or “Himalaya”) are exact matches for the image tags (“Mountaineer,” “Cappuccino,” or “Macaw”). In some embodiments, word/phrase parsing and processing techniques such as word stemming, word definition and/or synonym retrieval and analysis, etc., may be used to detect matches among related but non-matching terms. However, these techniques also may fail, even for related keywords/tags. Therefore, in some embodiments, the content processing/analysis component 420 and/or content recommendation engine 425 may perform a word vector comparison to address this problem. As illustrated in the example shown of FIG. 20, the keywords extracted from the text document in FIG. 19 may be analyzed within a 3D word vector space, and the distance between those keywords and each of the image tags may be computed. As shown in FIG. 21, the keyword-to-tag vector space analysis performed in FIG. 20 may determine that the image tag “Mountaineer” sufficiently close within the word vector space to the extracted keywords, and thus should be considered as an image tag match for the filtered feature space comparison.

[0087] Another potential problem that may arise in embodiments of retrieving and/or comparing tags associated with content resources, is caused by homographic keywords and/or resource tags. A homographic word or phrase (or homonym) is one that has different and non-related meaning with the same spelling. An example of homonym image tags is shown in FIG. 22, where a first image is tagged with the word “Crane,” meaning the long-legged and long-necked bird, and a second image tagged with the same word “Crane,” meaning the machine with a projecting arm used for moving heavy objects. In this case, the content processing/analysis component 420 and/or the content recommendation engine 425 may perform a word sense disambiguation process on the two image tags, to determine which sense of the word crane each refers to. In this example, the word sense disambiguation process may initially retrieve a Wordnet database entry (or other definition data) associated with the each tag, as shown in FIG. 22 for the two different “Crane” tags.

[0088] An illustrative word sense disambiguation process is shown in FIGS. 23-24. In this process, other keywords within the authored document may be used, and/or the specific context of the word “Crane” within the document (e.g., description, part of speech, tense, etc.), by the content processing/analysis component 420 and/or the content recommendation engine 425 to
5 determine the most likely meaning of the word crane within the authored text document, and thus which of the “Crane” image tags is related to the authored text document. For example, referring now to FIG. 23, an input text 2301 is shown, from which a number of relevant keywords 2302 have been extracted. The first extracted keyword (“Crane”) may be compared to the image tags within the content repository 440, and in this example two matching tags 2303 have been
10 identified corresponding to two “Crane” tagged images 2304a and 2304b within the content repository 440.

[0089] As shown in FIG. 24, to address this potential problem of word sense ambiguity, the disambiguation process may continue by comparing one or more additional keywords 2302 extracted from the input content 2301, to the other content tags of the two matching images
15 2304a and 2304b. In this example, the additional extracted keywords of “mechanical,” “machine,” “lifting,” and “construction” may be compared to the content tags and/or extracted features associated with each of the images 2304a and 2304b. As shown in FIG. 24, these additional comparisons may disambiguate the initial keyword match of “Crane,” so that avian crane image 2304a is not returned by the content recommendation system 425, but the
20 construction crane image 2304b is returned.

[0090] In other examples, a similar disambiguation process may be performed using image similarity. For instance, the content processing/analysis component 420 and/or the content recommendation engine 425 may identify common image features between the images associated with the authored content (e.g., drawn or authored images) and the two different
25 “Crane” images, in order to determine which crane is the appropriate related image. These disambiguation processes also may be combined in various ways, for example, comparing keywords extracted from the authored text document to the visual features extracted from the images, etc. Therefore, in an authored text document referring to a crane, the related words “boom” and “pulley” might be matched visually to the lower crane image if a boom and pulley
30 can be visually identified within that image. Similarly, if the authored text document refers to a

crane and includes the related words “beak” and “feathers,” then the “crane” keyword might be matched visually to the upper crane image if a beak and feathers can be visually identified within that image.

[0091] Referring now to FIGS. 25-28, an end-to-end example is shown of performing the process in FIG. 12, specifically for the embodiment retrieving a set of relevant images for an article authored by a user via the user interface 415. Initially, at 2501 (FIG. 25), a user types the text for an article into the Demo Editor (AIditor) user interface. At 2502, a number of keywords are extracted from the text of the article, and at 2503 the extracted keywords are compared by an AI Rest Service to the image tags stored for a library of images within an image content repository 440. FIGS. 26 and 27 illustrate the same example process of FIG. 25, but with additional details regarding the operation of the AI Rest Service. As shown in FIG. 26, using the techniques discussed above, an AI Rest Service identifies one or more tags (e.g., “mountaineer” as being related to the authored article. Then, as shown in FIG. 27, in some embodiments a combination of different software services may be used to perform this step, such as first cognitive text REST service used to determine a set of keywords from the text input, and a second internal REST service used to map the keywords to image tags. Each of these services may be implemented within the content recommendation engine 425 and/or via an external service provider. Then, in step 2505 (shown in FIG. 28), the content recommendation engine 425 may transmit the determined image tag(s) to a search API associated with an image content repository 440. In some cases, the search API may be implemented within a cloud-based content hub, for instance an Oracle Content and Experience Cloud (CEC). In step 2506, the Search API may retrieve a set of relevant images based on the tag matching, and in step 2507 the retrieved images (or scaled-down versions of the images) may be transmitted back and embedded within the user interface at 415 (at screen region 2810).

[0092] Although the example shown in FIGS. 25-28 illustrates a specific embodiment of retrieving a set of relevant images for an article authored by a user, it should be understood that the steps in FIG. 12 may be similarly executed to retrieve other types of content as well. For instance, similar steps may be performed to retrieve articles (or other text documents) related to the text input by the user via the user interface 415. Related content resources of other media types (e.g., audio files, video clips, graphics, social media posts, etc.) also may be retrieved in

other embodiments. Additionally, when the user imports/creates other types of input besides text into the user interface (e.g., a drawn or uploaded image, a spoken audio input, a video input, etc.), similar steps may be executed to retrieve related content resources of various different types (e.g., related articles, images, video, audio, social media, etc.), depending on the configuration of the content recommendation engine 425 and/or the user preferences.

[0093] For instance, referring now to FIGS. 29-35, another example embodiment is shown in which the process steps of FIG. 12 are performed to retrieve a set of related articles (or other text content resources), based on the original authored text input received via the user interface 415 (e.g., a user's blog post, email, article, etc.). As shown in FIG. 29, a user has authored a new article via the user interface 415, and a set of article topics has been identified by an AI-based REST service invoked by the content recommendation engine 425. As shown in FIG. 30, the identified article topics may be compared to the topics previously-identified for a set of articles in an article content repository 440. In these examples, FIG. 29 shown according to one embodiment, while Fig. 30 shows another embodiment. Figure 29 is just a subset of Fig-30, we can safely remove Fig-29 from the diagrams. Thus, the article topics may be determined and stored as metadata or other associated data objects, using similar techniques to the processes described above in FIG. 6), in which image features/tags were determined and stored associated with images. Similarly, an article content repository 440 may have metadata or other associated storage including article topics, dates, keywords, authors, publications, etc., for each article stored in the repository 440. In the example shown in FIG. 30, an article relating to deaths on Mount Everest has been identified as potentially related to the user's newly-created article, based on matching of the article topics. FIGS. 31-35 illustrate an end-to-end process of using system 400 to find related articles to a user input article, similar to the steps shown in FIG. 25-28 for finding related images. In step 3101 (FIG. 31), a user creates a new article via the user interface 415. In step 3102, the article text is transmitted to one or software services by the content recommendation engine 425 (e.g., AI-based REST services), and in step 3103 (FIG. 32) the software service(s) use cognitive text service functionality to analyze the text of the article to determine one or more topics of the article. In step 3104, the determined article topics are transmitted back to the content recommendation engine 425, and in step 3105 the recommendation engine 425 sends both the article text and the identified topics to a separate API (e.g., within a cloud-based content hub), where in step 3106, the article may be saved to a

repository 440 for future reference and indexed based on the identified topics. Also in step 3106 (FIG. 33), the existing repository 440 of articles may be searched via a search API to identify potentially related topics based on a topic matching process (FIG. 34). Finally, in step 3107 (FIG. 35), the articles identified as being potentially related to the newly-created article may be transmitted back (in whole, or just links) to be embedded within the user interface 415 (e.g., at user interface region 3510).

[0094] As illustrated in the above example of FIGS. 29-33, certain embodiments described herein may include identification of topics of newly-created text documents, and/or topics of text documents stored within content repositories 440, as well as comparisons and identifications of topic closeness and matches. In various embodiments described herein, various techniques including explicit semantic analyses may be used for text topic evaluation and topic “closeness” techniques. As shown in FIGS. 29-30, in some cases, such techniques may use a large-scale data source (e.g., Wikipedia), to provide fine-grained semantic representation of unrestricted natural language texts, to represent meaning in a high-dimensional space of natural concepts derived from the data sources. For instance, text classification techniques may be used to explicitly represent the meaning of any text in terms of Wikipedia-based concepts. The semantic representation may be the feature vector of the text snippet, that is converted through topic modeling. Wikipedia (or another large-scale data source) may be used to include a larger vocabulary (e.g., bag of words) to the system, to cover a large universe of words. Wikipedia-based concepts may be the titles of the Wikipedia pages which are used as the classes/categories while classifying a given text snippet. Given a text snippet, the nearest Wikipedia page title may be returned (e.g., “Mount Everest,” “Stephen Hawking,” “Car Accident,” etc.), which may be used as the class/category of the text. The effectiveness such techniques may be evaluated automatically by computing the degree of semantic relatedness between fragments of natural language text.

[0095] In these text classification/relatedness evaluation techniques, one advantage of using large-scale publically available knowledge sources (e.g., Wikipedia or other encyclopedias) is access to vast amounts of highly organized human knowledge pre-encoded into the publically available sources, which are in constant change/development. Machine learning techniques may be used based on Wikipedia and/or other sources to build a semantic interpreter that maps

fragments of natural language text into a weighted sequence of Wikipedia concepts ordered by their relevance to the input. Input texts therefore may be represented as weighted vectors of concepts, called interpretation vectors. The meaning of a text fragment is thus interpreted in terms of its affinity with a host of Wikipedia concepts. Semantic relatedness of texts then may
5 be computed by comparing their vectors in the space defined by the concepts, for example, using the cosine metric. Such a semantic analysis may be explicit in the sense the manifest concepts may be grounded in human cognition. Because, the user input may be received via the user interface 415 as plain text, conventional text classification algorithms may be used to rank the concepts represented by these articles according to their relevance to the given text fragment.
10 Thus, an online encyclopedia (e.g., Wikipedia) may be used directly, without the need for deep language understanding or pre-cataloged common-sense knowledge. In some embodiments, each Wikipedia concept may be represented as an attribute vector of words that occur in the corresponding article. Entries of these vectors may be assigned weights using, for example, a term frequency-inverse document frequency (TFIDF) scheme. These weights may quantify the
15 strength of association between words and concepts. To speed up semantic interpretation, an inverted index, which maps each word into a list of concepts in which it appears, may be used. The inverted index also may be used to discard insignificant associations between words and concepts, by removing those concepts whose weights for a given word are below a certain threshold. A semantic interpreter may be implemented as a centroid-based classifier, which may
20 rank Wikipedia concepts by relevance, based on a received text fragment. For instance, a semantic interpreter within the content recommendation engine may receive an input text fragment T, and represent the fragment as a vector (e.g., using the TFIDF scheme). The semantic interpreter may iterate over the text words, retrieve corresponding entries from the inverted index, and merges them into a weighted vector. Entries of the weighted vector may
25 reflect the relevance of the corresponding concepts to text T. To compute semantic relatedness of a pair of text fragments, their vectors may be compared using, e.g., the cosine metric.

[0096] In other examples, similar methods for generating features for text categorization may include supervised learning tasks, in which words occurring in the training documents may be used features. Therefore, in some examples Wikipedia concepts may be used to augment the bag
30 of words. On the other hand, computing semantic relatedness of a pair of texts is essentially a “one-off” task, therefore, the bag of words representation may be replaced with a representation

based on concepts. These and other related techniques are described in the more detail in the paper "Computing Semantic Relatedness using Wikipedia-based Explicit Semantic Analysis," by Evgeniy Gabrilovich and Shaul Markovitch, Department of Computer Science Technion – Israel Institute of Technology, which is fully incorporated by reference herein for all purposes, as well
5 as the other related discussed therein. Using the techniques described in this paper and others, with a filtered subset of Wikipedia, for each article there may be one concept which is the title of the article. When the content recommendation engine 425 receives a text document via user interface 415, the text may first be summarized. In these cases, each unique word in the text, after the stop words are removed and the words stemmed, may be given a weight based on
10 frequency and inverse frequency of the word in the article. Each word may then be compared to see what Wikipedia articles (concepts) it appears in, and thus for that word the content recommendation engine 425 may create a concept vector. The concept vectors for all the words in the text document can be combined to form a weighted concept vector for the text document. The content recommendation engine 425 then may measure the similarity between each word
15 concept vector and the text concept vector. All words above a certain threshold then may be selected as the "keywords" for the document.

[0097] With reference now to FIG. 36, an example semantic text analyzer system 3600 is shown, illustrating the techniques used by the analyzer system 3600 to perform semantic text summarization in certain embodiments described above. Such a system 3600 may be
20 incorporated into and/or separate from by accessed by the content recommendation engine 425 in various implementations.

[0098] In some implementations, explicit semantic analysis for computing semantic relatedness is repurposed for computing a text summary of a given text document. More specifically, the text summary is derived based on word embeddings. In other words, the context
25 of n-grams (e.g., words) is captured for purpose of determining semantic similarity, in contrast to typical similarity measures such as cosine on bag of words or edit distance on strings.

[0099] The given text document can be an article, a web page, or other piece text for which a text summary is desired. As with the classification approach described herein, the text is not limited to written language but may include other human-readable symbols, numbers, charts,
30 tables, equations, formulas, etc.

[0100] The text summary approach using explicit semantic analysis operates generally as follows: (1) grammatical units (e.g., sentences or words) are extracted from the given text document using any known technique for identifying and extracting such units, (2) each of the extracted grammatical units and the text document are represented as weighted vectors of
 5 knowledge base concepts, (3) the semantic relatedness between the text document as a whole and each grammatical unit is computed using the weighted vectors, and (4) one or more of the grammatical units most semantically relatedly to the text document as a whole are selected for inclusion in a text summary of the text document. In some cases, representation of weighted vectors of knowledge base concepts may correspond to topic modeling, in which each sentence
 10 or word first may be converted into a feature vector, and then the feature difference/similarity may be calculated in the high-dimensional vector space. There may be various methods to convert the words to vectors, for example, WORD2VEC or a Latent Dirichlet Allocation.

[0101] FIG. 36 illustrates text summarization using explicit semantic analysis. First a text summarizer is built based on a knowledge base 3602. The knowledge base 3602 can be general
 15 or domain-specific. An example of a general knowledge base is a collection of encyclopedia articles such as a collection of Wikipedia articles or other encyclopedic collection of text articles. However, the knowledge base 3602 can instead be domain-specific such as a collection of text articles specific to a particular technical field such as a collection of medical, scientific, engineering, or financial articles.

20 [0102] Each article of the knowledge base 3602 is represented as an attribute vector of n-grams (e.g., words) that occur in the article. Entries in the attribute vector are assigned weights. For example, the weights may be used using a term frequency-inverse document frequency scoring scheme. The weights in the attribute vector for an article quantify the strength of association between the n-grams (e.g., words) of the article and the article as a concept.

25 [0103] In some implementations, the term frequency-inverse document frequency scoring scheme computes a weight for a given n-gram t of a given article document d as represented by the following equation:

$$\frac{(k + 1)tf_{t,d}}{tf_{t,d} + k(1 - b + b \frac{L_d}{L_{avg}})} \cdot \log \frac{M - df_t + 0.5}{df_t + 0.5}$$

[0104] Here, $tf_{t,d}$ represents the frequency of n-gram t in document d . And df_t represents the document frequency of n-gram t in the knowledge base 3602. M represents the total number of documents in the training set, L_d represents the length of the document d in number of terms, and
5 L_{avg} represents the average length in the training corpus, and K and b are free parameters. In some implementations, k is approximately 1.5 and b is approximately 0.75.

[0105] The preceding is one example of a term frequency-inverse document frequency scoring scheme that may be used to weight attributes of an attribute vector. Other statistical measures that reflect how important an attribute (e.g., an n-gram) is to an article in the knowledge base
10 3602 may be used. For example, other TF/IDF variations such as BM25F that takes anchor text into account may be used with certain types knowledge bases such as, for example, a knowledge base of web pages or other set of hyperlinked documents.

[0106] A weighted inverted index builder computer 3604 builds a weighted inverted index 3606 from the attribute vectors representing the articles of the knowledge base 3602. The
15 weighted inverted index 3606 maps each distinct n-gram represented in the set of attribute vectors to a concept vector of concepts (articles) in which the n-gram appears. Each concept in the concept vector may be weighted according to the strength of association between the concept and the n-gram to which the concept vector is mapped by the weighted inverted index 3606. In some implementations, the indexer computer 3604 uses the inverted index 3606 to discard
20 insignificant associations between n-grams and concepts by removing from the concept vectors those concepts whose weights for a given n-gram are below a threshold.

[0107] To generate a text summary of a given text document 3610, grammatical units 3608 are extracted from a given text document 3610 and the semantic relatedness between each
grammatical unit and the given text document 3610 is computed. A number of grammatical units
25 having a high degree of semantic relatedness to the given text document 3610 are selected for inclusion in the text summary.

[0108] The number of grammatical units that are selected for inclusion in the text summary may vary based on a variety of different factors. One approach is to select a predefined number of grammatical units. For example, the predefined number may be configured by a user of the
30 system or learned by a machine learning process. Another approach is to select all grammatical

units having a degree of semantic relatedness to the given text document 3610 that is above a predefined threshold. The predefined threshold can be configured by a user of the system or learned by a machine learning process. Yet another possible approach is determine the grammatical unit having the highest degree of semantic relatedness to the given text document
5 3610 and then selecting all other grammatical units where the difference in the degree of semantic relatedness to the given text document 3610 of the grammatical unit and the highest degree is below a predefined threshold. The grammatical unit having the highest degree and any other grammatical units below the predefined threshold are selected for inclusion in the text summary. Again, the predefined threshold can be configured by a user of the system or learned
10 by a machine learning process.

[0109] In some implementations, a grammatical unit having the highest or a relatively high degree of semantic relatedness to the given text document 3610 is not always selected for inclusion in the text summary. For example, a first grammatical unit that has a lower degree of semantic relatedness to the given text document 3610 than a second grammatical unit may be
15 selected for inclusion in the text summary and the second grammatical unit may not be selected for inclusion in the text summary if the first grammatical unit is not sufficiently dissimilar with respect to the grammatical units that have already been selected for inclusion in the text summary. The extent of a grammatical unit's dissimilarity with respect to the existing text summary can be measured in a variety of different ways such as, for example, by using lexical
20 approach, a probabilistic approach, or a hybrid of the lexical approach and the probabilistic approach. Using a dissimilarity measure to select grammatical units for inclusion in the text summary can prevent multiple similar grammatical units from being included in the same text summary.

[0110] In some implementations, other techniques for selecting a number of grammatical units
25 for inclusion in the text summary as a function of the units' semantic relatedness to the given text document 3610 and its dissimilarity relative to one or more of the other units may be used and is not limited to any particular technique. For example, given a number of grammatical units with a semantic relatedness to the given text document 3610 above a threshold, the dissimilarity of each composite grammatical unit with respect the combination of the number of grammatical units
30 may be measured and a number of grammatical units that are most dissimilar to each other may

be selected for inclusion in the text summary. As a result, the grammatical units selected for inclusion in the text summary are highly semantically related to the text document, as a whole, yet dissimilar from each other. This is a more useful text summary than one containing highly semantically related but similar grammatical units because similar grammatical units are more likely to be redundant of each other in terms of the information conveyed by the grammatical units than dissimilar grammatical units.

[0111] Another possibility is to compute a composite similarity/dissimilarity measure for grammatical units and then select grammatical units for inclusion in the text summary based on their composite scores. For example, the composite measure might be a weighted average of the semantic relatedness measure and a dissimilarity measure. For example, a possible composite measure computed as a weighted average is:

$$(a * Similarity) + (b * Dissimilarity)$$

[0112] Here, the parameter *Similarity* represents the semantic relatedness of a grammatical unit to the input text 3610 as a whole. For example, the parameter *Similarity* could be the similarity estimate 3620 computed for the grammatical unit. The parameter *Dissimilarity* represents a dissimilarity measure of the dissimilarity of the grammatical unit to the set of one or more grammatical units. For example, the set of one or more grammatical units can be a set of one or more grammatical units that are already selected for inclusion in the text summary. The parameter *a* represents the weight applied to the similarity measure in the weighted average. The parameter *b* represents the weight application to the dissimilarity measure in the weighted average. The composite measure effectively balances the similarity measure and the dissimilarity measure against each other. They can be balanced against each other equally (e.g., $a = 0.5$ and $b = 0.5$). Alternatively, the similarity measure may be given more weight (e.g., $a = 0.8$ and $b = 0.2$).

[0113] A grammatical unit extracted from the given text document can be a sentence, a phrase, a paragraph, a word, an n-gram, or other grammatical unit. In the case, where the grammatical units 3608 extracted from the given text document 3610 are words or n-grams, the process may be viewed as keyword generation, as opposed to text summarization.

[0114] The text summarizer 3612 accepts a piece of text. The piece of text is the given text document 3610 or a grammatical unit thereof. The piece of text is represented as an “input” vector of weighted attributes (e.g., words or n-grams) of the piece of text. Each weight in the input vector is for a corresponding attribute (e.g., word or n-gram) identified in the piece of text and represents the strength of association between the piece of text and the corresponding attribute. For example, the weights may be calculated according to a TF-IDF scheme or the like.

[0115] In some implementations, the weight of an attribute in the input vector is computed as follows:

$$\frac{(k + 1)tf_{t,d}}{tf_{t,d} + k \left(1 - b + b \frac{L_d}{L_{avg}} \right)}$$

10

[0116] Here, $tf_{t,d}$ is the frequency of n-gram t in the piece of text d . Parameters k , b , L_d , and L_{avg} are as before except with respect to knowledge base 3602 instead of a classification training set. In some implementations, k is approximately 1.5 and b is approximately 0.75.

[0117] It should be noted that other weighting schemes are possible and embodiments are not limited to any particular weighting scheme when forming the input vector. Forming the input vector may also include unit length normalization such as described above with respect to the training data item vectors.

[0118] The text summarizer 3612 iterates over the non-zero weighted attributes of the input vector formed based on the piece of text, retrieves attribute vectors corresponding to the attributes from the weighted inverted index 3606, and merges the retrieved attribute vectors into a weighted vector of concepts that represents the piece of text. This weighted vector of concepts is referred to hereinafter as a “concept” vector.

[0119] The attribute vectors retrieved from the weighted inverted index 3606 corresponding to the attributes of the input vector are also each a vector of weights. However, the weights in the attribute vector quantify the strength of association between respective concepts of the knowledge base 3602 and the attribute mapped to the attribute vector by the inverted index 3606.

25

[0120] The text summarizer 3612 creates a concept vector for the piece of text. The concept vector is a vector of weights. Each weight in the concept vector represents the strength of association between a respective concept of the knowledge base 3602 and the piece of text. A concept weight in the concept vector is computed by the text summarizer 3612 as a sum of values, one value for each attribute non-zero weighted in the input vector. Each value for an attribute of the sum is computed as the product of (a) the weight of the attribute in the input vector and (b) the weight of the concept in the attribute vector for the attribute. Each concept weight in the concept vector reflects the relevance of the concept to the piece of text. In some implementations, the concept vector is normalized. For example, the concept vector may be normalized for unit length or concept length (e.g., like class length above).

[0121] The text summarizer 3612 may generate a concept vector 3616 for the input text 3610 and a concept vector 3614 for each of the grammatical units 3608. The vector comparator 3618 compares a concept vector 3614 generated for a grammatical unit to the concept vector 3616 generated for the input text 3610 using a similarity measure to generate a similarity estimate 3620. In some implementations, a cosine similarity measure is used. Implementations are not limited to any particular similarity measure and any similarity measure capable of measuring the similarity between two non-zero vectors may be used.

[0122] The similarity estimate 3620 quantifies the degree of semantic relatedness between a grammatical unit and the input text 3610 from which the grammatical unit was extracted. For example, the similarity estimate 3620 may be a value between 1 and 0 inclusive with values closer to 1 representing a higher degree of semantic relatedness and values closer to 0 representing a lower degree of semantic relatedness.

[0123] A similarity estimate 3620 may be computed for each of the grammatical units 3608. The similarity estimates 3620 generated for the grammatical units 3608 may be used to select one or more of the grammatical units 3608 for inclusion in a text summary of the input text 3610 (or to select one or more keywords for keyword generation for the input text 3610).

[0124] There a variety of applications of the above technique for text summarization to provide an accurate text summary of longer text such as, for example, news stories, blog posts, journal articles, web pages, etc.

[0125] In any or all the above embodiments, after one or more content resources (e.g., images, articles, etc.) have been identified as potentially related to the content currently being created by the user via the user interface 415, the related content resources are transmitted back to the content recommendation engine 425, where they may be retrieved, modified, and embedded into the user interface 415, for example, by the content retrieval/embedding component 445. Using the retrieval/embedding component 445, the potentially related content resources may be provided to the user via the user interface 415 in such a way that they may be optionally selected to be included with the content currently being created. Two example user interfaces are shown in FIGS. 37 and 38, in which image recommendations are provided to a user during the creation of content. In FIG. 37, a media recommendation pane is shown including images selected based on the text of the content currently being authored by the user via the user interface. In FIG. 38, a visual-feature analysis has been used to select a set of images that are potentially related to a first image ("filename.JPG) selected by the user. Similar techniques and user interface screens may be used to allow users to select, drag-and-drop, and embed images, links to articles and other text documents, audio/video files, etc. into the content currently being created by the user.

[0126] FIG. 39 depicts a simplified diagram of a distributed system 3900 in which the various examples discussed above can be implemented. In the illustrated example, the distributed system 3900 includes one or more client computing devices 3902, 3904, 3906, 3908, coupled to a server 3912 via one or more communication networks 3910. The client computing devices 3902, 3904, 3906, 3908 may be configured to run one or more applications.

[0127] In various embodiments, server 3912 may be adapted to run one or more services or software applications that enable one or more operations associated with the content recommendation system 400. For example, users may use the client computing devices 3902, 3904, 3906, 3908 (e.g., corresponding to content author device 410) to access one or more cloud-based services provided by via the content recommendation engine 425.

[0128] In various examples, the server 3912 may also provide other services or software application, and can include non-virtual and virtual environments. In some examples, these services may be offered as web-based or cloud services or under a Software as a Service (SaaS) model to the users of the client computing devices 3902, 3904, 3906, 3908. Users operating the

client computing devices 3902, 3904, 3906, 3908 may in turn use one or more client applications to interact with the server 3912 to use the services provided by these components.

[0129] In the configuration depicted in FIG. 39, the server 3912 may include one or more components 3918, 3920, 3922 that implement the functions performed by the server 3912. These components may include software components that may be executed by one or more processors, hardware components, or combinations thereof. It should be appreciated that various different system configurations are possible, which may be different from the example distributed system 3900.

[0130] The client computing devices 3902, 3904, 3906, 3908 may include various types of computing systems, such as portable handheld devices such as smartphones and tablets; general purpose computers such as personal computers and laptops; workstation computers; wearable devices such as a head-mounted display; gaming systems such as handheld gaming devices, gaming consoles, and Internet-enabled gaming devices; thin clients; various messaging devices; sensors and other sensing devices; and the like. These computing devices may run various types and versions of software applications and operating systems (e.g., Microsoft Windows[®], Apple Macintosh[®], UNIX[®] or UNIX-like operating systems, Linux or Linux-like operating systems such as Google Chrome[™] OS) including various mobile operating systems (e.g., Microsoft Windows Mobile[®], iOS[®], Windows Phone[®], Android[™], BlackBerry[®], Palm OS[®]). The client devices may be capable of executing various different applications such as various Internet-related apps, communication applications (e.g., e-mail applications, short message service (SMS) applications) and may use various communication protocols. A client device may provide an interface that enables a user of the client device to interact with the client device. The client device may also output information to the user via this interface. Although FIG. 39 depicts only four client computing devices, any number of client computing devices may be supported.

[0131] Network(s) 3910 in the distributed system 3900 may be any type of network familiar to those skilled in the art that can support data communications using any of a variety of available protocols, including without limitation TCP/IP (transmission control protocol/Internet protocol), SNA (systems network architecture), IPX (Internet packet exchange), AppleTalk, and the like. Merely by way of example, the network(s) 3910 can be a local area network (LAN), networks based on Ethernet, Token-Ring, a wide-area network, the Internet, a virtual network, a virtual

private network (VPN), an intranet, an extranet, a public switched telephone network (PSTN), an infra-red network, a wireless network (e.g., a network operating under any of the Institute of Electrical and Electronics (IEEE) 802.11 suite of protocols, Bluetooth®, and/or any other wireless protocol), and/or any combination of these and/or other networks.

5 **[0132]** The server 3912 may be composed of one or more general purpose computers, specialized server computers (including, by way of example, PC (personal computer) servers, UNIX® servers, mid-range servers, mainframe computers, rack-mounted servers, etc.), server farms, server clusters, or any other appropriate arrangement and/or combination. The server 3912 can include one or more virtual machines running virtual operating systems, or other computing
10 architectures involving virtualization, such as one or more flexible pools of logical storage devices that can be virtualized to maintain virtual storage devices for the server. In various examples, the server 3912 may be adapted to run one or more services or software applications that perform the operations as described above.

[0133] The server 3912 may run an operating system including any of those discussed above,
15 as well as any commercially available server operating system. The server 3912 may also run any of a variety of additional server applications and/or mid-tier applications, including HTTP (hypertext transport protocol) servers, FTP (file transfer protocol) servers, CGI (common gateway interface) servers, JAVA® servers, database servers, and the like. Examples of database servers include, without limitation, those commercially available from Oracle, Microsoft,
20 Sybase, IBM (International Business Machines), and the like.

[0134] In some implementations, the server 3912 may include one or more applications to analyze and consolidate data feeds and/or event updates received from users of the client computing devices 3902, 3904, 3906, 3908. As an example, data feeds and/or event updates may include, but are not limited to, Twitter® feeds, Facebook® updates or real-time updates received
25 from one or more third-party information sources and continuous data streams, which may include real-time events related to sensor data applications, financial tickers, network performance measuring tools (e.g., network monitoring and traffic management applications), clickstream analysis tools, automobile traffic monitoring, and the like. The server 3912 may also include one or more applications to display the data feeds and/or real-time events via one or
30 more display devices of the client computing devices 3902, 3904, 3906, 3908.

[0135] The distributed system 3900 may also include one or more data repositories 3914, 3916. These data repositories may provide a mechanism for storing information various types of information, such as the information described by the various examples discussed above. The data repositories 3914, 3916 may reside in a variety of locations. For example, a data repository
5 used by the server 3912 may be local to server 3912 or may be remote from server 3912 and in communication with server 3912 via a network-based or dedicated connection. The data repositories 3914, 3916 may be of different types. In some examples, a data repository used by the server 3912 may be a database, for example, a relational database, such as databases provided by Oracle Corporation® and other vendors. One or more of these databases may be adapted to
10 enable storage, update, and retrieval of data to and from the database in response to SQL-formatted commands.

[0136] In some examples, one or more of the data repositories 3914, 3916 may also be used by applications to store application data. The data repositories used by applications may be of different types such as, for example, a key-value store repository, an object store repository, or a
15 general storage repository supported by a file system.

[0137] In some examples, a cloud environment may provide one or more services such as those discussed above. FIG. 40 is a simplified block diagram of one or more components of a system environment 4000 in which these and other services can be offered as cloud services. In the example illustrated in in FIG. 40, a cloud infrastructure system 4002 may provide one or
20 more cloud services that may be requested by users using one or more client computing devices 4004, 4006, and 4008. The cloud infrastructure system 4002 may include one or more computers and/or servers that may include those described above for server 3912 of FIG. 39. The computers in cloud infrastructure system 4002 of FIG. 40 may be organized as general purpose computers, specialized server computers, server farms, server clusters, or any other appropriate arrangement
25 and/or combination.

[0138] Network(s) 4010 may facilitate communication and exchange of data between the clients 4004, 4006, 4008 and the cloud infrastructure system 4002. The network(s) 4010 may include one or more networks. The networks may be of the same or different types. The network(s) 4010 may support one or more communication protocols, including wired and/or
30 wireless protocols, for facilitating the communications.

[0139] The example depicted in FIG. 40 is only one example of a cloud infrastructure system and is not intended to be limiting. It should be appreciated that, in other examples, the cloud infrastructure system 4002 may have more or fewer components than those depicted in FIG. 40, may combine two or more components, or may have a different configuration or arrangement of components. For example, although FIG. 40 depicts three client computing devices, any number of client computing devices may be supported in other examples.

[0140] The term cloud service is generally used to refer to a service that is made available to users on demand and via a communication network such as the Internet by systems (e.g., the cloud infrastructure system 4002) of a service provider. Typically, in a public cloud environment, servers and systems that make up the cloud service provider's system are different from the customer's own on-premises servers and systems. The cloud service provider's systems are managed by the cloud service provider. Customers can thus avail themselves of cloud services provided by a cloud service provider without having to purchase separate licenses, support, or hardware and software resources for the services. For example, a cloud service provider's system may host an application, and a user may, via the Internet, on demand and self-service, order and use the application without the user having to buy infrastructure resources for executing the application. Cloud services are designed to provide easy, scalable access to applications, resources and services. Several providers offer cloud services. For example, several cloud services are offered by Oracle Corporation® of Redwood Shores, California, such as middleware services, database services, Java cloud services, and others.

[0141] In various examples, the cloud infrastructure system 4002 may provide one or more cloud services using different models, such as under a Software as a Service (SaaS) model, a Platform as a Service (PaaS) model, an Infrastructure as a Service (IaaS) model, and others, including hybrid service models. The cloud infrastructure system 4002 may include a suite of applications, middleware, databases, and other resources that enable provision of the various cloud services.

[0142] A SaaS model enables an application or software to be delivered to a customer over a communication network like the Internet, as a service, without the customer having to buy the hardware or software for the underlying application. For example, a SaaS model may be used to provide customers access to on-demand applications that are hosted by the cloud infrastructure

system 4002. Examples of SaaS services provided by Oracle Corporation® include, without limitation, various services for human resources/capital management, customer relationship management (CRM), enterprise resource planning (ERP), supply chain management (SCM), enterprise performance management (EPM), analytics services, social applications, and others.

5 [0143] An IaaS model is generally used to provide infrastructure resources (e.g., servers, storage, hardware and networking resources) to a customer as a cloud service to provide elastic compute and storage capabilities. Various IaaS services are provided by Oracle Corporation®.

[0144] A PaaS model is generally used to provide, as a service, platform and environment resources that enable customers to develop, run, and manage applications and services without
10 the customer having to procure, build, or maintain such resources. Examples of PaaS services provided by Oracle Corporation® include, without limitation, Oracle Java Cloud Service (JCS), Oracle Database Cloud Service (DBCS), data management cloud service, various application development solutions services, and others.

[0145] In some examples, resources in the cloud infrastructure system 4002 may be shared by
15 multiple users and dynamically re-allocated per demand. Additionally, resources may be allocated to users in different time zones. For example, the cloud infrastructure system 4002 may enable a first set of users in a first time zone to utilize resources of the cloud infrastructure system for a specified number of hours and then enable the re-allocation of the same resources to another set of users located in a different time zone, thereby maximizing the utilization of
20 resources.

[0146] The cloud infrastructure system 4002 may provide the cloud services via different deployment models. In a public cloud model, the cloud infrastructure system 4002 may be owned by a third party cloud services provider and the cloud services are offered to any general public customer, where the customer can be an individual or an enterprise. In certain other
25 embodiments, under a private cloud model, the cloud infrastructure system 4002 may be operated within an organization (e.g., within an enterprise organization) and services provided to customers that are within the organization. For example, the customers may be various departments of an enterprise such as the Human Resources department, the Payroll department, etc. or even individuals within the enterprise. In certain other embodiments, under a community
30 cloud model, the cloud infrastructure system 4002 and the services provided may be shared by

several organizations in a related community. Various other models such as hybrids of the above mentioned models may also be used.

[0147] The client computing devices 4004, 4006, 4008 may be devices similar to those described above for the client computing devices 3902, 3904, 3906, 3908 of FIG. 39. The client computing devices 4004, 4006, 4008 of FIG. 40 may be configured to operate a client application such as a web browser, a proprietary client application (e.g., Oracle Forms), or some other application, which may be used by a user of the client computing device to interact with the cloud infrastructure system 4002 to use services provided by the cloud infrastructure system 4002.

[0148] In various examples, the cloud infrastructure system 4002 may also provide “big data” and related computation and analysis services. The term “big data” is generally used to refer to extremely large data sets that can be stored and manipulated by analysts and researchers to visualize large amounts of data, detect trends, and/or otherwise interact with the data. The analysis the cloud infrastructure system 4002 can perform may involve using, analyzing, and manipulating large data sets to detect and visualize various trends, behaviors, relationships, etc. within the data. This analysis may be performed by one or more processors, possibly processing the data in parallel, performing simulations using the data, and the like. The data used for this analysis may include structured data (e.g., data stored in a database or structured according to a structured model) and/or unstructured data (e.g., data blobs (binary large objects)).

[0149] As depicted in the embodiment in FIG. 40, cloud infrastructure system 4002 may include infrastructure resources 4030 that are used for facilitating the provision of various cloud services offered by cloud infrastructure system 4002. Infrastructure resources 4030 may include, for example, processing resources, storage or memory resources, networking resources, and the like.

[0150] In some examples, to facilitate efficient provisioning of these resources for supporting the various cloud services provided by cloud infrastructure system 4002 for different customers, the resources may be bundled into sets of resources or resource modules (also referred to as “pods”). Each resource module or pod may comprise a pre-integrated and optimized combination of resources of one or more types. In some examples, different pods may be pre-provisioned for different types of cloud services. For example, a first set of pods may be provisioned for a

database service, a second set of pods, which may include a different combination of resources than a pod in the first set of pods, may be provisioned for Java service, and the like. For some services, the resources allocated for provisioning the services may be shared between the services.

5 [0151] The cloud infrastructure system 4002 may itself internally use the services 4032 that are shared by different components of cloud infrastructure system 4002 and which facilitate the provisioning of services by cloud infrastructure system 4002. These internal shared services may include, without limitation, a security and identity service, an integration service, an enterprise repository service, an enterprise manager service, a virus scanning and white list service, a high
10 availability, backup and recovery service, service for enabling cloud support, an email service, a notification service, a file transfer service, and the like.

[0152] In various examples, the cloud infrastructure system 4002 may include multiple subsystems. These subsystems may be implemented in software, or hardware, or combinations thereof. As depicted in FIG. 40, the subsystems may include a user interface subsystem w that
15 enables users or customers of cloud infrastructure system 4002 to interact with cloud infrastructure system 4002. The user interface subsystem 4012 may include various different interfaces such as a web interface 4014, an online store interface 4016 where cloud services provided by cloud infrastructure system 4002 are advertised and are purchasable by a consumer, and other interfaces 4018. For example, a customer may, using a client device, request (service
20 request 4034) one or more services provided by cloud infrastructure system 4002 using one or more of interfaces 4014, 4016, 4018. For example, a customer may access the online store, browse cloud services offered by cloud infrastructure system 4002, and place a subscription order for one or more services offered by cloud infrastructure system 4002 that the customer wishes to subscribe to. The service request may include information identifying the customer and
25 one or more services that the customer desires to subscribe to. For example, a customer may place a subscription order for services such as those discussed above. As part of the order, the customer may provide information identifying the amount of resources the customer needs and/or for what time frame, among other things.

[0153] In some examples, such as the example depicted in FIG. 40, the cloud infrastructure
30 system 4002 may include an order management subsystem (OMS) 4020 that is configured to

process the new order. As part of this processing, OMS 4020 may be configured to: generate an account for the customer, if not done already; receive billing and/or accounting information from the customer that is to be used for billing the customer for providing the requested service to the customer; verify the customer information; upon verification, book the order for the customer;
5 and orchestrate various workflows to prepare the order for provisioning, among other operations.

[0154] Once properly validated, OMS 4020 may then invoke the order provisioning subsystem (OPS) 4024 that is configured to provision resources for the order including processing, memory, and networking resources. The provisioning may include allocating resources for the order and configuring the resources to facilitate the service requested by the customer order. The manner in
10 which resources are provisioned for an order and the type of the provisioned resources may depend upon the type of cloud service that has been ordered by the customer. For example, according to one workflow, OPS 4024 may be configured to determine the particular cloud service being requested and identify a number of pods that may have been pre-configured for that particular cloud service. The number of pods that are allocated for an order may depend
15 upon the size/amount/level/scope of the requested service. For example, the number of pods to be allocated may be determined based upon the number of users to be supported by the service, the duration of time for which the service is being requested, and the like. The allocated pods may then be customized for the particular requesting customer for providing the requested service.

[0155] The cloud infrastructure system 4002 may send a response or notification 4044 to the requesting customer to indicate when the requested service is now ready for use. In some instances, information (e.g., a link) may be sent to the customer that enables the customer to start using and availing the benefits of the requested services.

[0156] The cloud infrastructure system 4002 may provide services to multiple customers. For
25 each customer, the cloud infrastructure system 4002 is responsible for managing information related to one or more subscription orders received from the customer, maintaining customer data related to the orders, and providing the requested services to the customer. The cloud infrastructure system 4002 may also collect usage statistics regarding a customer's use of subscribed services. For example, statistics may be collected for the amount of storage used, the
30 amount of data transferred, the number of users, and the amount of system up time and system

down time, and the like. This usage information may be used to bill the customer. Billing may be done, for example, on a monthly cycle.

[0157] The cloud infrastructure system 4002 may provide services to multiple customers in parallel. The cloud infrastructure system 4002 may store information for these customers, including possibly proprietary information. In some examples, the cloud infrastructure system 4002 includes an identity management subsystem (IMS) 4028 that is configured to manage customers information and provide the separation of the managed information such that information related to one customer is not accessible by another customer. IMS 4028 may be configured to provide various security-related services such as identity services, such as information access management, authentication and authorization services, services for managing customer identities and roles and related capabilities, and the like.

[0158] FIG. 41 illustrates an example of a computer system 4100 that may be used to implement the various examples discussed above. In some examples, the computer system 4100 may be used to implement any of the various servers and computer systems described above. As shown in FIG. 41, the computer system 4100 includes various subsystems including a processing subsystem 4104 that communicates with a number of other subsystems via a bus subsystem 4102. These other subsystems may include a processing acceleration unit 4106, an I/O subsystem 4108, a storage subsystem 4118 and a communications subsystem 4124. The storage subsystem 4118 may include non-transitory computer-readable storage media 4122 and a system memory 4110.

[0159] The bus subsystem 4102 provides a mechanism for letting the various components and subsystems of the computer system 4100 communicate with each other as intended. Although the bus subsystem 4102 is shown schematically as a single bus, alternate examples of the bus subsystem may utilize multiple buses. The bus subsystem 4102 may be any of several types of bus structures including a memory bus or memory controller, a peripheral bus, and a local bus using any of a variety of bus architectures. For example, such architectures may include an Industry Standard Architecture (ISA) bus, Micro Channel Architecture (MCA) bus, Enhanced ISA (EISA) bus, Video Electronics Standards Association (VESA) local bus, and Peripheral Component Interconnect (PCI) bus, which can be implemented as a Mezzanine bus manufactured to the IEEE P1386.1 standard, and the like.

[0160] The processing subsystem 4104 controls the operation of the computer system 4100 and may comprise one or more processors, application specific integrated circuits (ASICs), or field programmable gate arrays (FPGAs). The processors may include be single core or multi-core processors. The processing resources of computer system 4100 can be organized into one or more processing units 4132, 4134. A processing unit may include one or more processors, including single core or multi-core processors, one or more cores from the same or different processors, a combination of cores and processors, or other combinations of cores and processors. In some examples, the processing subsystem 4104 can include one or more special purpose co-processors such as graphics processors, digital signal processors (DSPs), or the like. In some examples, some or all of the processing units of the processing subsystem 4104 can be implemented using customized circuits, such as application specific integrated circuits (ASICs), or field programmable gate arrays (FPGAs).

[0161] In some examples, the processing units in the processing subsystem 4104 can execute instructions stored in the system memory 4110 or on the computer readable storage media 4122. In various examples, the processing units can execute a variety of programs or code instructions and can maintain multiple concurrently executing programs or processes. At any given time, some or all of the program code to be executed can be resident in the system memory 4110 and/or on the computer-readable storage media 4122 including potentially on one or more storage devices. Through suitable programming, the processing subsystem 4104 can provide various functionalities described above. In instances where computer system 4100 is executing one or more virtual machines, one or more processing units may be allocated to each virtual machine

[0162] In some examples, a processing acceleration unit 4106 may be provided for performing customized processing or for off-loading some of the processing performed by the processing subsystem 4104 so as to accelerate the overall processing performed by the computer system 4100.

[0163] The I/O subsystem 4108 may include devices and mechanisms for inputting information to the computer system 4100 and/or for outputting information from or via the computer system 4100. In general, use of the term “input device” is intended to include all possible types of devices and mechanisms for inputting information to the computer system

4100. User interface input devices may include, for example, a keyboard, pointing devices such as a mouse or trackball, a touchpad or touch screen incorporated into a display, a scroll wheel, a click wheel, a dial, a button, a switch, a keypad, audio input devices with voice command recognition systems, microphones, and other types of input devices. User interface input devices
5 may also include motion sensing and/or gesture recognition devices such as the Microsoft Kinect® motion sensor that enables users to control and interact with an input device, the Microsoft Xbox® 360 game controller, devices that provide an interface for receiving input using gestures and spoken commands. User interface input devices may also include eye gesture recognition devices such as the Google Glass® blink detector that detects eye activity (e.g.,
10 “blinking” while taking pictures and/or making a menu selection) from users and transforms the eye gestures as input into an input device (e.g., Google Glass®). Additionally, user interface input devices may include voice recognition sensing devices that enable users to interact with voice recognition systems (e.g., Siri® navigator), through voice commands.

[0164] Other examples of user interface input devices include, without limitation, three
15 dimensional (3D) mice, joysticks or pointing sticks, gamepads and graphic tablets, and audio/visual devices such as speakers, digital cameras, digital camcorders, portable media players, webcams, image scanners, fingerprint scanners, barcode reader 3D scanners, 3D printers, laser rangefinders, and eye gaze tracking devices. Additionally, user interface input devices may include, for example, medical imaging input devices such as computed tomography,
20 magnetic resonance imaging, position emission tomography, medical ultrasonography devices. User interface input devices may also include, for example, audio input devices such as MIDI keyboards, digital musical instruments and the like.

[0165] In general, use of the term “output device” is intended to include all possible types of devices and mechanisms for outputting information from computer system 4100 to a user or
25 other computer. User interface output devices may include a display subsystem, indicator lights, or non-visual displays such as audio output devices, etc. The display subsystem may be a cathode ray tube (CRT), a flat-panel device, such as that using a liquid crystal display (LCD) or plasma display, a projection device, a touch screen, and the like. For example, user interface output devices may include, without limitation, a variety of display devices that visually convey

text, graphics and audio/video information such as monitors, printers, speakers, headphones, automotive navigation systems, plotters, voice output devices, and modems.

[0166] The storage subsystem 4118 provides a repository or data store for storing information that is used by the computer system 4100. The storage subsystem 4118 provides a tangible non-
5 transitory computer-readable storage medium for storing the basic programming and data constructs that provide the functionality of some examples. Software (e.g., programs, code modules, instructions) that when executed by the processing subsystem 4104 provide the functionality described above may be stored in the storage subsystem 4118. The software may be
10 executed by one or more processing units of the processing subsystem 4104. The storage subsystem 4118 may also provide a repository for storing data used in accordance with the present disclosure.

[0167] The storage subsystem 4118 may include one or more non-transitory memory devices, including volatile and non-volatile memory devices. As shown in FIG. 41, the storage subsystem 4118 includes a system memory 4110 and a computer-readable storage media 4122. The system
15 memory 4110 may include a number of memories including a volatile main random access memory (RAM) for storage of instructions and data during program execution and a non-volatile read only memory (ROM) or flash memory in which fixed instructions are stored. In some implementations, a basic input/output system (BIOS), containing the basic routines that help to transfer information between elements within the computer system 4100, such as during start-up,
20 may typically be stored in the ROM. The RAM typically contains data and/or program modules that are presently being operated and executed by the processing subsystem 4104. In some implementations, the system memory 4110 may include multiple different types of memory, such as static random access memory (SRAM) or dynamic random access memory (DRAM), and the like.

[0168] By way of example, and not limitation, as depicted in FIG. 41, the system memory 4110 may load application programs 4112 that are being executed, which may include client applications, Web browsers, mid-tier applications, relational database management systems (RDBMS), etc., program data 4111, and an operating system 4116. By way of example, the operating system 4116 may include various versions of Microsoft Windows®, Apple
25 Macintosh®, and/or Linux operating systems, a variety of commercially-available UNIX® or
30

UNIX-like operating systems (including without limitation the variety of GNU/Linux operating systems, the Google Chrome® OS, and the like) and/or mobile operating systems such as iOS, Windows® Phone, Android® OS, BlackBerry® 10 OS, and Palm® OS operating systems.

[0169] The computer-readable storage media 4122 may store programming and data constructs that provide the functionality of some examples. The computer-readable media 4122 may provide storage of computer-readable instructions, data structures, program modules, and other data for the computer system 4100. Software (programs, code modules, instructions) that when executed by the processing subsystem 4104 provides the functionality described above that may be stored in the storage subsystem 4118. By way of example, the computer-readable storage media 4122 may include non-volatile memory such as a hard disk drive, a magnetic disk drive, an optical disk drive such as a CD ROM, DVD, a Blu-Ray® disk, or other optical media. The computer-readable storage media 4122 may include, but is not limited to, Zip® drives, flash memory cards, universal serial bus (USB) flash drives, secure digital (SD) cards, DVD disks, digital video tape, and the like. Computer-readable storage media 4122 may also include, solid-state drives (SSD) based on non-volatile memory such as flash-memory based SSDs, enterprise flash drives, solid state ROM, and the like, SSDs based on volatile memory such as solid state RAM, dynamic RAM, static RAM, DRAM-based SSDs, magnetoresistive RAM (MRAM) SSDs, and hybrid SSDs that use a combination of DRAM and flash memory based SSDs. The computer-readable storage media 4122 may provide storage of computer-readable instructions, data structures, program modules, and other data for the computer system 4100.

[0170] In some examples, the storage subsystem 4118 may also include a computer-readable storage media reader 4120 that can further be connected to the computer-readable storage media 4122. The reader 4120 may receive and be configured to read data from a memory device such as a disk, a flash driver, etc.

[0171] In some examples, the computer system 4100 may support virtualization technologies, including but not limited to virtualization of processing and memory resources. For example, the computer system 4100 may provide support for executing one or more virtual machines. The computer system 4100 may execute a program such as a hypervisor that facilitates the configuring and managing of the virtual machines. Each virtual machine generally runs independently of the other virtual machines. A virtual machine may be allocated memory,

compute (e.g., processors, cores), I/O, and networking resources. Each virtual machine typically runs its own operating system, which may be the same as or different from the operating systems executed by other virtual machines executed by the computer system 4100. Accordingly, multiple operating systems may potentially be run concurrently by the computer system 4100.

5 [0172] The communications subsystem 4124 provides an interface to other computer systems and networks. The communications subsystem 4124 serves as an interface for receiving data from and transmitting data to other systems from the computer system 4100. For example, the communications subsystem 4124 may enable the computer system 4100 to establish a communication channel to one or more client computing devices via the Internet for receiving
10 and sending information from and to the client computing devices.

[0173] The communication subsystem 4124 may support both wired and/or wireless communication protocols. For example, in some examples, the communications subsystem 4124 may include radio frequency (RF) transceiver components for accessing wireless voice and/or data networks (e.g., using cellular telephone technology, advanced data network technology,
15 such as 3G, 4G or EDGE (enhanced data rates for global evolution), WiFi (IEEE 802.11 family standards, or other mobile communication technologies, or any combination thereof), global positioning system (GPS) receiver components, and/or other components. In some examples, the communications subsystem 4124 can provide wired network connectivity (e.g., Ethernet) in addition to or instead of a wireless interface.

20 [0174] The communication subsystem 4124 can receive and transmit data in various forms. For example, in some examples, the communications subsystem 4124 may receive input communication in the form of structured and/or unstructured data feeds 4126, event streams 4128, event updates 4130, and the like. For example, the communications subsystem 4124 may be configured to receive (or send) data feeds 4126 in real-time from users of social media
25 networks and/or other communication services such as Twitter® feeds, Facebook® updates, web feeds such as Rich Site Summary (RSS) feeds, and/or real-time updates from one or more third party information sources.

[0175] In some examples, the communications subsystem 4124 may be configured to receive data in the form of continuous data streams, which may include event streams 4128 of real-time
30 events and/or event updates 4130, that may be continuous or unbounded in nature with no

explicit end. Examples of applications that generate continuous data may include, for example, sensor data applications, financial tickers, network performance measuring tools (e.g. network monitoring and traffic management applications), clickstream analysis tools, automobile traffic monitoring, and the like.

5 [0176] The communications subsystem 4124 may also be configured to output the structured and/or unstructured data feeds 4126, event streams 4128, event updates 4130, and the like to one or more databases that may be in communication with one or more streaming data source computers coupled to the computer system 4100.

10 [0177] The computer system 4100 can be one of various types, including a handheld portable device (e.g., an iPhone® cellular phone, an iPad® computing tablet, a PDA), a wearable device (e.g., a Google Glass® head mounted display), a personal computer, a workstation, a mainframe, a kiosk, a server rack, or any other data processing system.

15 [0178] Due to the ever-changing nature of computers and networks, the description of the computer system 4100 depicted in FIG. 41 is intended only as a specific example. Many other configurations having more or fewer components than the system depicted in FIG. 41 are possible. Based on the disclosure and teachings provided herein, a person of ordinary skill in the art will appreciate other ways and/or methods to implement the various examples.

II. SMART CONTENT – SEARCH RESULT RANKING

20 [0179] For a content recommendation system that is configured to recommend specific content (content items) in response to input from a client or user, the process of evaluating and ranking the content items relative to each other plays an important role in the overall recommendation. For example, a user may input one or more search terms into a search engine query interface, the recommendation system may recommend one or more relevant content items (e.g., web pages,
25 documents, images, etc.) that most closely match the input search terms. In other examples, a user may author original content such as an email, a document, an on-line article, a blog post, etc., using a content authoring system, which may be configured to provide recommendations of related content items that may be relevant to the content authored by the user. For instance, as described above, a content recommendation system may recommend relevant images or links to

web pages, etc. to the user such that the author, if so desired, can incorporate one or more of the recommended content items into the content being authored by the user. Several examples of such techniques are described in the sections above. As these examples illustrate, the ranking and recommendations of specific content items may apply to many different use cases, in response to receiving manually entered user input and/or automated input from other processes. The content items to be ranked and/or recommended may correspond to images, web pages, other media files, documents, digital objects, etc., which are available to the recommendation system for searching. These content items may be stored in one or more repositories, which may be private or public repositories (e.g., the Internet), that are accessible to the recommendation system.

[0180] However, ranking of content items is a non-trivial task. For example, consider a recommendation system that uses tag matching techniques for performing the recommendations. In such a system, the content items that are available for searching by the recommendation system are tagged and the content items along with the tags may be stored in one or more repositories. The tagging may be performed by a content item tagging service/application. For a content item, the one or more tags associated with the content item are indicative of the content contained by the content item. A value (also sometimes referred to as a tag probability) may also be associated with each tag, where the value provides a measure (e.g., a probability) of the content indicated by the tag occurring in the content item. Upon receiving user input (e.g., a search term/phrase, content being authored by a user) for which recommendation of content items is to be provided, the user input may be analyzed to identify one or more tags to be associated with the user input. The recommendation system may then use a tag matching technique to identify a set of content items, from the content items available for searching, whose associated tags match the tags associated with the user input. The recommendation system may then use some ranking algorithm to rank the content items within the identified set and display the results to the user.

[0181] The effectiveness of the ranking algorithm used by the recommendation system may however be limited in certain use cases and may not produce optimal results. For example, consider a situation where there are multiple tags associated with the user input. For example, if the user types the words "coffee" and "human" into an image search engine, then tags "coffee"

and "human" may be associated with the user input. In certain embodiments, the search terms themselves may be treated as tags associated with the user input. For purposes of simplicity, assume that the collection of content items available for searching includes tagged images.

Using these two tags, the recommendation system may retrieve a set of matching content items
5 from the collection of content items (e.g., images), where a content item is deemed to be matching if at least one tag associated with the content item matches a tag associated with the user input. Consider a scenario (Case 1) where multiple matching content items are retrieved by the recommendation system and each content item has both the "coffee" and "human" tags associated with it. One possible way of ranking these retrieved content items is to (a) for each
10 matching content item, add the values associated with the tags "coffee" and "human" for that content item, and then (b) rank the content items based upon their associated added sums. This however presents a problem because, for multiple matching content items, the values associated with matching tags can add up to the same value, which is a very likely situation since most tagging services normalize probabilities on a scale of one. For example, three matching images
15 may have associated tag values as follows: Image A (("coffee", 0.5), ("human", 0.5)), Image B (("coffee", 0.2), ("human", 0.8)), and Image C (("coffee", 0.7), ("human", 0.3)). As can be seen, the sum of the values of the matching tags for each of these matching images is one ("1"), and thus there is no way, using conventional summing techniques, to rank one of these images over the other. Accordingly, simply ranking the images based upon the added sum of their matching
20 tag values cannot be used for ranking the images.

[0182] Extending the above example, it is also possible that the set of matching images includes images that matched only one tag (e.g., only matched "coffee," or only matched "human") and have the same associated tag value. This again presents a problem for ranking the images. For example, consider , three matching images may have associated tag values as
25 follows: Image A (("coffee", 0.5)), and Image B (("human", 0.5)). There is again, using conventional techniques, no way to rank one of these images over the other.

[0183] The situation is further aggravated when there are more than two tags associated with the user input. For example, if the user types the words "coffee," "human," and "café" into an image search engine, then there are three search tags "coffee", "human", and "café" associated
30 with the user input. The recommendation system may retrieve a set of matching content items

from the collection of content items (e.g., images), where a content item is deemed to be matching if at least one tag associated with the content item matches a tag associated with the user input. The number of matching tags for a matching content item may vary from only one matching tag to multiple matching tags (maximum of three matchings tags for the "coffee," "human," and "café" example). Again in this scenario, for multiple matching content items, the values associated with matching tags can add up to the same value. For example, three matching images may have associated tag values as follows: Image A (("human", 0.8), ("coffee", 0.2)), Image B (("human", 0.2), ("coffee", 0.2), ("café", 0.6)), and Image C (("human", 0.5), ("café", 0.5)). As can be seen, the sum of the values of the matching tags for each of these matching images is one ("1"), and thus there is no way, using conventional summing techniques, to rank one of these images over the other.

[0184] Accordingly, in many cases, simple tag matching techniques might not return optimal content recommendations for a variety of reasons. For example, certain images (or other content items) in the repository may be tagged with only one or two content tags, while other images/content items may be tagged with large numbers of tags, including potentially dozens or even hundreds of tags for a single content item. In such cases, conventional tag matching techniques might over recommend a heavily tagged content item (e.g., because it more frequently includes at least one tag matching the input terms) and/or might under recommend such an item (e.g., because even when matching one or more content tags the majority of its tags will still fail to match the input terms). Similarly, the input content provided by the user or client system might include only a few input terms (e.g., expressly input search terms or topics extracted from a larger input text), or might include a relatively large number of input terms, depending on the input data received. In such cases, conventional tag matching techniques might either fail to identify certain relevant content items in the repository (e.g., because too few input terms match the tags of related content items), or might mistakenly recommend less relevant content items (e.g., because the less relevant content items include one or more matching tags).

[0185] In certain embodiments, improved techniques for evaluating, ranking, and recommending tagged content items are described herein. In some embodiments, a content recommendation system may receive input content from a client device, such as a search query, newly authored text input, or the like. One or more tags may be included in or associated with

the input content received from the client device, and/or may be determined and extracted from the input content based on preprocessing and analyses techniques performed on the input content. Additionally, the content recommendation system may access a content repository storing a plurality of tagged content items, such as images, media content files, links to web pages and/or other documents. In some cases, the content repository may store the data identifying tagged content items, and for each of the tagged content item, may further store associated tag information for each item, where the tag information for a content item includes information identifying one or more tags associated with the content item and a tag value for each associated tag.

10 **[0186]** In response to receiving the input data for which a recommendation is to be made, the content recommendation system may retrieve, from the collection of searchable content items, a set of matching tagged content items from the content repository, where a content item is deemed to be a matching content item if at least one content tag associated with the content item matches a tag associated with the input content. Then, for each of the matching tagged content items
15 retrieved from the content repository, the content recommendation system may calculate two scores: (1) a first score (also referred to as tag count score) based on the number of tags associated with the content item that match tags associated with the input content, and (2) a second score (also referred to as tag value based score or TVBS) based on the tag value for each of the matching tags for the content item. The content recommendation system then calculates a
20 final ranking score for each of the matching content items, based on the first score and the second score for that matching content item. The final ranking scores computed for the set of matching content items are then used to generate a ranked list of the matching content items. This ranked list is then used for identifying a recommended subset of the matching content items to be output to the user or client system.

25 **[0187]** Referring now to FIG. 42, a block diagram is shown illustrating a computing environment 4200 with a content recommendation system 4220 implemented to evaluate and rank content items from a content repository 4230, in response to input content received from user or client systems 4210 according to certain embodiments. Various components and subsystems within the content recommendation system 4220 are also shown in this example,
30 including a graphical user interface (GUI) 4215 through which client systems 4210 may interact

with the content recommendation system 4220, to provide input content and to receive data identifying a subset of recommended content items. In certain embodiments, the GUI 4215 may be a GUI of a separate client application 4215 (e.g., a web browser application) used by a user to author content. In this embodiment, the content recommendation system 4220 may receive the user provided or authored content from the client application. The content may be received by the content recommendation system 4220 using an application programming interface (API) that enables the client application and the content recommendation system 4220 to interact and exchange information with each other.

[0188] The embodiment depicted in FIG. 42 is merely an example and is not intended to unduly limit the scope of claimed embodiments. One of ordinary skill in the art would recognize many possible variations, alternatives, and modifications. For example, in some implementations, the content recommendation system 4220 may have more or fewer systems or subsystems than those shown in FIG. 42, may combine two or more systems, or may have a different configuration or arrangement of systems. The content recommendation system 4220 may be implemented one or more computing systems including, in some embodiments, separate systems using independent computing and network infrastructures with dedicated and specialized hardware and software. Alternatively or additionally, one or more of these components and subsystems may be integrated into a single system performing separate functionalities. The various systems, subsystems and components depicted in FIG. 42 may be implemented in software (e.g., code, instructions, program) executed by one or more processing units (e.g., processors, cores) of the respective systems, hardware, or combinations thereof. The software may be stored on a non-transitory storage medium (e.g., on a memory device).

[0189] At a high level, the content recommendation system 4220 is configured to receive user input content and then make content item recommendations responsive to and based upon the user input content. The recommendations are made from a collection of content items that are available and accessible to the content recommendation system 4220 for making recommendations. The collection of content items can include images, various types of documents, media content, digital objects, and the like. Based upon tags information associated with the user input content and tags information associated with the collection of content items, the content recommendation system 4220 is configured to, using a tag matching technique,

identify a set of matching content items for the user input content. The content recommendation system 4220 is then configured to rank the content items in the set of matching content items using innovative ranking techniques described in this disclosure. Based upon the ranking, the content recommendation system 4220 is then configured to identify and recommend a subset of the matching content items to be output to the user or client system.

[0190] The content recommendation system 4220 includes a content tagger subsystem 4222 that is configured to receive or retrieve content items that are available for recommendation by the content recommendation system 4220. The content items can include, without limitation, images, web-pages, documents, media files, etc. The content items may be received or retrieved from one or more content repositories 4230. The content repositories 4230 may include various public or private content repositories such as libraries or databases, including image libraries, document stores, local or wide area networks of web-based resources (e.g., the Internet), and the like. One or more content repositories 4230 may be stored locally to the content recommendation system 4220, while other content repositories may be separate, remote from the content recommendation system 4220 and accessible to the content recommendation system 4220 via one or more computer networks.

[0191] In certain embodiments, for each content item, the content tagger subsystem 4222 is configured to retrieve and analyze the contents of the content item and identify one or more content tags (tags) to be associated with the content item. For each tag that is associated with a content item, the content tagger 4222 may also determine a tag value associated with the tag, where the value provides a measure (e.g., a probability) of the content indicated by the tag occurring in the content item. The tag value for a tag may correspond to a numerical value measure representing how applicable the particular content tag is to that content item. One or more tags and corresponding tag values may be associated with a content item. For a content item having multiple associated tags and corresponding tag values, the tag values may represent the relative prominence of the image topic or theme indicated by the tags in that image. For example, a first tag associated with a content item having a relatively high tag value may indicate that the content or feature indicated by that first tag is particularly relevant and prominent in the content item. In contrast, a second tag associated with the same content item having a lower tag value may indicate that the content or feature indicated by the second tag is not as prominent or

prevalent in the content item relative to the content indicated by the first content tag. For example, an image content item may have two associated tags and values as follows: ("human", 0.8), ("coffee", 0.2)). This indicates that image contains content related to coffee (e.g., a coffee cup) and a human (e.g., a human drinking the coffee), and the human is more prominently depicted in the image as compared to the depiction of the coffee (e.g., a large portion of the image may depict a human and the coffee cup may cover a small area of the image. The tag values may be represented using different formats. For example, in some implementations, that tag values may be expressed as floating point numbers between 0.0 and 1.0. In some implementations, the sum of all tag values for tags associated with a particular content item may sum to a fixed and uniform value (e.g., add up to 1).

[0192] In some embodiments, the content tagger 4222 may use the services of a content tagging service to perform the tagging task for the content items, including identifying one or more tags to be associated with a content item and a tag value for each tag. In certain embodiments, content tagger 4222 is implemented using one or more predictive machine-learning models that have been trained to take a content item as input and predict tags for the content item and the associated tag values. In some embodiments, the tags may be selected from a set of preconfigured tags that are used to train the models. Various machine learning techniques using pre-trained machine learning models and/or other artificial intelligence based tools, including AI-based text or image classification systems, topic or feature extractions, and/or any other combination of techniques described above, may be used for determining the tags to be associated with a content item and the corresponding tag values.

[0193] In some embodiments, the content items retrieved from content repositories 4230 might already include associated content tags and tag values. When the retrieved content items do not include the tags information, and/or when the content recommendation system 4220 is configured to determine additional tags for the content items, then the content tagger 4222 may be used to update or generate new tags for the retrieved content items. The content tagger 4222 may use various different techniques to generate the tag information (e.g., one or more tags and associated tag values) for a content item. For example, the content tagger 4222 may use any or all of the previously described techniques for analyzing the retrieved content items and determining content tags, such as parsing, processing, feature extraction, and/or other analysis techniques. The type of parsing, processing, feature extraction, and/or analysis may depend on

the type of the content item. For instance, for text-based content items, such as blog posts, letters, emails, articles, documents, etc., the analysis may include keyword extraction and processing tools (e.g., stemming, synonym retrieval, etc.), topic analysis tools, etc. For content items that are images, artificial intelligence-based image classification tools may be used to
5 identify particular image features and/or generate image tags. For instance, analysis of an image may identify a plurality of image features, and the image may be tagged with each of these identified features. One or both types of analyses (i.e., tag extraction from images, and keyword/topic extraction from text content), may be performed via REST-based services or other web services, using analytics, machine-learning algorithms, and/or artificial intelligence (AI)
10 based techniques, such as an AI-based cognitive image analysis service, or a similar AI/REST cognitive text service to be used for text content. Similar techniques may be used for other types of content items, such as video files, audio files, graphics, or social media posts, where a specialized web-service may be used to extract and analyze specific features (e.g., words, objects within images/video, facial expressions, etc.) depending on the media type of the content item.

15 **[0194]** In some embodiments, the content tagger 4222 may use one or more machine-learning and/or artificial intelligence-based pre-trained models trained with training data to identify and extract content features to be used for determining the tags and the tag values for the content items. For example, a model training system may generate one or more models, which may be trained in advance using machine-learning algorithms based on training data sets including a
20 training data set of previous input data (e.g., text input, images, etc.), and the corresponding tags for the previous input data. In various embodiments, one or more different types of trained models may be used, including classification systems that execute supervised or semi-supervised learning techniques, such as a Naïve Bayes model, a Decision Tree model, a Logistic Regression model, or a Deep Learning Model, or any other machine learning or artificial intelligence based
25 prediction system that may execute supervised or unsupervised learning techniques. For each machine learning model or model type, the trained models may be executed by one or more computing systems during which a content item is provided as input to the one or more models and the output from the models may identify the one or more tags to be associated with the content item, or the output of the models may be used to identify the one or more tags to be
30 associated with the content item. Accordingly, content tagger 4222 may use various different tools or techniques such as, without limitation, keyword extraction and processing (e.g.,

stemming, synonym retrieval, etc.), topic analysis, feature extraction from images, machine learning and AI-based modeling tools and text or image classification systems, and/or any other combination of techniques described above to determine or generate the tag information (e.g., one or more tags and associated tag values) for each content item that is available for
5 recommendation.

[0195] In certain embodiments, the content items available for recommendations and their associated tag information (e.g., for each content item, one or more tags associated with the content item and the corresponding tag values) may be stored in a data store 4223. In some
10 embodiments, the content/tag information data store 4223 may store data identifying the content items retrieved from content repositories 4230, which may include the items themselves (e.g., images, web-pages, documents, media files, etc.), or additionally/alternatively may include references to the items (e.g., item identifiers, network addresses from which the content items can be retrieved, descriptions of the items, thumbnails of the items, etc.). An example illustrating the type of data that may be stored in the content/tag information data store 4223 is
15 shown in FIG. 45 and discussed in more detail below.

[0196] The content recommendation system 4220 includes a tag identifier subsystem 4221 that is configured to receive the user input content from devices 4210 and determine one or more tags to be associated with the user content. In some embodiments, the user content received from devices 4210 may include the associated tags. In some other embodiments, the tag identifier
20 4221 may be configured to process the input data to determine one or more tags to be associated with the input data. As one example, the tag identifier 4221 may use a data tagging service to identify a set of one or more tags to be associated with the input data. The tag identifier 4221 may provide then provide the tags associated with the input data (and also the user content in some implementations) to a recommended content item identifier and ranker subsystem 4224
25 (which may be referred to for brevity as the content item ranker 4224) for further processing.

[0197] In some embodiments, tag identifier 4221 may use the various techniques used by the content tagger 4222, and described above, to determine the one or more tags to be associated with the received user content. In some embodiments, tag identifier 4221 and content tagger 4222 may both use the same superset of tags from which tags to be associated with the user input
30 and the content items are determined. In certain embodiments, tag identifier 4221 and content

tagger 4222 may use the same data tagging service for identifying tags to be associated with the user content, and the content items, respectively. In yet other embodiments, the tag identifier 4221 and content tagger 4222 subsystems may be implemented as a single subsystem configured to perform similar (or even identical) processes on the content items received from repositories 4230 and the input content received from client systems 4210.

[0198] As described, when the content items are tagged by content tagger 4222, for each content item, one or more tags to be associated with the content item are identified along with a tag value for each tag. With respect to tagging for user contents, in some embodiments, tag identifier 4221 is configured to only determine the tags to be associated with the user input without any associated tag values. In such embodiments, each tag is given equal weight with respect to the ranking performed by the content item ranker 4224 based upon the tags associated with the user input. In some other embodiments, both tags and associated tag values may be determined for the user content and used by the content item ranker 4224 for ranking the content item recommendations.

[0199] As described above, the user content that is received and processed by tag identifier 4221 can come in different forms. For example, the user content can include contents of a document (e.g., emails, articles, blog posts, documents, social media posts, images, etc.) being authored by a user, content created or selected by a user (e.g., multi-media files), and the like. As another example, the user input can be a document (e.g., a webpage) accessed by a user. As yet another example, the user content could be search terms input by a user (e.g., a browser-based search engine) for performing a search. In certain embodiments, for example, for search terms, the terms themselves may be used as tags.

[0200] As depicted in FIG. 42 and described above, the content item ranker 4224 receives as input, from the tag identifier 4221, information identifying a set of one or more tags associated with the user content. Based upon this tag information for the user content and based upon the content items available for recommendations, the content item ranker 4224 is configured use tag matching techniques to identify one or more content items that are most related to and/or relevant to the input content. In cases where multiple content items are identified as being related or relevant to the user input, the content item ranker 4224 is further configured to use innovative ranking techniques described herein to rank the content items. More details related to

various techniques used by the content item ranker 4224 for scoring and ranking the content items are described in more detail below. The content item ranker 4224 is configured to generate a ranked list of content items to be recommended to the user in response to the user input received for the user. The ranked list of the content items is then provided to a recommendation selector subsystem 4225 for further processing.

[0201] Using the ranked list of content items received from the content item ranker 4224, the recommendation selector 4225 is configured to select one or more particular content items to be recommended to user responsive to the input content received from the client system 4210. In certain scenarios, all the content items in the ranked list may be selected for recommendation. In some other scenarios, a subset of the ranked content items may be selected for recommendation, where the subset includes less than all the content items in the ranked list and the one or more content items included in the subset are selected based upon the rankings of the content items in the ranked list. For example, the recommendation selector 4225 may select the top "X" ranked (e.g., top 5, top 10, etc.) content items from the ranked list for the recommendation, where X is some integer that is less than or equal to the number of ranked items. In certain embodiments, the recommendation selector 4225 may select the content items to be included in the subset to be recommended to the user based upon the scores associated with the content items in the ranked list. For example, only those content items having associated scores above a user-configurable threshold score may be selected to be recommended to the user.

[0202] Information identifying the content items selected for recommendation by the recommendation selector 4225 may then be communicated from the content recommendation system 4220 to a user client device 4210 of the user. The information regarding the recommended content items may then be output to the user via the user client device. For example, information regarding the recommendations may be output via a GUI 4215 displayed on the user client device or via an application 4215 executed by the user client device. For example, if the user input corresponded to a search query input by a user via a webpage displayed by a browser executed by the user device, the information regarding the recommendations may also be output to the user via that webpage or additional webpages displayed by the browser. In certain embodiments, for each recommended content item, the information output to the user may include information identifying the content item (e.g., text information, a thumbnail of an image,

etc.) and information for accessing the content item. For example, the information accessing the content item may be in the form of a link (e.g., an URL), which when selected by the user (e.g., by a mouse clicking action) causes the corresponding content item to be accessed and displayed to the user via the user client device. In some embodiments, the information identifying the content item and information for accessing the content item may be combined (e.g., a thumbnail representation of a recommended image that both identifies the image content item and can also be selected by the user to access the image itself.)

[0203] In various embodiments, the content recommendation system 4220, including its associated hardware/software components 4221-4225 and services may be implemented as back-end services remote from the front-end client devices 4210. Interactions between client devices 4210 and the content recommendation system 4220 may be Internet-based web browsing sessions, or client-server application sessions, during which users access may input user content (e.g., search terms, original authored content, etc.) via client devices 4210, and may receive content item recommendations from the content recommendation system 4220. Additionally or alternatively, the content recommendation system 4220 and/or content repositories 4230 and related services may be implemented as specialized software components executing directly on a client device 4210.

[0204] In some embodiments, system 4200 shown in FIG. 42 may be implemented as a cloud-based multi-tiered system, in which upper-tier user devices 4210 may request and receive access to the network-based resources and services via the content recommendation system 4220 residing on back-end application servers deployed and executed upon an underlying set of resources (e.g., cloud-based, SaaS, IaaS, PaaS, etc.). Some or all of the functionality described herein for the content recommendation system 4220 may be performed by or accessed using Representational State Transfer (REST) services and/or web services including Simple Object Access protocol (SOAP) web services or APIs, and/or web content exposed via Hypertext Transfer Protocol (HTTP) or HTTP Secure protocols. Thus, although not shown in FIG. 42 so as not to obscure the components shown with additional details, the computing environment 4200 may include additional client devices, one or more computer network(s), one or more firewalls, proxy servers, routers, gateways, load balancers, and/or other intermediary network devices,

facilitating the interactions between the client devices 4210, content recommendation system 4220, and content repositories 4230.

[0205] In various implementations, the systems depicted in computing environment 4200 may be implemented using one or more computing systems and/or networks, including specialized
5 server computers (such as desktop servers, UNIX servers, mid-range servers, mainframe computers, rack-mounted servers, etc.), server farms, server clusters, distributed servers, or any other appropriate arrangement and/or combination of computing hardware. For instance, content recommendation system 4220 may run an operating system and/or a variety of additional server
10 servers, File Transport Service (FTP) servers, Common Gateway Interface (CGI) servers, Java servers, database servers, and other computing systems. Any or all of the components or subsystems within the content recommendation system 4220 may include at least one memory, one or more processing units (e.g., processor(s)), and/or storage. The subsystems and/or
15 modules in the content recommendation system 4220 may be implemented in hardware, software (e.g., program code or instructions executable by a processor) executing on hardware, or combinations thereof. In some examples, the software may be stored in a memory (e.g., a non-transitory computer-readable medium), on a memory device, or some other physical memory, and may be executed by one or more processing units (e.g., one or more processors, one or more
20 processor cores, one or more Graphics Process Units (GPUs), etc.). Computer-executable instructions or firmware implementations of the processing unit(s) can include computer-executable or machine-executable instructions written in any suitable programming language, which can perform the various operations, functions, methods, and/or processes described herein. The memory may store program instructions that are loadable and executable on the processing unit(s), as well as data generated during the execution of these programs. The memory may be
25 volatile (such as random access memory (RAM)) and/or non-volatile (such as read-only memory (ROM), flash memory, etc.). The memory may be implemented using any type of persistent storage device, such as computer readable storage media. In some examples, computer-readable storage media may be configured to protect a computer from an electronic communication containing malicious code. .

[0206] FIG. 43 depicts a simplified flowchart 4300 depicting processing performed by a content recommendation system for identifying and ranking content items that are relevant to user content according to certain embodiments. The processing depicted in FIG. 43 may be implemented in software (e.g., code, instructions, program) executed by one or more processing units (e.g., processors, cores) of the respective systems, hardware, or combinations thereof. The software may be stored on a non-transitory storage medium (e.g., on a memory device). The method presented in FIG. 43 and described below is intended to be illustrative and non-limiting. Although FIG. 43 depicts the various processing steps occurring in a particular sequence or order, this is not intended to be limiting. In certain alternative embodiments, the processing may be performed in some different order or some steps may also be performed in parallel. The processing depicted in FIG. 43 may be performed by one or more systems depicted in FIG. 42, such as by the content recommendation system 4220. As an example, for the embodiment depicted in FIG. 42, the processing in 4302 and 4304 may be performed by tag identifier 4221, the processing in 4306 to 4316 may be performed by the content item ranker 4224, and the processing in 4318 and 4320 may be performed by recommendation selector 4225. However, it should be understood that the techniques and functionalities described in connection with FIG. 43 need not be limited only to implementations within the specific computing infrastructure shown in FIG. 42, but may be implemented using the other compatible computing infrastructures described herein.

[0207] In 4302, the content recommendation system 4220 may receive input content from one or more user or client systems 4210. As discussed above in reference in to FIG. 42, the input content may be received from a client device 4210, through a graphical user interface 4215 (e.g., a web-based GUI) provided by the content recommendation system 4220. In other examples, the input content may be received by a web server or backend service executing within the content recommendation system 4220, based on data transmitted by a front-end application (e.g., mobile application) installed on the client device 4210.

[0208] In some embodiments, the input content received in step 4302 may correspond to a set of search terms or phrases, input by a user into a search engine user interface. In other embodiments, the input content may correspond to original content authored by a user and input into a specialized user interface. For example, new original content may include on-line articles,

newsletters, emails, blog entries, and the like, and such content may be input by the user via software-based word processor tools, email client applications, web-development tools, etc. In still other examples, the input content received in step 4302 may be an image, graphic, audio input, or any other text and/or multimedia content which is generated or selected by the user via a client device 4310.

[0209] Referring briefly to FIG. 44, an example user interface 4400 is shown, including a user interface screen 4410 that allows users to input original authored content. In this example, the user interface screen 4410 is generically labeled “Content Authoring User Interface,” however, in various embodiments the user interface 4410 may correspond to an interface for a word processor, an article designer or blog entry creator, an email client application, or the like. In this example, the user interface 4410 includes a first textbox 4411 where the user may input a title or subject for the authored content, and a second textbox 4412 where the user may input the full text (e.g., article, email body, document, etc.) for the input content. Additionally, user interface 4410 includes a selectable button 4413 that allow the user to initiate a search for relevant content items (e.g., images, related articles, etc.) that may be incorporated into the newly authored content. In some cases, the selection of button 4413, or a similar user interface component, may initiate the process shown in FIG. 43, by initially analyzing and transmitting the user content received via the user interface 4410 (e.g., user content entered in 4411 and/or 4412) to the content recommendation system 4220. In other embodiments, a background process may run continuously (or periodically) within the front-end user interface, to continuously (or periodically) analyze the new text input received from the user (e.g., content input by the user in 4411 and/or 4412) and to re-initiate the process of FIG. 43 in response to text updates, so that the content item recommendations may be continuously or periodically updated in real time.

[0210] Referring back for FIG. 43, in 4304, the content recommendation system 4220 determines one or more tags for the input content received in step 4302. In some embodiments, the input content received in 4302 may already have tags associated with it or embedded in the input content, and the tag identifier 4221 may identify and extract a set of predetermined tags associated with the input content. If the input content received in step 4302 corresponds to search terms, then the tag identifier 4221 may simply use the search terms inputted by the user as the tags (e.g., excluding certain words such as articles, connective words, prepositions,

quantifiers, etc.). In the case of originally authored text content, or any other non-tagged input content received by the content recommendation system 4220, the tag identifier 4221 may be configured to analyze the various features of the input content received in 4302, and determine one or more tags to be associated with the received input content based on the analysis. As
5 previously described, the tag identifier 4221 may use various different techniques to determine one or more tags to be associated with the input content received in 4302.

[0211] Returning to the example user interface 4400 shown in FIG. 44, for this example, the input content received in 4302 may correspond to text input by the user in the subject/topic box 4411, namely, “Is Coffee Healthier For You Than Tea?” and text entered in box 4412. Based on
10 an analysis of the subject content in 4411 and the body text of the authored article in 4411, and any additional input content that may be provided, the content recommendation system 4220 may determine in 4304 that tags “coffee,” “tea,” and “human” are to be associated with the user content.

[0212] In 4306, based upon the tags determined for the input content in 4302, the content
15 recommendation system 4220 uses a tag matching technique to identify a set of matching content items from the collection of content items available for recommendation, where a content item is considered to be a match and identified in 4306 if at least one tag associated with the content item matches a tag determined for the input content in 4304. The set of content items identified in 4306 may also be referred to the matching set of content items and includes content items that
20 are candidates for being recommend to the user responsive to the input content received in 4302. In various embodiments, data processing techniques such as stemming and synonym retrieval/comparison may be used as part of the matching process in 4306 to identify matching tags.

[0213] For example, in the embodiment depicted in FIG. 42, the collection of content items
25 available for recommendations along with their associated tag information (e.g., tags associated with the content items and associated tag values) may be stored in the content/tag information data store 4223. As part of the processing in 4306, the content recommendation system 4220 may compare the one or more tags determined in 4304 to the tags associated with content items available for recommendation and identify those content items from the collection having at least
30 one associated tag that matches a tag identified in 4304.

[0214] Continuing with the example of FIG. 44 and assuming that tags "coffee," "human," and "tea" have been determined in 4304 for the user input content received in 4302, example table 4500 in FIG. 45 shows a matching set of content items that have been identified by the content recommendation system 4220 (e.g., by the content item ranker 4224) from the collection of content items available for making recommendations. As can be seen from table 4500, eight different image content items have been identified as having at least one associated tag that matches at least one of tags "coffee," "human," or "tea." As shown in this example, each matched image content item is identified using an image identifier 4501. The image description 4502 provided in FIG. 45 is a description of the contents of the matched images and has been provided in FIG. 45 so that that actual images do not have to be shown. Each matched image has one or more associated tags 4503 and a tag value 4504 is associated with each tag. In the example in table 4500, the tag values are floating point values with a predetermined range (e.g., between 0.0 and 1.0), and where the sum of the tag values for each content item sum to the same total (e.g., 1.0). Such embodiments provide additional technical advantages, for example, the uniformity in the sums of the tag count score and the TVBS may assure that content items having more associated tags are not ranked higher artificially or over-recommended based on the larger number of associated tags. Further, as discussed below, by having tag values ranging between 0.0 and 1.0, when they are multiplied (for example, in a use case described below), the resultant values allow content items to be ranked within a particular group or bucket, while assuring that the highest ranked content items within one group do not outrank the lowest ranked content items within the next higher group.

[0215] It can be seen from FIG. 45 that a content item (in this example, an image) is deemed to be a match if at least one tag associated with the content item matches a tag associated with the input content. A matched content item may have associated tags that match one or more of the tags associated with the input content. A matched content item may also have other tags associated with it that are different from the tags for the input content (e.g., Image_1, Image_3, etc.)

[0216] In 4308, a tag count score (first score) is computed for each matching content item identified in 4306 based upon the number of tags associated with the content item that match tags determined for the input content. In some scenarios, each matching tag of the content item is

given a value of one, and so the tag count score computed in 4308 for a content item equals the number of tags of the content item that match the tags associated with the input content. For the example matching images identified in FIG. 45, table 4600 in FIG. 46 identifies that tag count scores 4602 for each of the matching images identified in FIG. 45 (and also in FIG. 46). For
5 example, the tag count score for Image_1 is "1" (one) since one tag ("human") associated with the image matched tags associated with the input content. As another example, the tag count score for Image_2 is "2" (two) since two tags ("coffee" and "human") associated with the image matched tags associated with the input content. As yet another example, the tag count score for Image_8 is "1" (one) since one tag ("coffee") associated with the image matched tags associated
10 with the input content.

[0217] In 4310, the matching content items are grouped (bucketed) into groups or buckets based upon the tag counts scores computed for the content items in 4308. In certain embodiments, a group (or bucket) contains all content items having the same tag count score. . In scenarios where each matching tag of the content item is given a value of one, each group or
15 bucket includes content items having the same number of matched tags. Processing in 4310 may be optional and may not be performed in certain embodiments.

[0218] Continuing with the example shown in FIG. 46, the eight matching image may be grouped into two groups or buckets: a first group or bucket including content items having a tag count score of 1 (one) and a second group of bucket having a tag count score of 2 (two). The
20 first group will include images {Image_1, Image_3, Image_6, and Image_8} and the second group will include images {Image_2, Image_4, Image_5, and Image_7}. Notably, in this example, none of the matched images matched all three tags ("coffee," "human," "tea") associated with the input content.

[0219] In 4312, for each group identified in 4310, a tag value based score (second score) is
25 calculated for each of the candidate content items in that group. In some embodiments, the tag value based score (TVBS) for a particular content item is based upon and calculated using the tag values associated with the matching tags for the content item.

[0220] In certain embodiments, the TVBS for a content item is calculated by multiplying the tag values associated with the tags of the content item that matched the tags associated with the
30 input content. For example:

TVBS for Image_1 in FIG. 46 = 0.93

TVBS for Image_2 in FIG. 46 = 0.5 * 0.5 = 0.25

TVBS for Image_3 in FIG. 46 = 0.65

TVBS for Image_4 in FIG. 46 = 0.35 * 0.60 = 0.21, and so on.

5 Table 4600 shows TVBSs 4603 calculated for the various matching images using the above-described technique.

[0221] In certain embodiments, a Naïve Bayes approach is used for computing the tag value based score for a content item. For example, assuming that two tags tag_1 and tag_2 are determined for the input content, then the tag value based score (TVBS) for an image content
 10 item can be expressed as:

TVBS for $Image_i = P(Image_i | tag_1, tag_2) =$ probability of $Image_i$ given tags tag_1 and tag_2

Extending this for "n" tags:

$P(Image_i | tag_1, tag_2 \dots, tag_n) =$ probability of $Image_i$ given (tags tag_1 and tag_2 and ... tag_n)

Assuming the tags $tag_1, tag_2 \dots, tag_n$ to be mutually independent:

15 $TVBS \text{ for } Image_i = TVBS_i = P(Image_i | tag_1, tag_2 \dots, tag_n)$
 $= P(Image_i | tag_1) * P(Image_i | tag_2) * \dots * P(Image_i | tag_n)$

Equation 1

Now by Simple Naïve Bayes: (Equation 2)

$$P(Image_i | tag_t) = \frac{P(tag_t | Image_i) * P(Image_i)}{P(tag_t)}$$

20 Where:

$P(Image_i | tag_t) =$ probability of tag_t for the $Image_i$

$P(Image_i) =$ consider every image as unique and this term can be discarded or ignored

$P(tag_t) =$ frequency of the tag in the collection of content items (i.e., the number of content items (e.g., images) in the collection of content items available for recommendation that are tagged

with tag_i). As this term is in the denominator, the less frequent the tag is present in the collection of content items (i.e., the lesser number of content items having this associated tag), the higher will be the TVBS score for an image having the tag.

Upon expansion, the above formula becomes: Equation 3

$$5 \quad P(\text{Image}_i | tag_1, tag_2 \dots, tag_n) = \frac{P(tag_1 | \text{Image}_i)}{P(tag_1)} * \frac{P(tag_2 | \text{Image}_i)}{P(tag_2)} * \dots * \frac{P(tag_n | \text{Image}_i)}{P(tag_n)}$$

The numerator in Equation 3 is the multiplication of the probabilities that will result in higher score when the probabilities are equally likely (given the same tags matched for multiple images, the denominator will remain the same).

[0222] The following example illustrates the application of Equation 3 for calculating TVBSs for content items. Assume that the tags determined for input content are "human" and "coffee." Further assume that a collection of content items (images) available for recommendations contains three images with the following tags and tag values:

Image A: ("human", 0.5), ("coffee", 0.5)

Image B: ("human", 0.1), ("coffee", 0.9)

15 Image C: ("human", 0.8), ("coffee", 0.2)

For Image A:

$$P(\text{human} | \text{image}) = 0.5 \text{ and } P(\text{coffee} | \text{image}) = 0.5$$

Applying Equation 3,

$$\begin{aligned} \text{TVBS for Image A} &= P(\text{Image A} | \text{human, coffee}) \\ &= P(\text{Image A} | \text{human}) * P(\text{Image A} | \text{coffee}) \\ &= \frac{P(\text{human} | \text{Image A})}{P(\text{human})} * \frac{P(\text{coffee} | \text{Image A})}{P(\text{coffee})} \\ &= \frac{0.5}{\text{Frequency}(\text{human})} * \frac{0.5}{\text{Frequency}(\text{coffee})} \\ &= \frac{0.5}{\text{Frequency}(\text{human})} * \frac{0.5}{\text{Frequency}(\text{coffee})} \end{aligned}$$

20

Where "Frequency(human)" is the number of content items in the collection of content items available for recommendation that have a "human" tag associated with them, and

25

"Frequency(coffee)" is the number of content items in the collection of content items available for recommendation that have a "coffee" tag associated with them

$$= \frac{0.5}{3} * \frac{0.5}{3}$$

$$\text{TVBS for Image A} = 0.028$$

5 Using a similar technique,

$$\text{TVBS for Image B} = (0.1/3) * (0.9/3) = 0.01$$

$$\text{TVBS for Image C} = (0.8/3) * (0.2/3) = 0.018$$

As can be seen from this example, where the frequencies are the same, the TVBS is higher when the probabilities are equally likely (given the same tags matched for multiple images, the denominator will remain the same).

10 **[0223]** As per the expansion of Equation 3, the frequencies of content items in the collection of content items having a particular associated tag are taken into consideration for calculating TVBSs (and thus also impact the overall ranking score as described below). The fewer number of content items having a particular associated tag (i.e., lower frequency), the higher the TVBS score will be for an image having the particular tag. In some embodiments this is desirable to increase the likelihood of such content items with "rare" or less frequent tags to be ranked higher than content items with more frequent tags ranked higher and thus increase the likelihood of them being included in the list of content items that is recommended to a user. The value of the TVBS for a content item is thus inversely proportional to the frequency of occurrences of a particular tag appearing in the collection of content items (i.e., the number of content items in the collections having the associated particular tag).

15 **[0224]** In 4314, an overall ranking score is calculated for each matching content item identified in 4306 based upon the tag count score calculated for the content item in 4308 and the TVBS calculated for the content item in 4312. In some embodiments, the overall ranking score for a candidate content item may be calculated as the sum of the tag count score (calculated in 4308) and the TVBS (calculated in 4312) calculated for the content item. That is, for the image content item *Image_j*:

$$\text{Ranking Score (Image)} = \text{TagsCountScore}_i + \text{TVBS}_i \text{ ---Equation 4}$$

[0225] For the example depicted in FIG. 46, column 4604 indicates the overall ranking score computed for each matching image by, for each matching image, adding the tag count score for that image (indicated in column 4602) and the TVBS for the image (indicated in column 4603). For example, the calculated overall ranking scores are as follows:

5 Image_1: $1 + 0.93 = 1.93$

Image_2: $2 + 0.25 = 2.25$

Image_3: $1 + 0.65 = 1.65$, and so on.

[0226] In step 4316, the content recommendation system 4220 (e.g., the content items ranker 4224 in the content recommendation system 4220) generates a ranked list of the matching
10 content items based upon the overall ranking scored computed for the content items in 4314. Thus, continuing with the example of FIGS. 44-46, based upon the overall ranking scores (in column 4604) computed for the matching images, the Images may be ranked, from highest to lowest, as follows: (1) Image_2, (2) Image_5, (3) Image_4, (4) Image_7, (5) Image_6, (6) Image_1, (7) Image_8, (8) Image_3.

15 [0227] In certain embodiments wherein the tag values for the tags are in the range between 0.0 and 1.0, the computation of the overall ranking score using the (TagCountScore + TVBS) approach ensures that a content item having a higher number of associated tags that matched tags associated with the input content is ranked higher than a content item with a lower number of tag matches. For example, in the example in FIG. 46, an image with a tag count score of 2
20 (corresponding to two tags associated with the image matched tags associated with the input content) will always have an overall ranking score and thus be ranked higher than an image with a tag count score of 1 (corresponding to two tags associated with the image matched tags associated with the input content). This is because, given that the tag values range between zero and one, the TVBS for the image, which is computed by multiplying the tag values associated
25 with the matching tags, cannot exceed one. This also implies that, for a first group or bucket of content items corresponding to a first tag count score and a second group or bucket of content items corresponding to a second tag count score, if the first tag count score is higher than the second tag count score, then each content item in the first group will be ranked higher (sue to higher overall ranking scores) than the content items in the second group. Thus, in the example

where three tags ("coffee," "human," and "tea") were determined for the input content, content items having three content tag matches to the tags for the input content will always outrank the content items having two matching content tags, each of which will always outrank the content items having one matching content tag, and so on. Within each group or bucket, the content
5 items may be ranked based on their TVBSs, which favor both higher and more equal parameters for the matching content tags. However, it should be understood that in other embodiments, different equations or logic may be used for calculating tag count scores, TVBSs, and overall ranking scores, in order to implement different content item ranking priorities and policies.

[0228] In 4318, the content recommendation system 4220 (e.g., the recommendation selector
10 4225) may use the ranked list generated in 4316 to select one or more content items to be recommended to the user. In certain scenarios, all the content items in the ranked list may be selected for recommendation. In some other scenarios, a subset of the ranked content items may be selected for recommendation, where the subset includes less than all the content items in the ranked list and the one or more content items included in the subset are selected based upon the
15 rankings of the content items in the ranked list. For example, the recommendation selector 4225 may select the top "X" ranked (e.g., top 5, top 10, etc.) content items from the ranked list for the recommendation, where X is some integer that is less than or equal to the number of ranked items in the list. In certain embodiments, the recommendation selector 4225 may select the content items to be included in the subset to be recommended to the user based upon the overall
20 ranking scores associated with the content items in the ranked list. For example, only those content items having associated scores above a user-configurable threshold score may be selected to be recommended to the user.

[0229] In 4320, the content recommendation system 4220 may communicate information regarding the content items selected in 4318 to the user device. This information may be
25 referred to as recommendation information since it includes information regarding the content items to be recommended to the user. The recommendation information communicated in 4320 may also include ranking information (e.g., overall ranking scores associated with the selected content items). This information may be used on the user device to determine how (e.g., the order) information regarding the recommended content items is to be displayed to the user via
30 the user device. In some embodiments, as part of the recommendation information, the

recommendation selector 4225 may transmit either the content items themselves, or certain information identifying the content items (e.g., content item identifier and description, thumbnail image, network path or link for downloading, etc.) to the client device 4310 from which the input content was received in step 4302.

5 [0230] Information regarding the selected recommendations may then be output to the user via the user device. For example, information regarding the recommendations may be output via a GUI 4215 displayed on the user client device or via an application 4215 executed by the user client device. For example, if the user input corresponded to a search query input by a user via a webpage displayed by a browser executed by the user device, the information regarding the
10 recommendations may also be output to the user via a webpage showing results of the search or additional webpages displayed by the browser. In certain embodiments, for each recommended content item, the information output to the user may include information identifying the content item (e.g., text information, a thumbnail of an image, etc.) and information for accessing the content item. For example, the information accessing the content item may be in the form of a
15 link (e.g., an URL), which when selected by the user (e.g., by a mouse clicking action) causes the corresponding content item to be accessed and displayed to the user via the user client device. In some embodiments, the information identifying the content item and information for accessing the content item may be combined (e.g., a thumbnail representation of a recommended image that both identifies the image content item and can also be selected by the user to access the
20 image itself.)

[0231] For example, referring to FIG. 47, an example user interface 4700 is shown, corresponding to an update of the user interface screen 4400 from FIG. 44 that displays information related to recommended images. In this example, based on the title/subject 4711, the body text 4712, and/or any other input content, the content recommendation system 4220 has
25 selected, from the ranked list of images, the four highest-ranked content item images for recommendation to the user. Information related to these top four ranked images is displayed in order of their rank within a dedicated portion 4714 of the user interface 4700 for showing content item recommendations. In certain embodiments, thumbnail representations of the recommended images may be displayed in 4714. The user interface 4700 may support drag-and-drop

functionality or other techniques to allow the user to incorporate one or more of the suggested images displayed in 4714 into the body 4712 of the authored content.

[0232] The processing depicted in FIG. 43 and described above is not intended to be limiting. Various variations may be provided in different embodiments. For example, for the embodiment depicted in FIG. 43 and described above, the processing in 4312, 4314, and 4316 is performed for all the matching content items identified in 4306. In certain variations, the tag count scores calculated in 4308 may be used to filter out certain content items from further processing. For example, where the content items have different tag count scores, the content items with the lowest tag count score (or some other threshold) may be filtered out from further processing in the flowchart. In some other embodiments, only those content items having the highest tag count scores may be selected for further processing, filtering out the other content items from further processing. For example, the content item ranker 4224 might only calculate TVBSs for the highest tag score group (as determined in step 4304), or may calculate TVBSs in order from the highest to the lowest tag count score groups, during which the calculation process may stop upon reaching a threshold number of candidate content items or a threshold tag count score. Such filtering reduces the number of content items to be processed and may make the overall recommendation operation be performed faster and more efficiently using lesser processing resources (e.g., processor, memory, networking resources).

[0233] In the method depicted in FIG. 43 and described above, each tag associated with or determined for the input content in 4304 is given equal weight with respect to the ranking performed by the content recommendation system 4220. Based upon this assumption, the tag score for each matching content item was determined as the number of its content tags matching the tags of the input content. Accordingly, each of the matching content tags associated with a matching content item were valued/weight equally with respect to the determination of the tag count score. However, in other embodiments, different weights may be given to the tags associated with the input content. For example, for two tags determined for the input content, one tag may be indicated as "more important" for the input content by giving it a higher weight than the other tag. For example, for the example depicted in FIGS. 44-47 and described above, in which tags "human," "coffee," and "tea") were determined for the input content, instead of giving equal importance to the three tags, the tags be instead be weighed as follows: human = 1,

coffee = 2, and tea = 4. This weighing may indicate a relative importance of the tags to the input contents, for example, "tea" is weighed heavier than "coffee," which is weighed heavier than "human." In certain embodiments, the logic used for computing tag count scores for each content item may be modified to take into consideration the different weights assigned to the tags
 5 for the input content. According to one such modified logic, the contribution of each matching tag of the content item is multiplied by the weight associated with that same tag for the input content. For example, using the (human = 1, coffee = 2, and tea = 4) weighing for the input content, the tag count score for the matching images in FIG. 45 would be as follows:

Image_1:

- 10 - TagCountScore(without weighing) = "human" tag match = 1
 - TagCountScore(with weighing) = "human" tag match = $1 * 1 = 1$

Image_2:

- TagCountScore(without weighing) = "coffee" and "human" tags match = $1 + 1 = 2$
 15 - TagCountScore(with weighing) = "coffee" and "human" tags match = $2(1) + 1(1) = 3$

Image_3:

- TagCountScore(without weighing) = "tea" tag match = 1
 - TagCountScore(with weighing) = "tea" tag match = $4(1) = 4$

20

Image_4:

- TagCountScore(without weighing) = "coffee" and "human" tags match = $1 + 1 = 2$
 - TagCountScore(with weighing) = "coffee" and "human" tags match = $2(1) + 1(1) = 3$

25 Image_5:

- TagCountScore(without weighing) = "tea" and "human" tags match = $1 + 1 = 2$
 - TagCountScore(with weighing) = "tea" and "human" tags match = $4(1) + 1(1) = 5$

Image_6:

- TagCountScore(without weighing) = "coffee" tag match = 1
- TagCountScore(with weighing) = "coffee" tag match = $2(1) = 2$

5 Image_7:

- TagCountScore(without weighing) = "tea" and "human" tags match = $1 + 1 = 2$
- TagCountScore(with weighing) = "tea" and "human" tags match = $4(1) + 1(1) = 5$

Image_8:

- 10
- TagCountScore(without weighing) = "coffee" tag match = 1
 - TagCountScore(with weighing) = "coffee" tag match = $2(1) = 2$

[0234] As a result of the different tag count scores, the grouping or bucketing of the content items performed in 4310 would be different. Thus, content items having content tags matching both the “Tea” and “Human” tags would be grouped together and assigned a tag score of 5, while content items having only one matching content tag for “Tea” would be grouped together and assigned a tag score of 4, and content items having content tags matching both the “Coffee” and “Human” tags would be grouped together and assigned a tag score of 3, and so on. Thus, in such embodiments, the overall ranking of the candidate content items is impacted by not just how many tags of the content item matched the tags of the input content but also by which specific tags of the content items match the input content tags and the weights of relative importance given to those tags. In this example, images Image_5 and Image_7 would be the highest ranking overall content items, based on their highest tag count score of 5 (Tea = 4 + Human = 1).

[0235] Although specific implementations have been described, various modifications, alterations, alternative constructions, and equivalents are possible. Implementations described in the present disclosure are not restricted to operation within certain specific data processing environments, but are free to operate within a plurality of data processing environments. Additionally, although implementations have been described using a particular series of transactions and steps, it should be apparent to those skilled in the art that this is not intended to

be limiting. Although some flowcharts describe operations as a sequential process, many of the operations can be performed in parallel or concurrently. In addition, the order of the operations may be rearranged. A process may have additional steps not included in the figure. Various features and aspects of the above-described implementations may be used individually or jointly.

5 [0236] Further, while implementations described in the present disclosure have been described using a particular combination of hardware and software, it should be recognized that other combinations of hardware and software are also possible. Some implementations described herein may be implemented only in hardware, or only in software, or using combinations thereof. The various processes described herein can be implemented on the same processor or different
10 processors in any combination.

[0237] Where devices, systems, components, or modules are described as being configured to perform certain operations or functions, such configuration can be accomplished, for example, by designing electronic circuits to perform the operation, by programming programmable electronic circuits (such as microprocessors) to perform the operation, such as by executing
15 computer instructions or code, or processors or cores programmed to execute code or instructions stored on a non-transitory memory medium, or any combination thereof. Processes can communicate using a variety of techniques including, but not limited to, conventional techniques for inter-process communication, and different pairs of processes may use different techniques, or the same pair of processes may use different techniques at different times.

20 [0238] Specific details are given in this disclosure to provide a thorough understanding of the embodiments. However, embodiments may be practiced without these specific details. For example, well-known circuits, processes, algorithms, structures, and techniques have been shown without unnecessary detail in order to avoid obscuring the embodiments. This description provides example embodiments only, and is not intended to limit the scope, applicability, or
25 configuration of other embodiments. Rather, the preceding description of the embodiments will provide those skilled in the art with an enabling description for implementing various embodiments. Various changes may be made in the function and arrangement of elements.

[0239] The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense. It will, however, be evident that additions, subtractions, deletions, and
30 other modifications and changes may be made thereunto without departing from the broader

spirit and scope of the disclosure. Thus, although specific implementations have been described, these are not intended to be limiting; various modifications and equivalents are within the scope of the disclosure.

WHAT IS CLAIMED IS:

1. A method of selecting content based on vector comparisons within a vector space, comprising:
- receiving, by a computing device, and via a user interface, text input data;
- 5 determining, by the computing device, at least one of a keyword, a topic, or a feature of the text input data, based on an analysis of the text input data;
- executing, by the computing device, a transformation algorithm, wherein the at least one determined keyword, topic, or feature are provided as input to the transformation algorithm, and wherein the transformation algorithm outputs a feature vector corresponding to
- 10 the text input data;
- comparing, by the computing device, the feature vector corresponding to the text input data, to each of a plurality of additional feature vectors stored within a vector space;
- selecting, by the computing device, one or more of the additional feature vectors, based on the comparisons of the feature vector to the plurality of additional feature vectors;
- 15 retrieving, by the computing device, and from a content repository, one or more content files corresponding to the one or more selected additional feature vectors; and
- rendering, by the computing device, selectable representations of the one or more content files via the user interface.
2. The method of claim 1, further comprising:
- 20 receiving, via the user interface, a selection of a first user input component corresponding to a first content file;
- retrieving the first content file from the content repository; and
- embedding a representation of first content file within a user interface region including the text input data.
- 25 3. The method of claim 1, wherein comparing the feature vector corresponding to the text input data to the plurality of additional feature vectors stored within the vector space comprises:
- retrieving one or more tags associated with the feature vector;

determining a subset of the plurality of additional feature vectors having one or more tags matching the one or more tags associated with the feature vector; and
comparing the feature vector corresponding to the text input data, to the subset of the plurality of additional feature vectors stored within the vector space.

5 4. The method of claim 1, wherein comparing the feature vector corresponding to the text input data to the plurality of additional feature vectors stored within the vector space comprises:

 for each particular feature vector of the plurality of additional feature vectors stored within the vector space, calculating a Euclidean distance between the particular feature
10 vector and the feature vector corresponding to the text input data.

 5. The method of claim 1, wherein the retrieved content files comprise a plurality of image files, and wherein the method further comprises, prior to receiving the text input data via the user interface:

 receiving and storing each of the plurality of image files in the content repository;
15 using an image classification software tool to identify one or more image features within each of the plurality of image files;

 generating one or more image tags for each of the plurality of image files, based on the image features identified within the image files; and

 generating the plurality of additional feature vectors stored within the vector
20 space, corresponding to the plurality of image files in the content repository, based on the based on the image features identified within the image files.

 6. The method of claim 1, wherein the vector space comprises a plurality of different vector spaces, each different vector space storing vectors corresponding to a different type of media content, and wherein the method further comprises:

25 receiving a selection via the user interface identifying a type of media content to be embedded with the text input data; and

 accessing a particular vector space from the plurality of different vector spaces, corresponding to the identified type of media content, wherein the one or more additional feature vectors are selected from the particular vector space.

7. The method of claim 6, wherein the plurality of different vector spaces includes at least:

a first vector space storing a plurality of feature vectors corresponding to image files; and

5 a second vector space storing a plurality of feature vectors corresponding to web pages.

8. The method of claim 1, wherein the analysis of the text input data comprises (1) a keyword extraction process, (2) a stemming process, and (3) a synonym retrieval process.

10 9. A method comprising:

receiving, by a content recommendation system, input content from a client computing device;

determining, by the content recommendation system, a set of one or more tags associated for the input content;

15 based upon the set of tags for the input content, identifying by the content recommendation system, a matching set of content items from a plurality of content items, wherein a content item from the plurality of content items is included in the matching set of content items if at least one tag associated with the content item matches a tag in the one or more tags for the input content;

20 for each content item in the matching set of content items:

(a) calculating, by the content recommendation system, a first score for the content item, wherein the first score is based on a number of content tags associated with the content item matching the one or more tags for the input content;

25 (b) calculating, by the content recommendation system, a second score for the content item, wherein the second score is based upon one or more tag values corresponding to one or more tags associated with the content item that match the one or more tags for the input content; and

(c) calculating, by the content recommendation system, an overall item ranking score for the content item based upon the first score and the second score calculated for the content item;

5 generating, by the content recommendation system, a ranked list of content items from the matching set of content items based upon the overall item ranking score calculated for each of the set of matching content items;

selecting, by the content recommendation system, one or more content items from the matching set of content items, based on the ranked list of content items; and

10 communicating, by the content recommendation system and to the client computing device, recommendation information regarding the one or more content items selected from the matching content items, wherein the recommendation information enables information regarding the selected one or more content items to be output using the client computing device.

10. The method of claim 9, wherein:

15 a weight is associated with at least one tag in the one or more tags for the input content; and

calculating the first score for the content item comprises:

determining that a tag associated with the content item matches the at least one tag in the one or more tags for the input content; and

20 calculating the first score for the content item by using the weight associated with the at least one tag.

11. The method of claim 9, wherein calculating the second score for the content item comprises:

25 identifying the one or more content tags associated with the content item that match the one or more tags for the input content;

determining, for each tag in the one or more content tags associated with the content item, a tag value for the tag; and

computing the second score for the content item by multiplying the one or more tag values determined for the one or more tags associated with the content item.

12. The method of claim 9, wherein calculating the second score for the content item comprises using a naïve Bayes technique based upon the one or more tag values corresponding to one or more tags associated with the content item that match the one or more tags for the input content.

5 13. The method of claim 9, wherein calculating the second score for the content item comprises:

for at least a first tag in the one or more tags associated with the content item that match the one or more tags for the input content, determining a frequency of content items in the plurality of content items that are associated with the first tag;

10 calculating the second score using the frequency wherein the second score is inversely proportional to the frequency.

14. The method of claim 9, wherein calculating the overall item ranking score for the content item comprises adding the first score calculated for the content item and the second score calculated for the content item.

15 15. The method of claim 9, further comprising:

for each content item in the plurality of content items, determining a set of one or more tags to be associated with the content item and, for each tag, a tag value for the tag.

20 16. The method of claim 9, wherein the input content received from the client computing device comprises at least one of: content authored by a user of the client computing device, or one or more search terms provided as input to a search engine.

17. The method of claim 9, wherein determining the one or more tags associated with the input content comprises:

performing a keyword extraction process on the input content;

performing a stemming process on the output of the keyword extraction process;

25 and

performing a synonym retrieval process on the output of the stemming process.

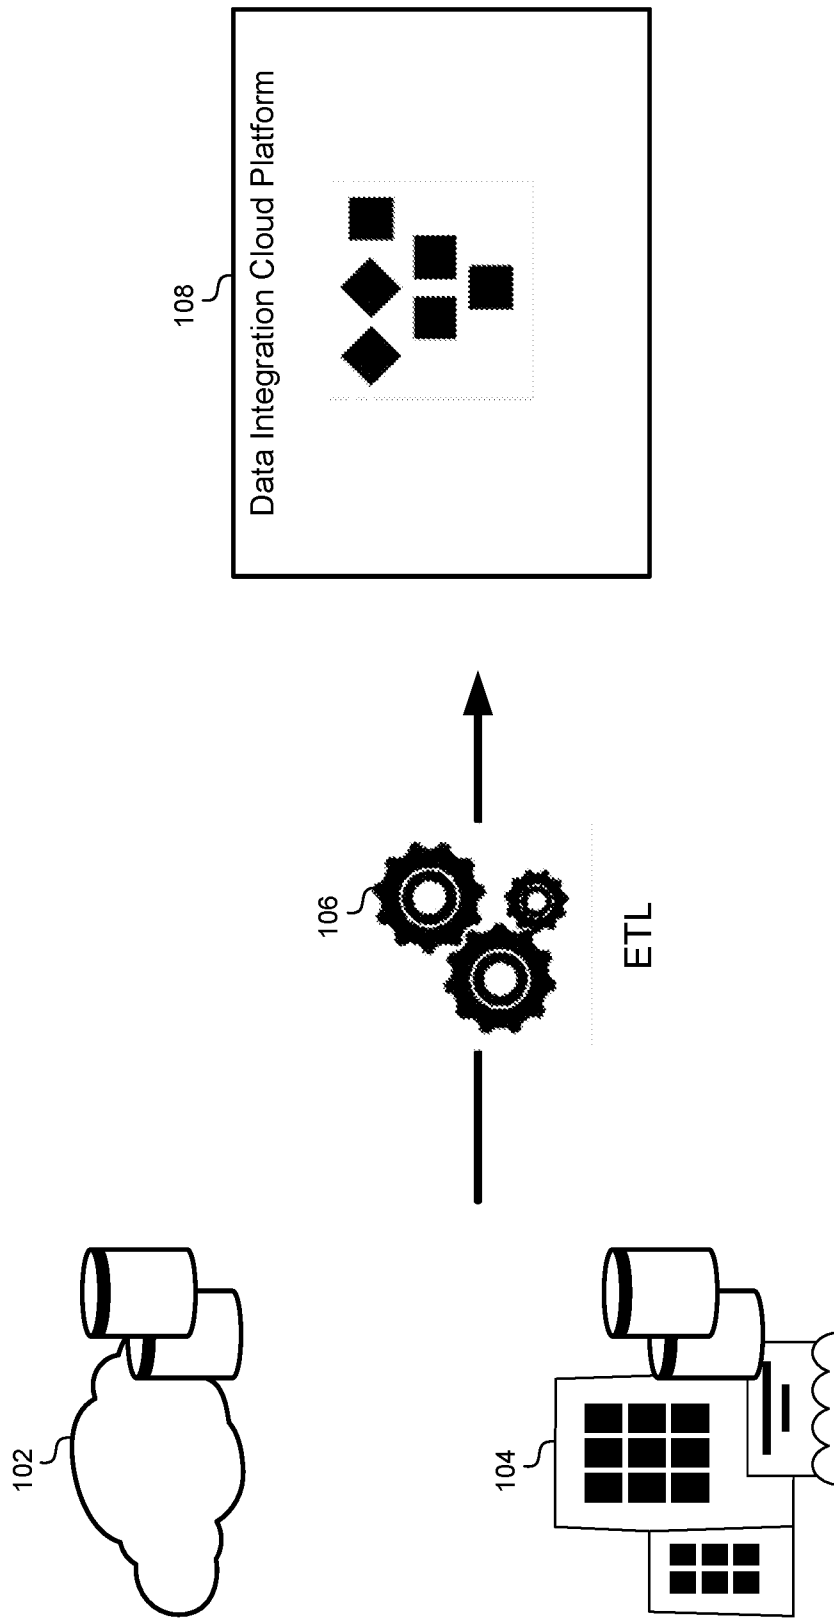


FIG. 1

202

Data Integration Platform Cloud Service

Welcome!

Instances
Activity
SSH Access

As of Dec 8, 2017 5:04 PM UTC

Summary

2

Instances

2

OCPUs

30

GB
Memory

874

GB
Storage

2

Public IPs

Instances

204

[Create instance](#)

DIPCDEF

Version: 18.1.3

Created On: Dec 8, 2017 12:39:24 PM UTC

OCPUs: 1

Memory: 15 GB

Storage: 437 GB

DIPCGHI

Version: 17.4.5

Created On: Dec 8, 2017 9:36:22 AM UTC

OCPUs: 1

Memory: 15 GB

Storage: 437 GB

One or more patch(es) available

206

FIG. 2

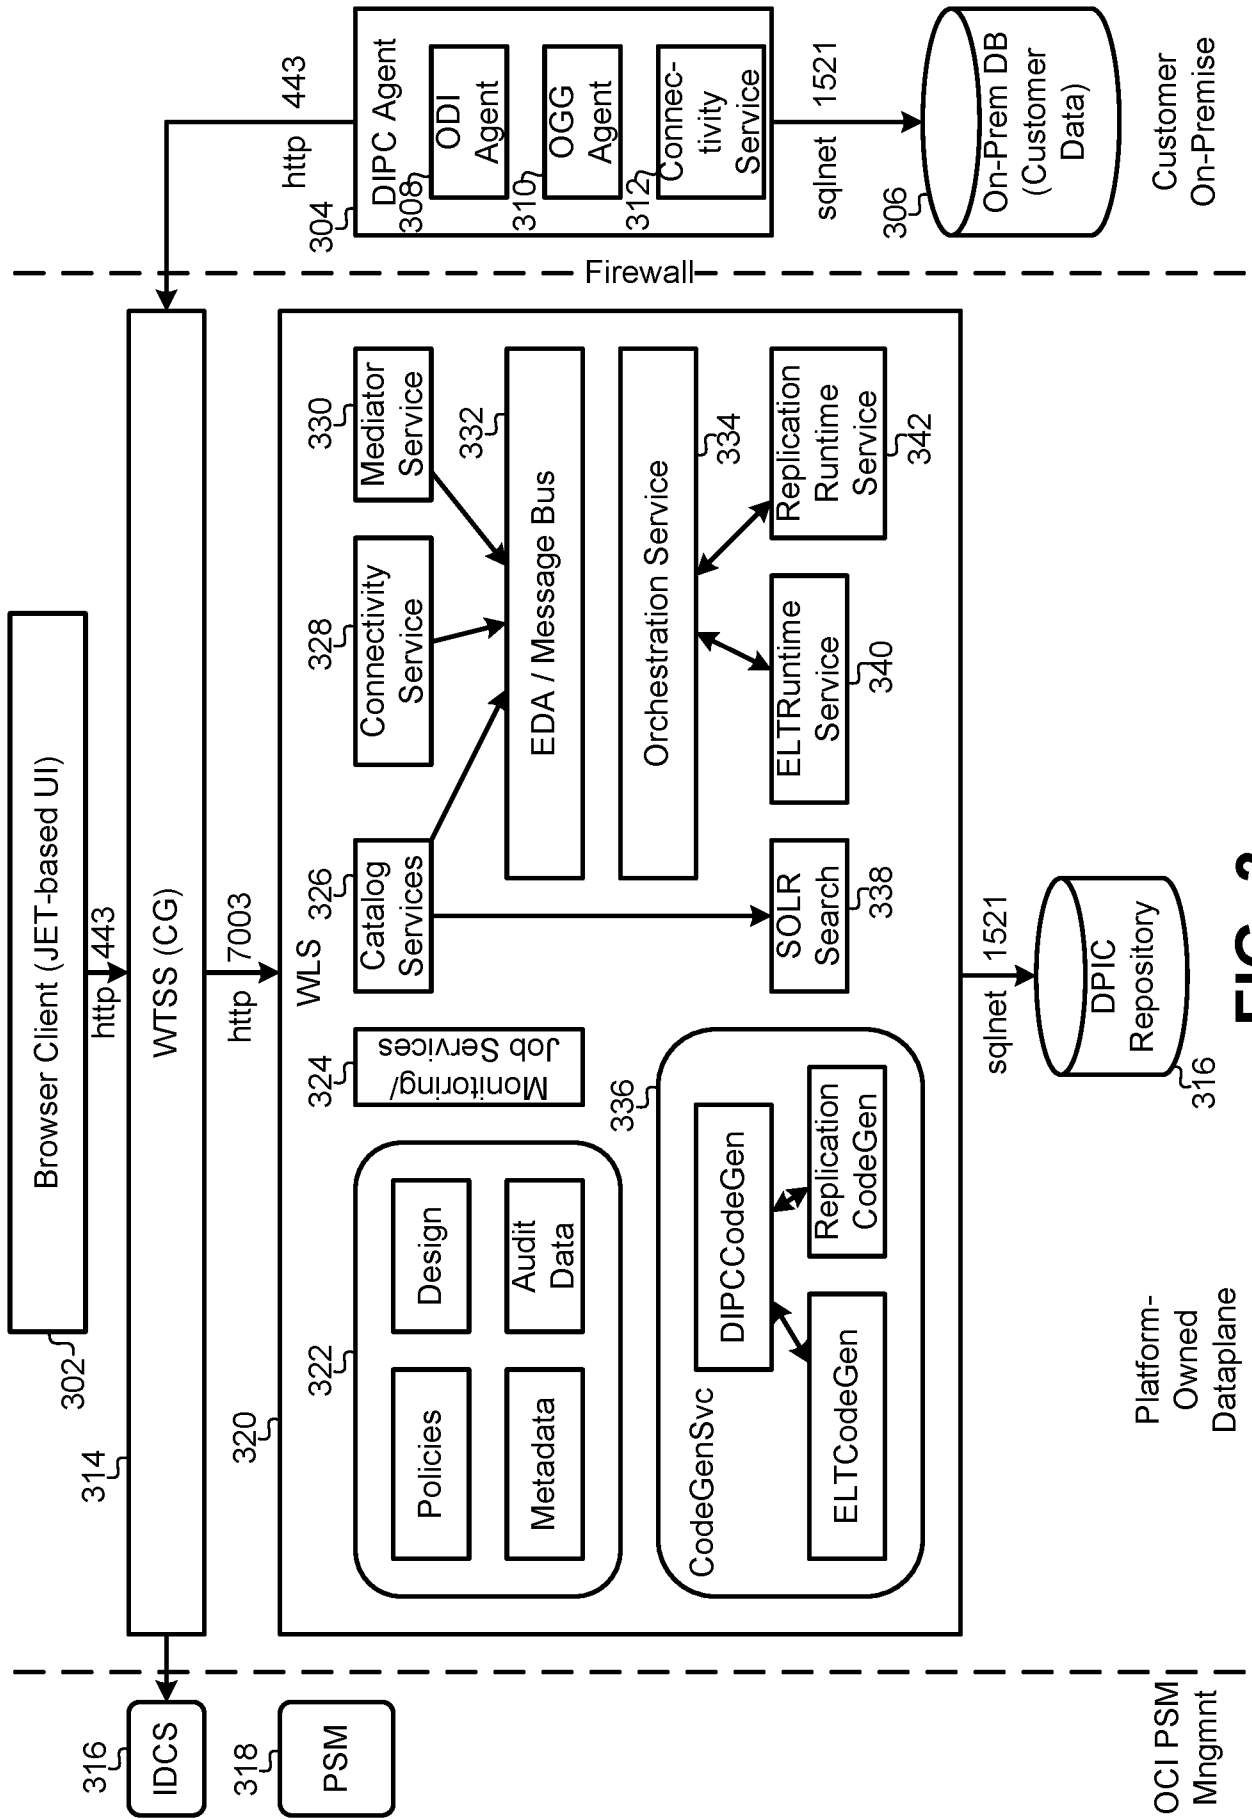


FIG. 3

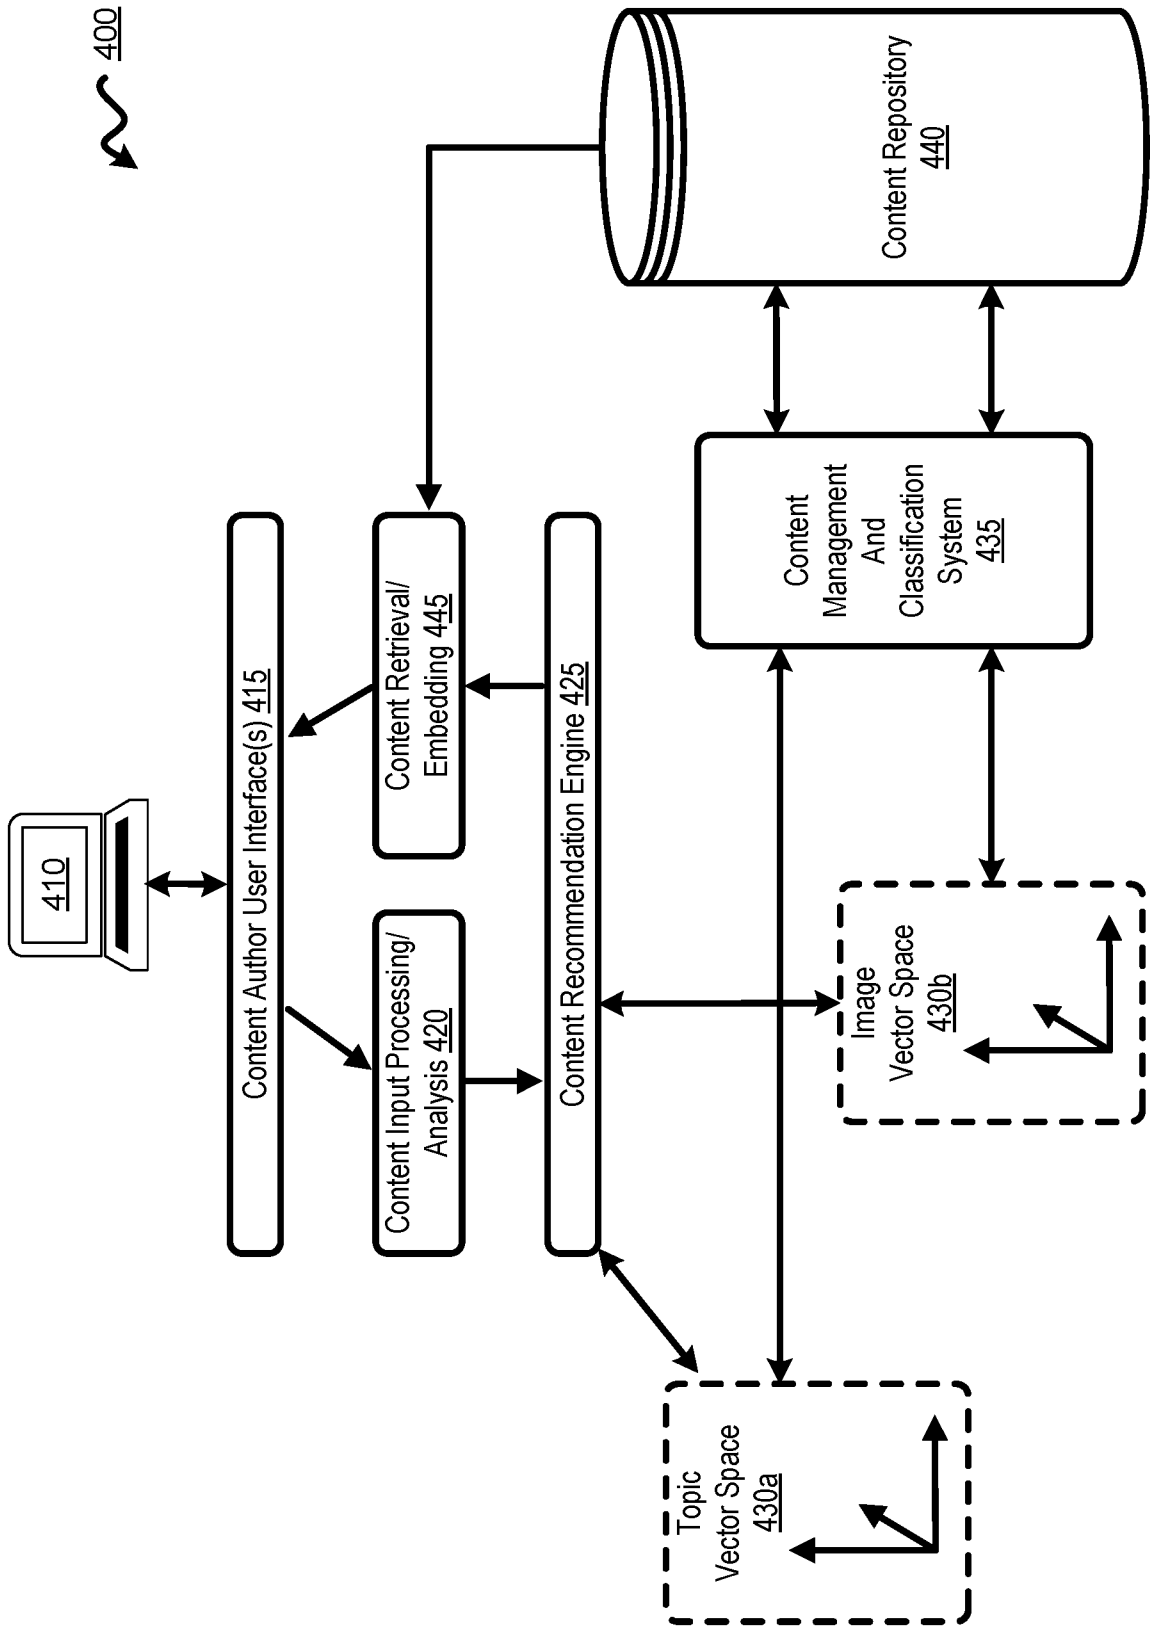


FIG. 4

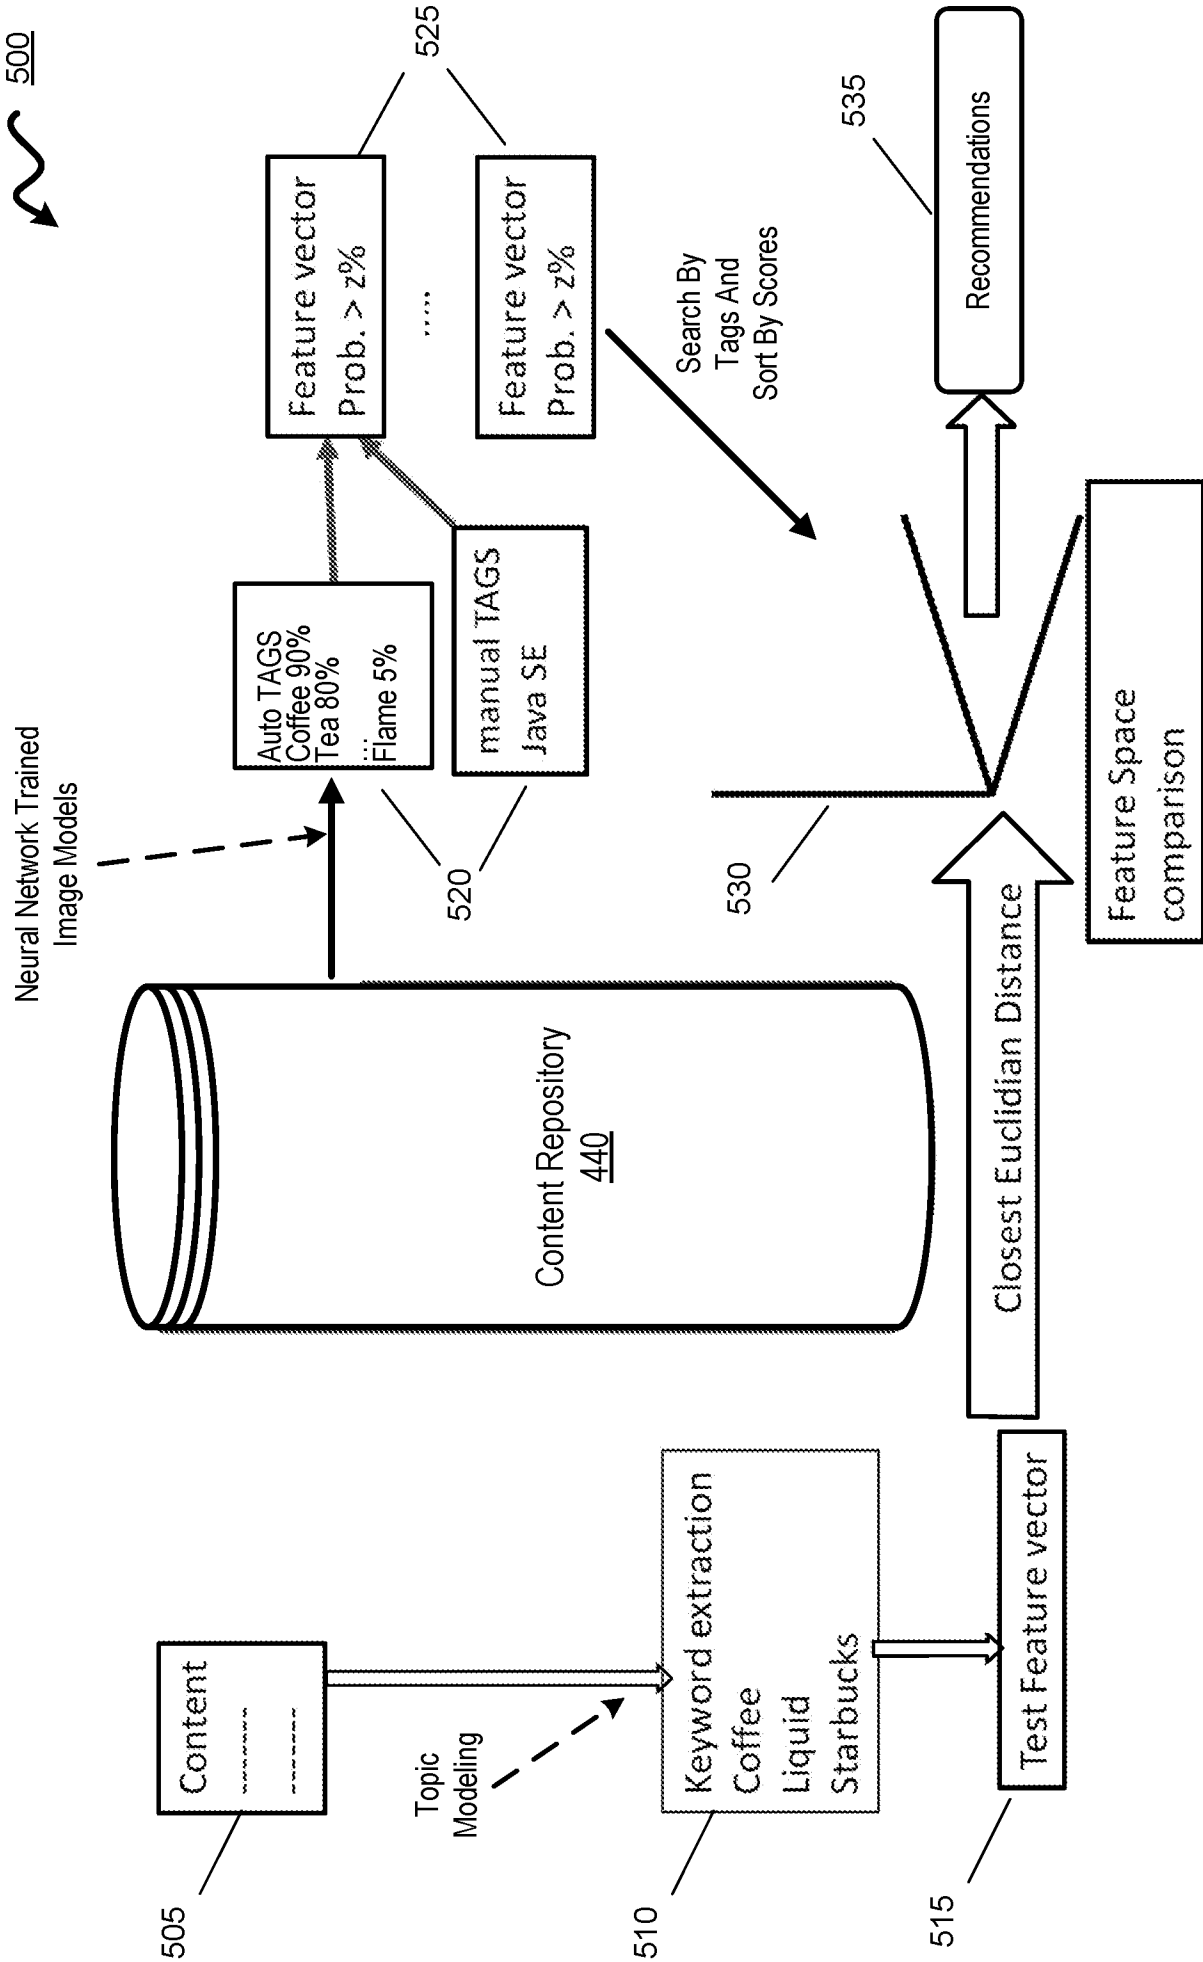


FIG. 5

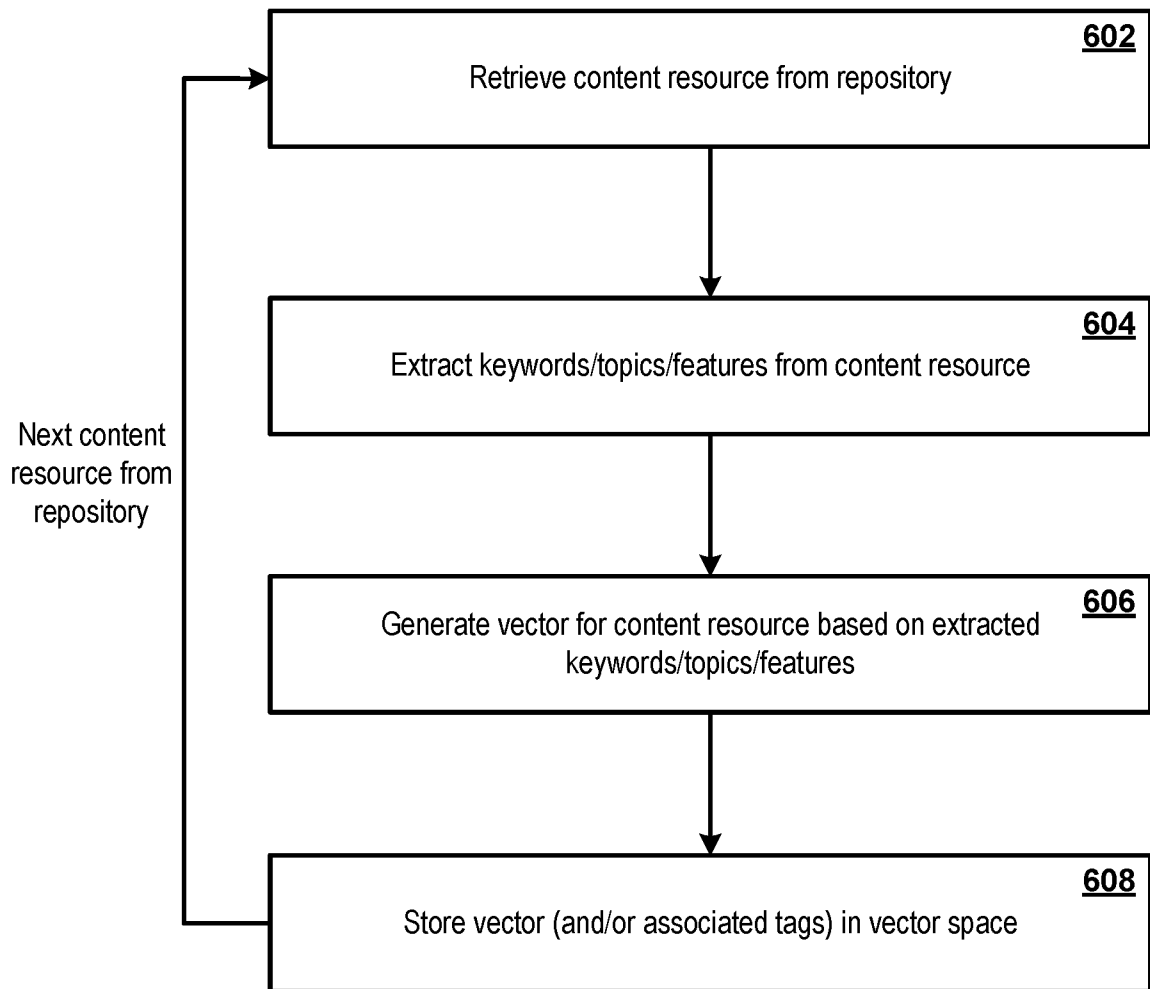


FIG. 6

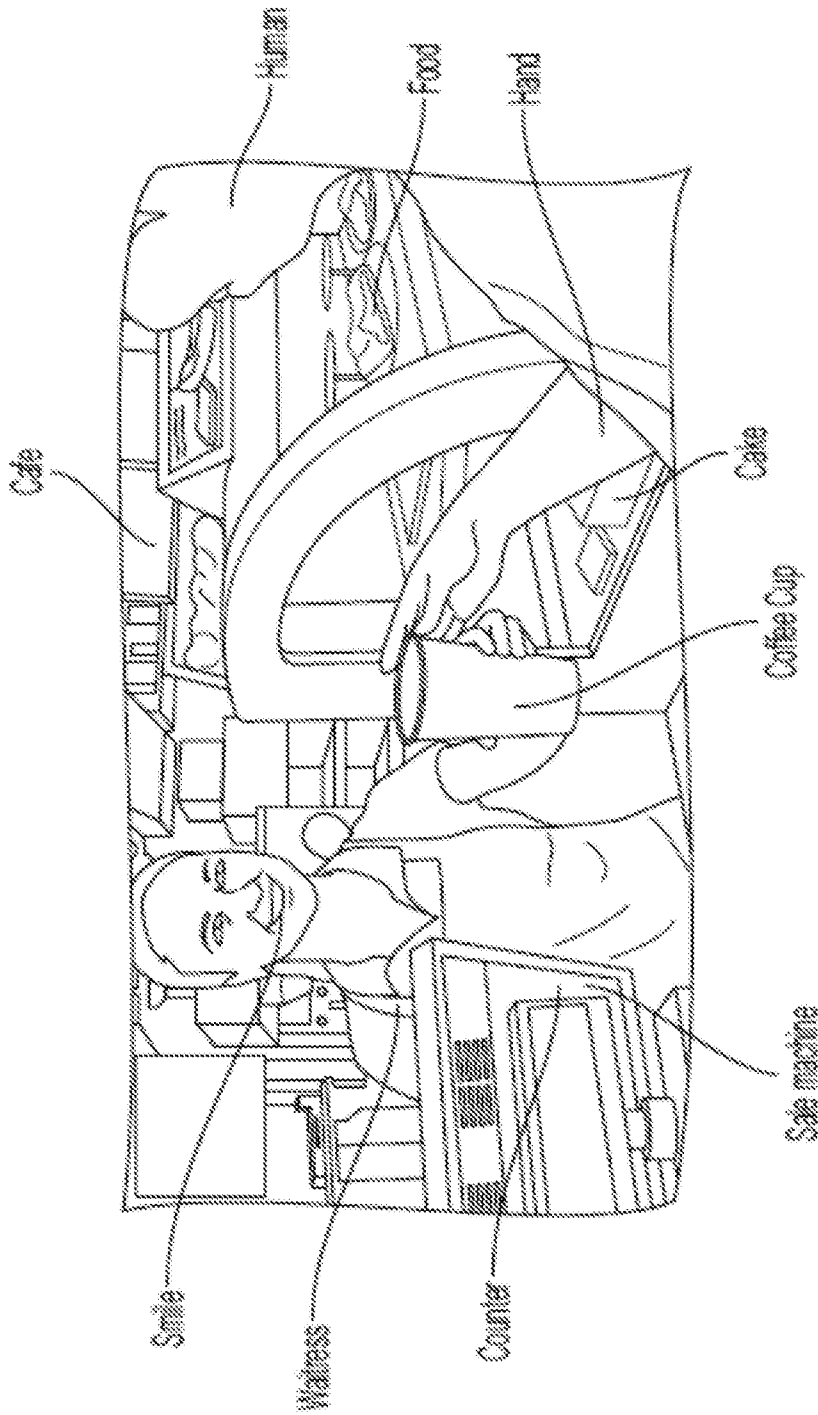


FIG. 7

Keywords Extraction

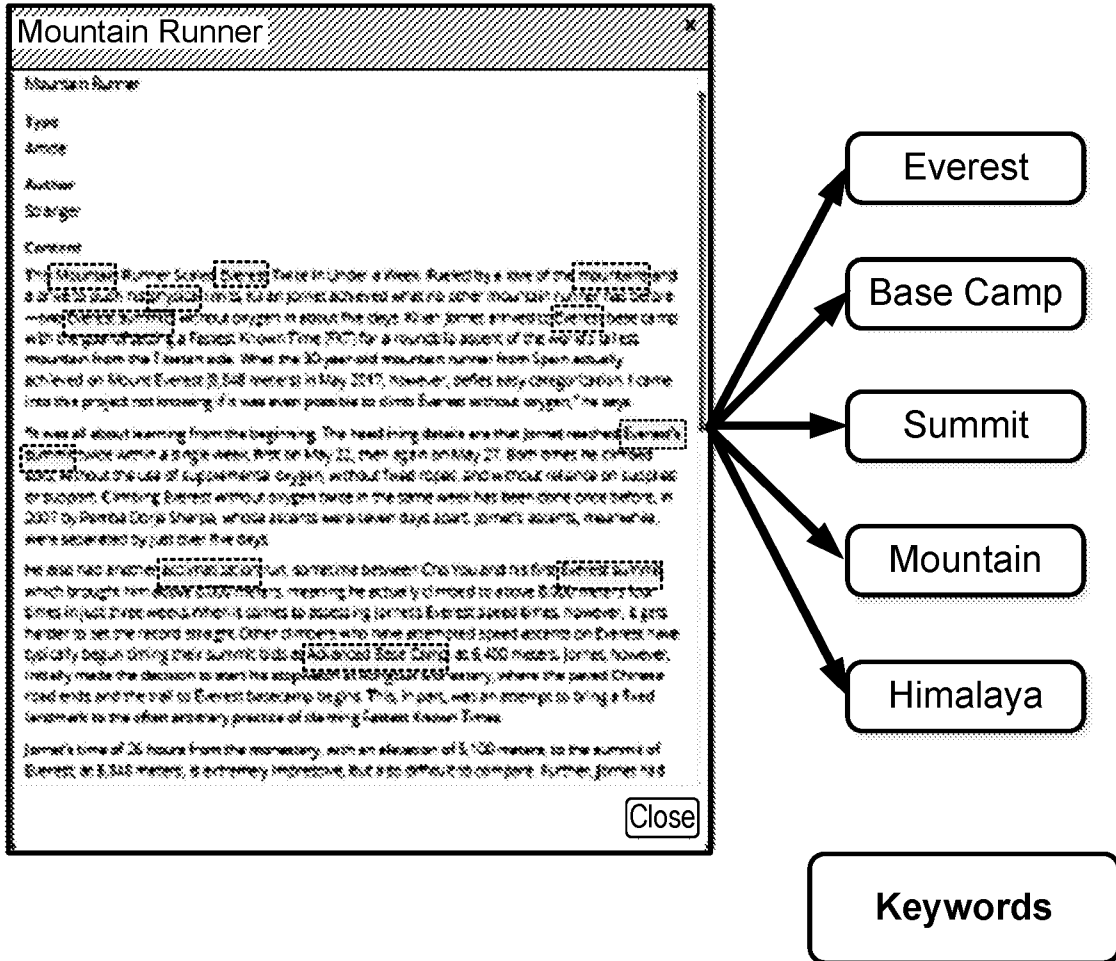


FIG. 8

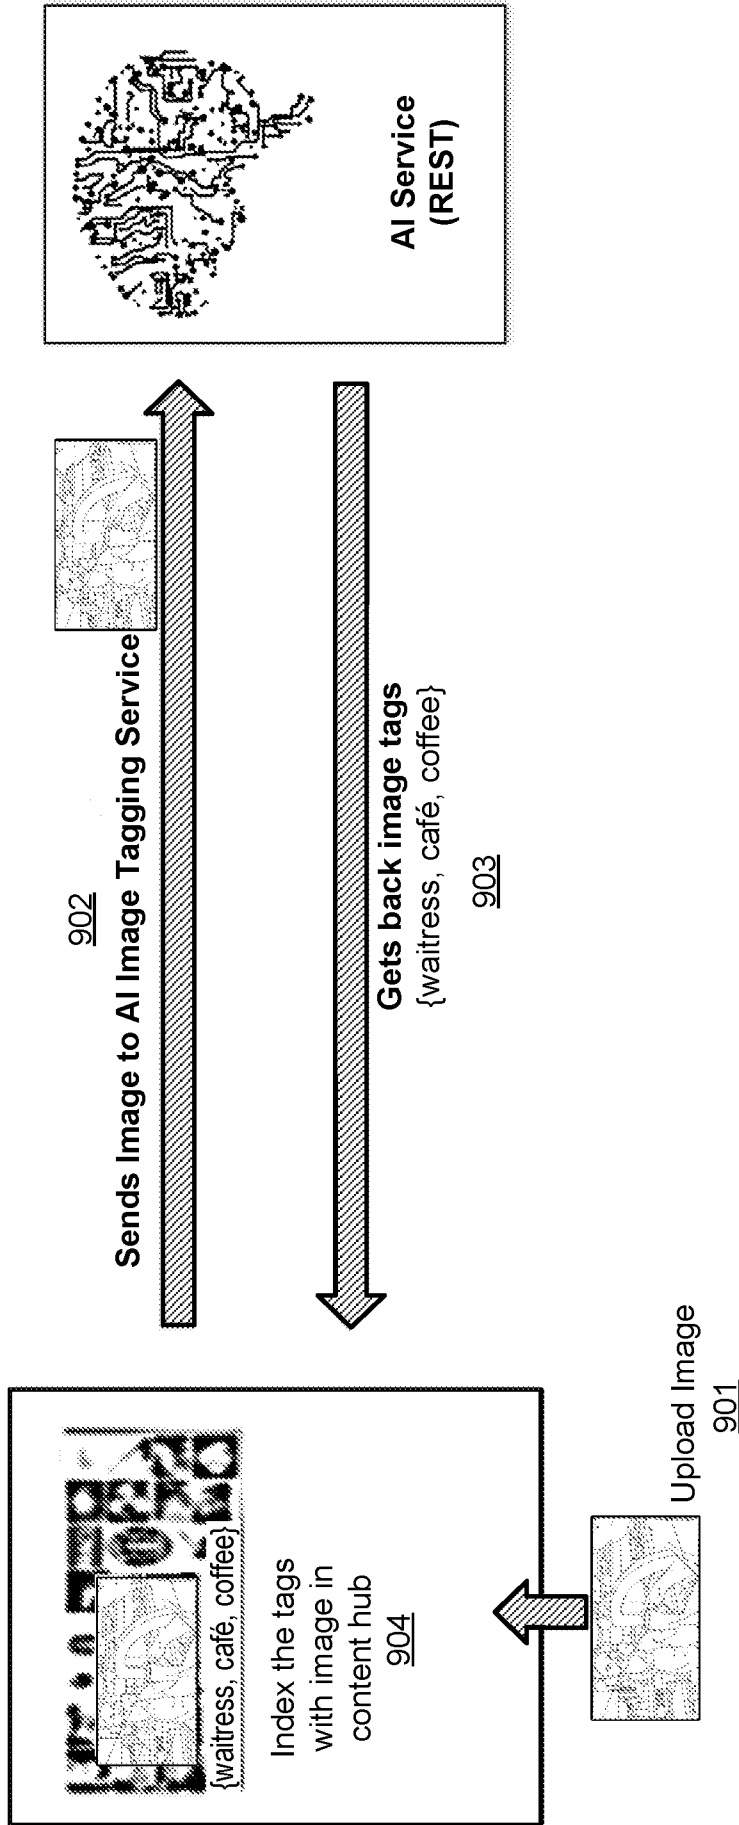


FIG. 9

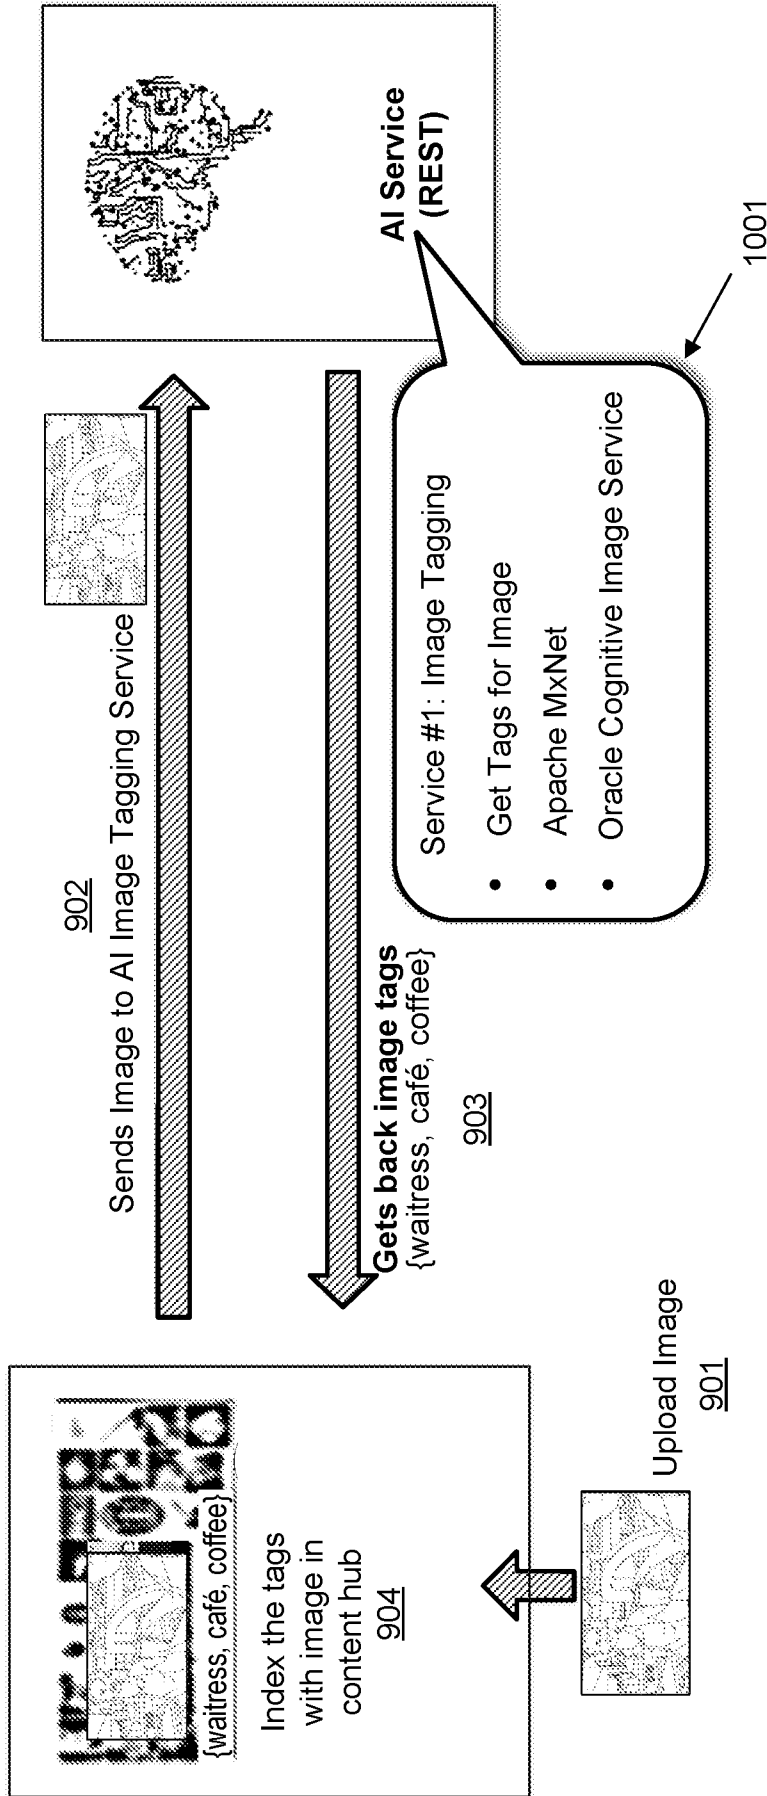


FIG. 10

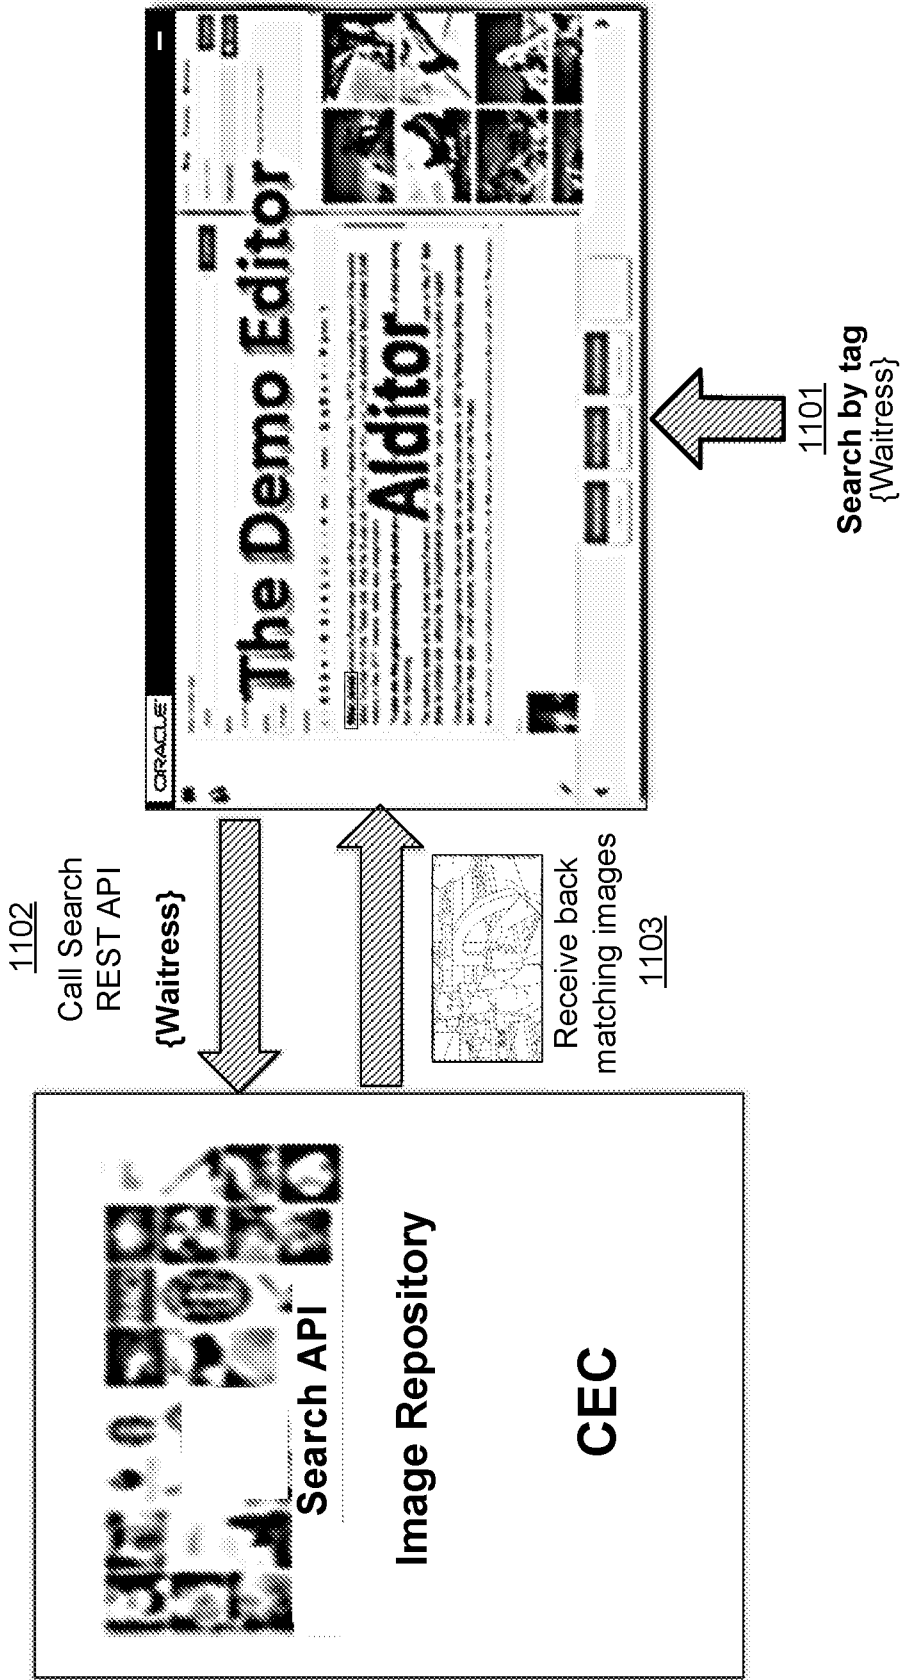
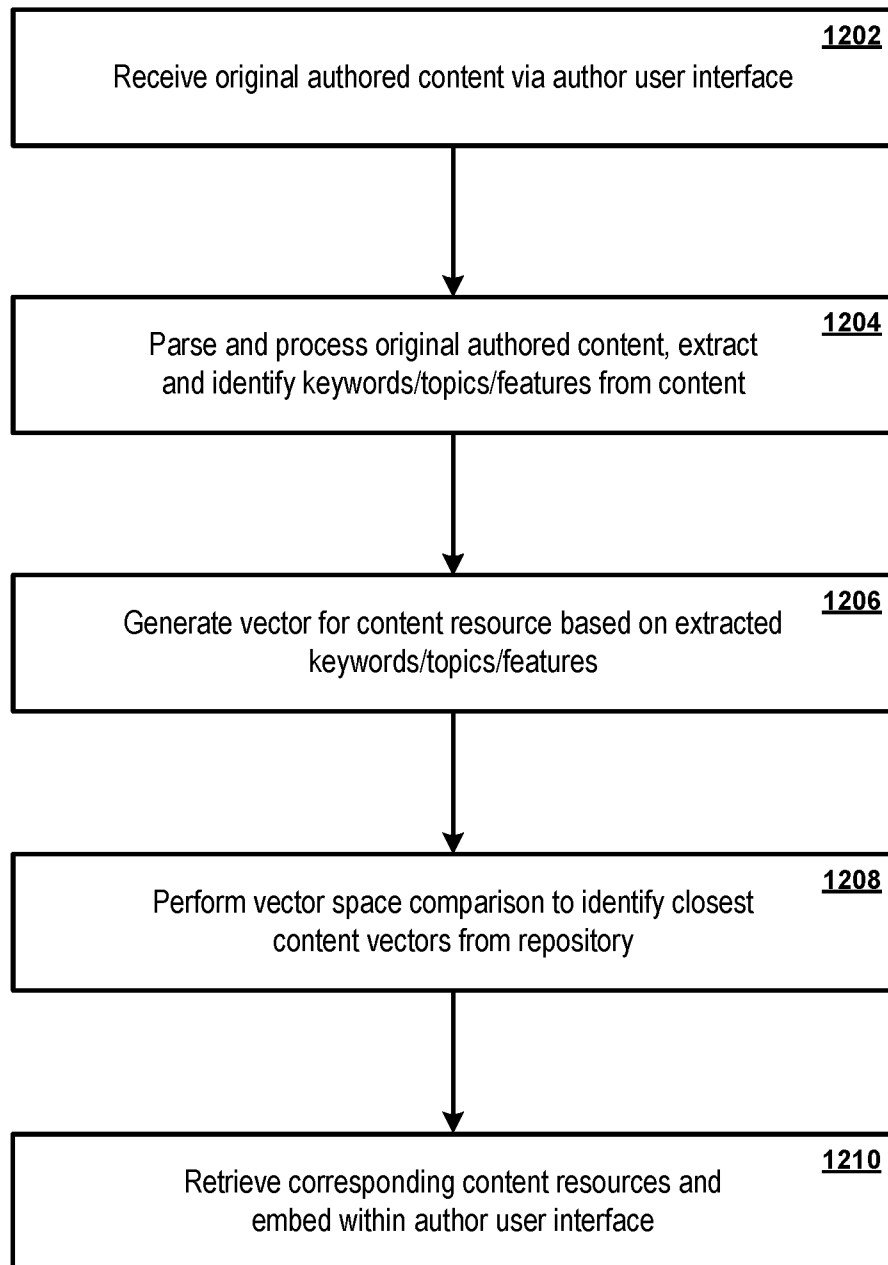


FIG. 11

**FIG. 12**

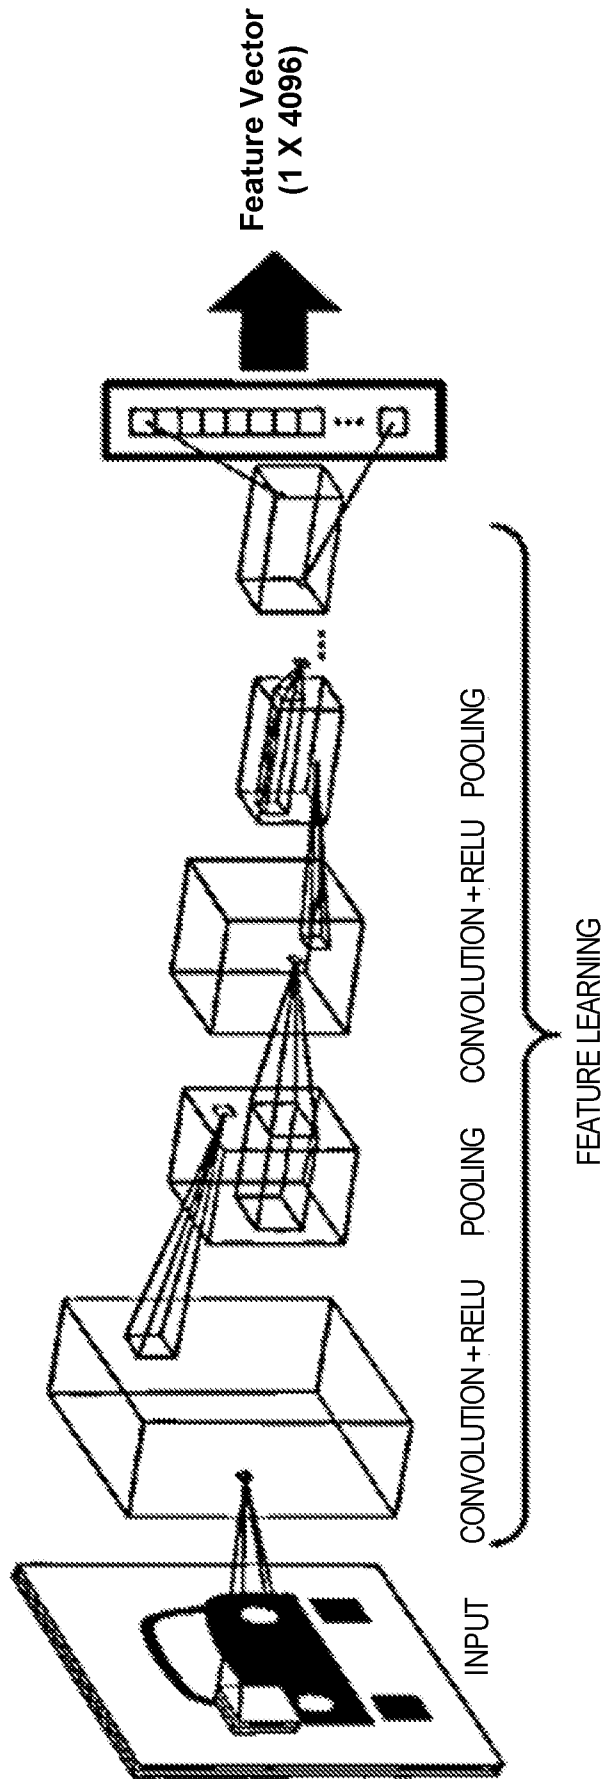


FIG. 13

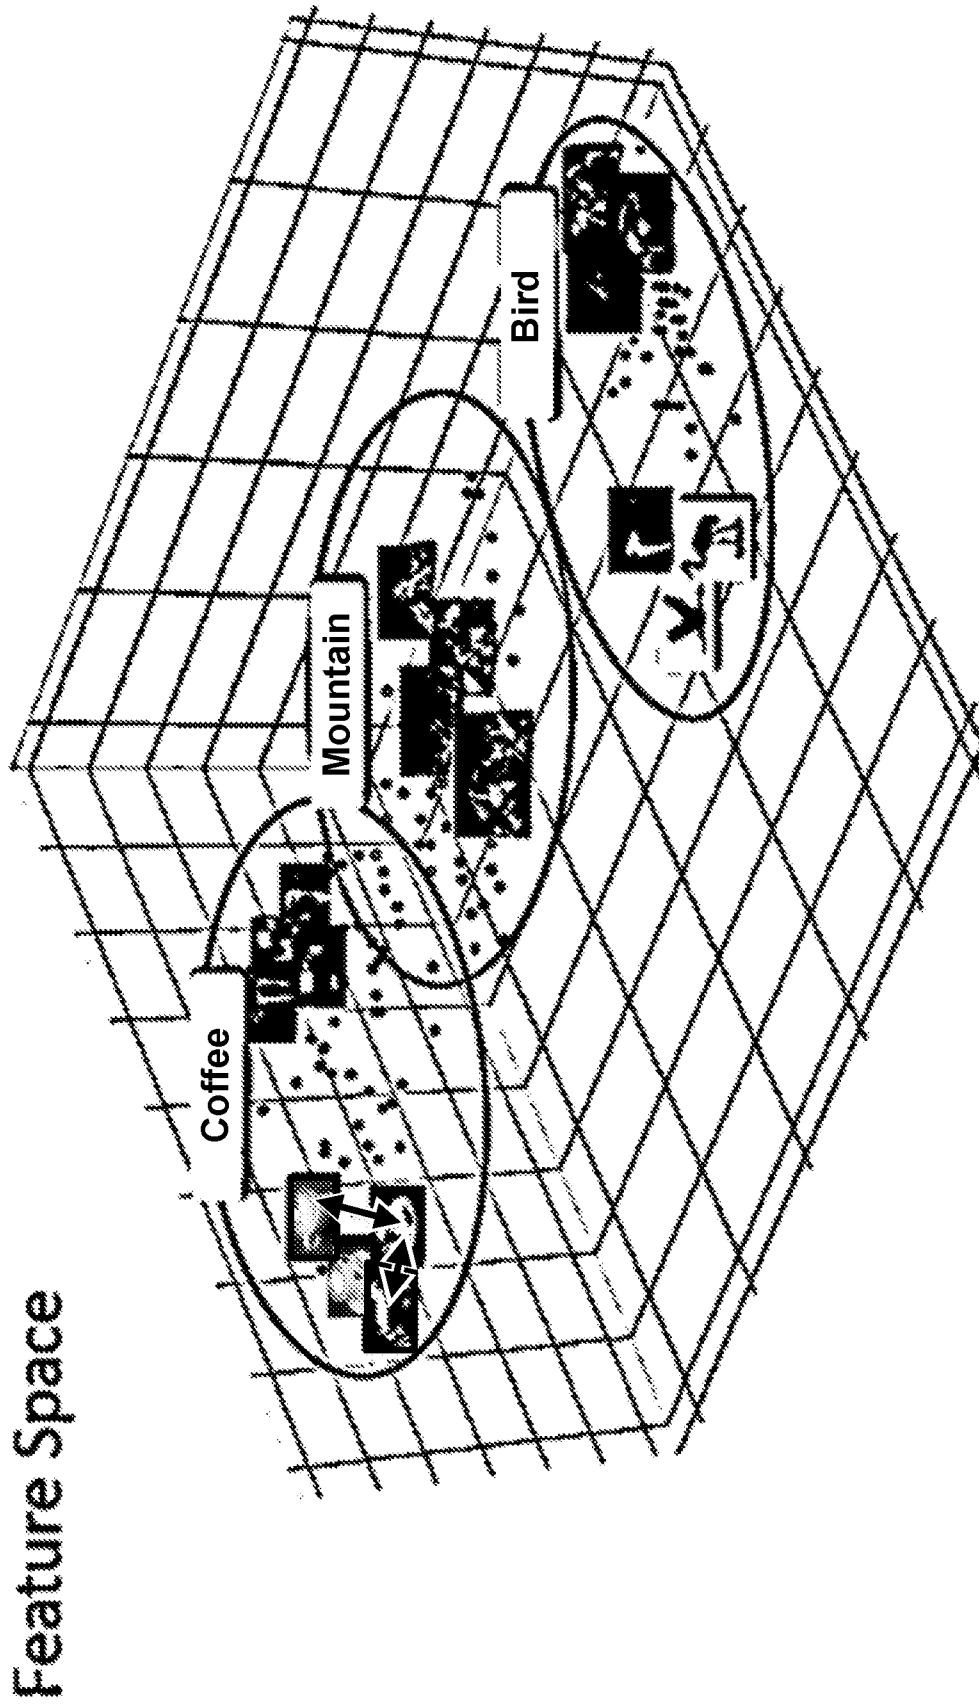


FIG. 14

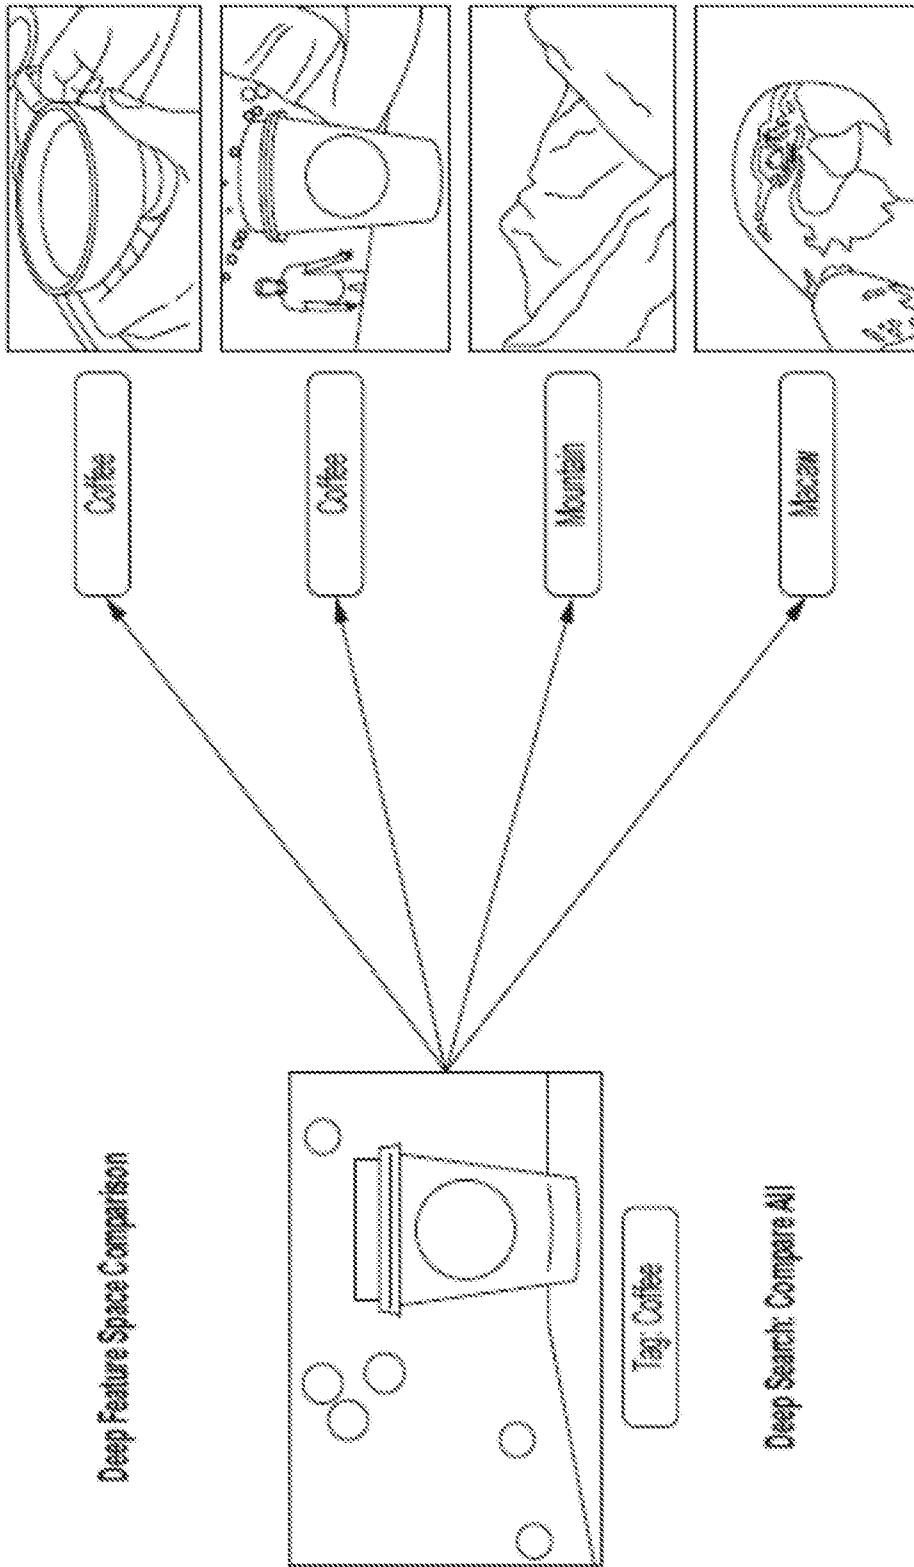
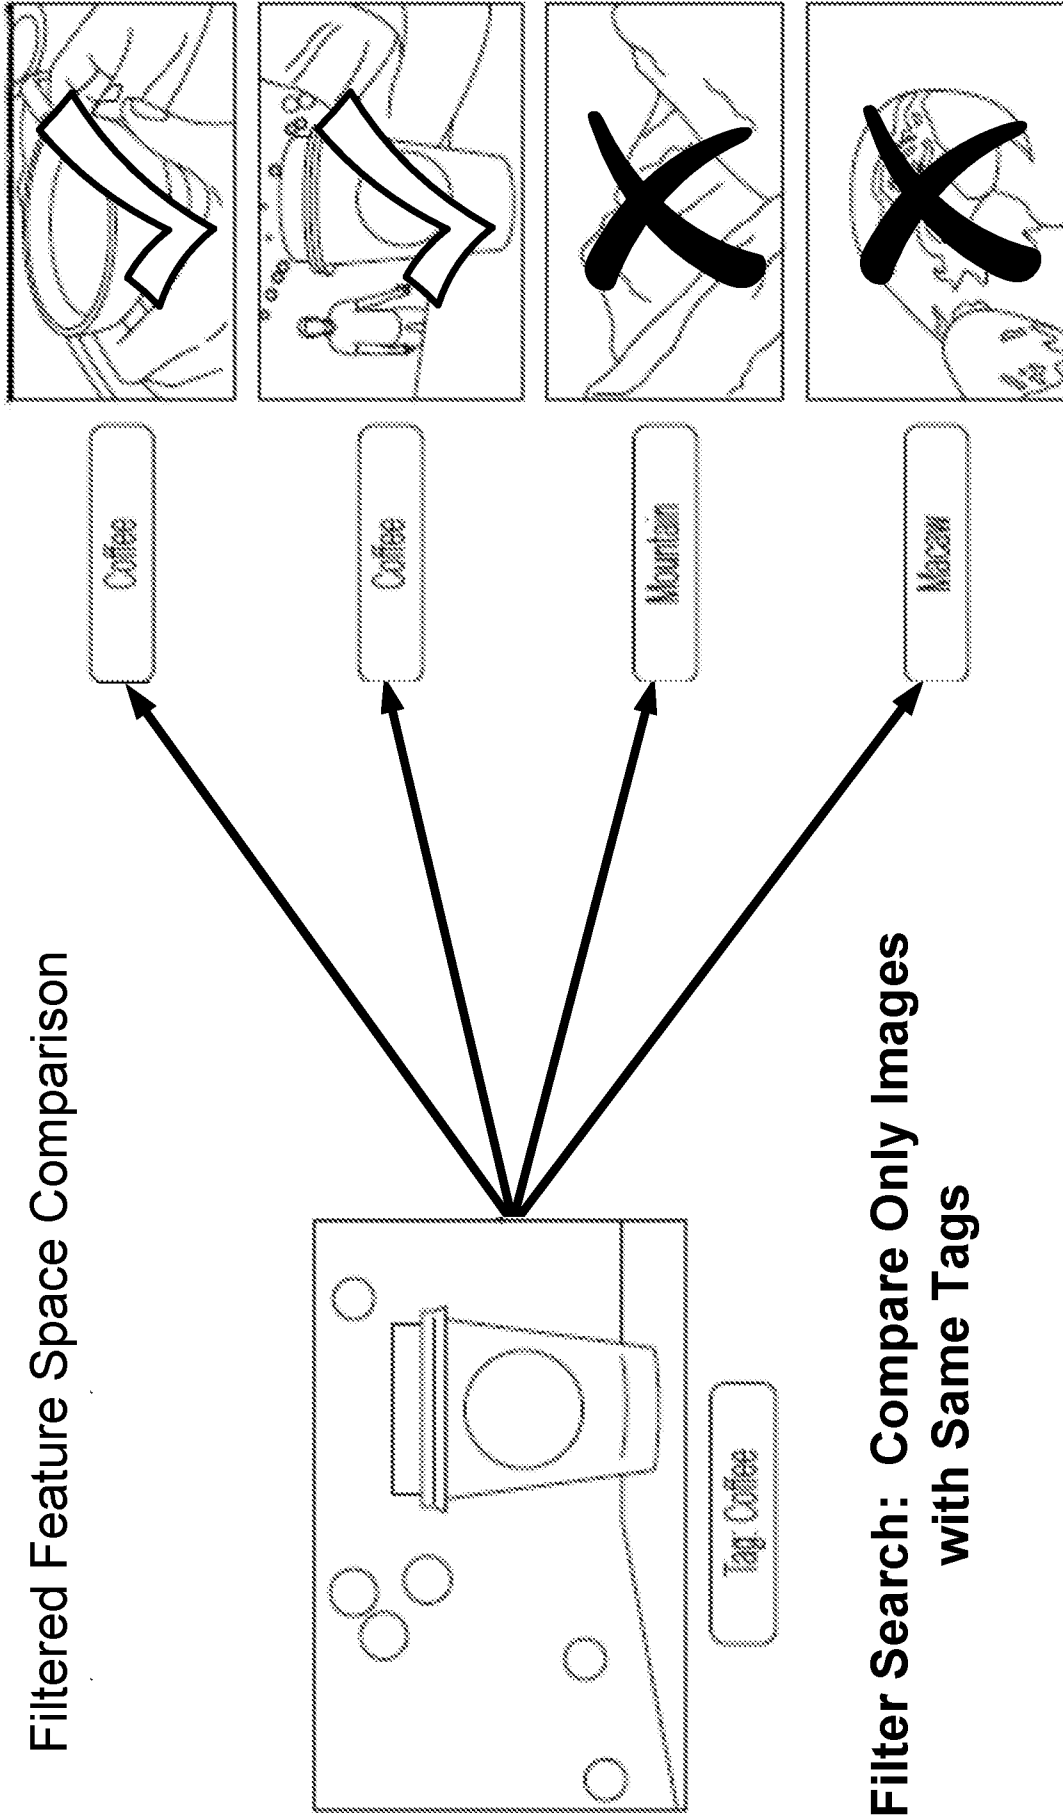


FIG. 15



Filtered Feature Space Comparison

Filter Search: Compare Only Images with Same Tags

FIG. 16

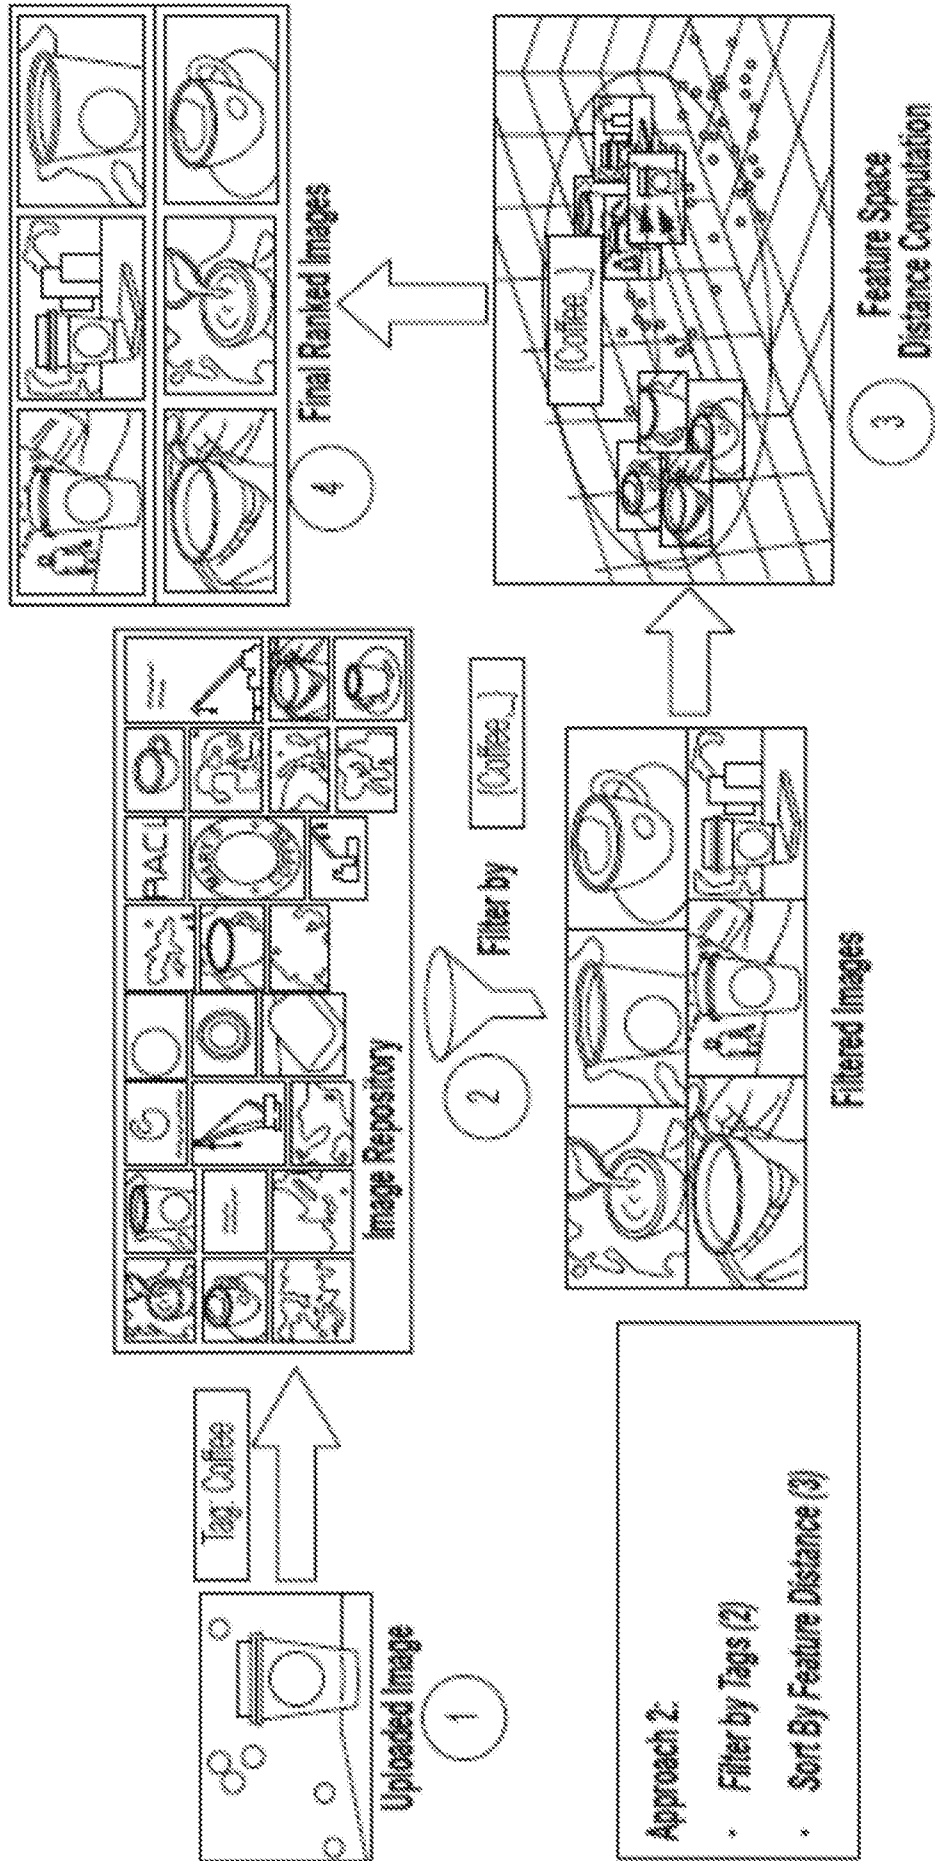


FIG. 17

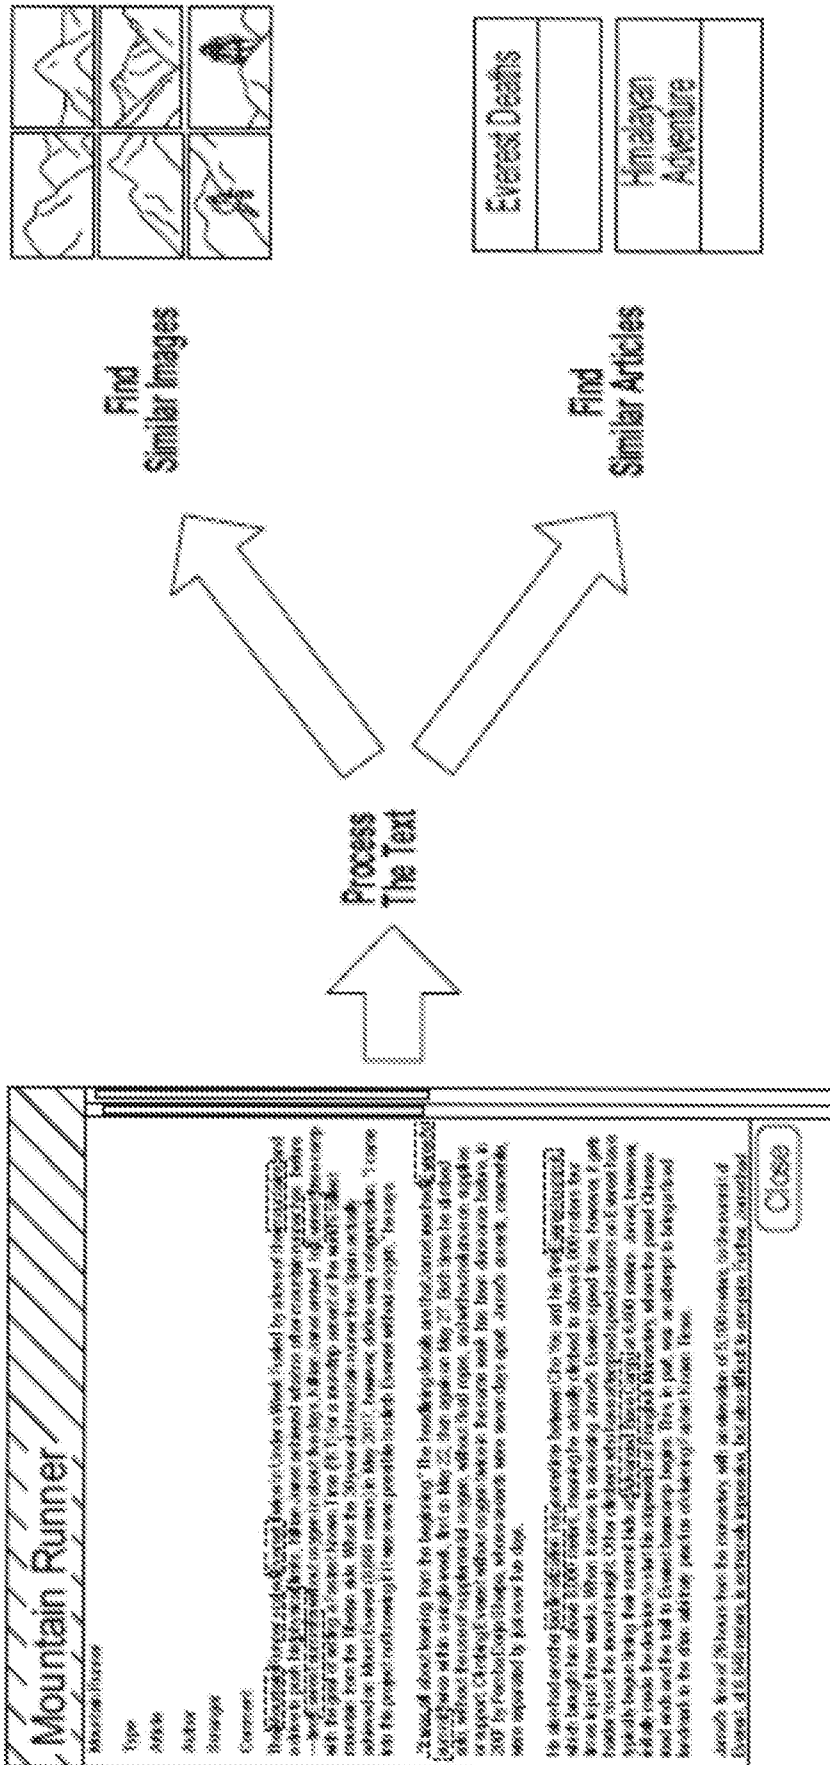


FIG. 18

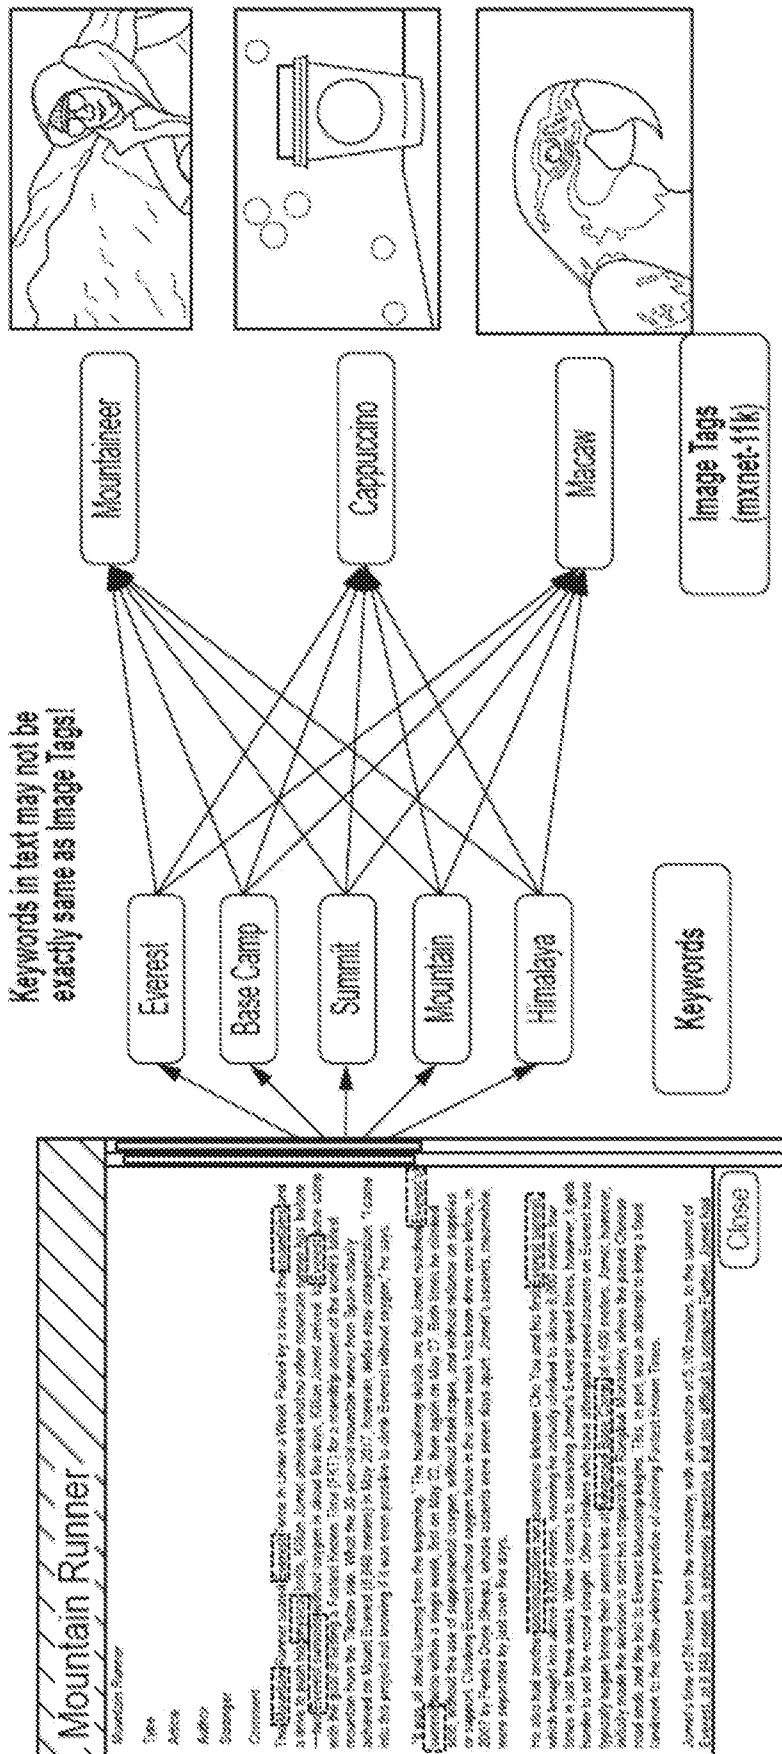


FIG. 19

Word Vectors: Feature Space(3D)

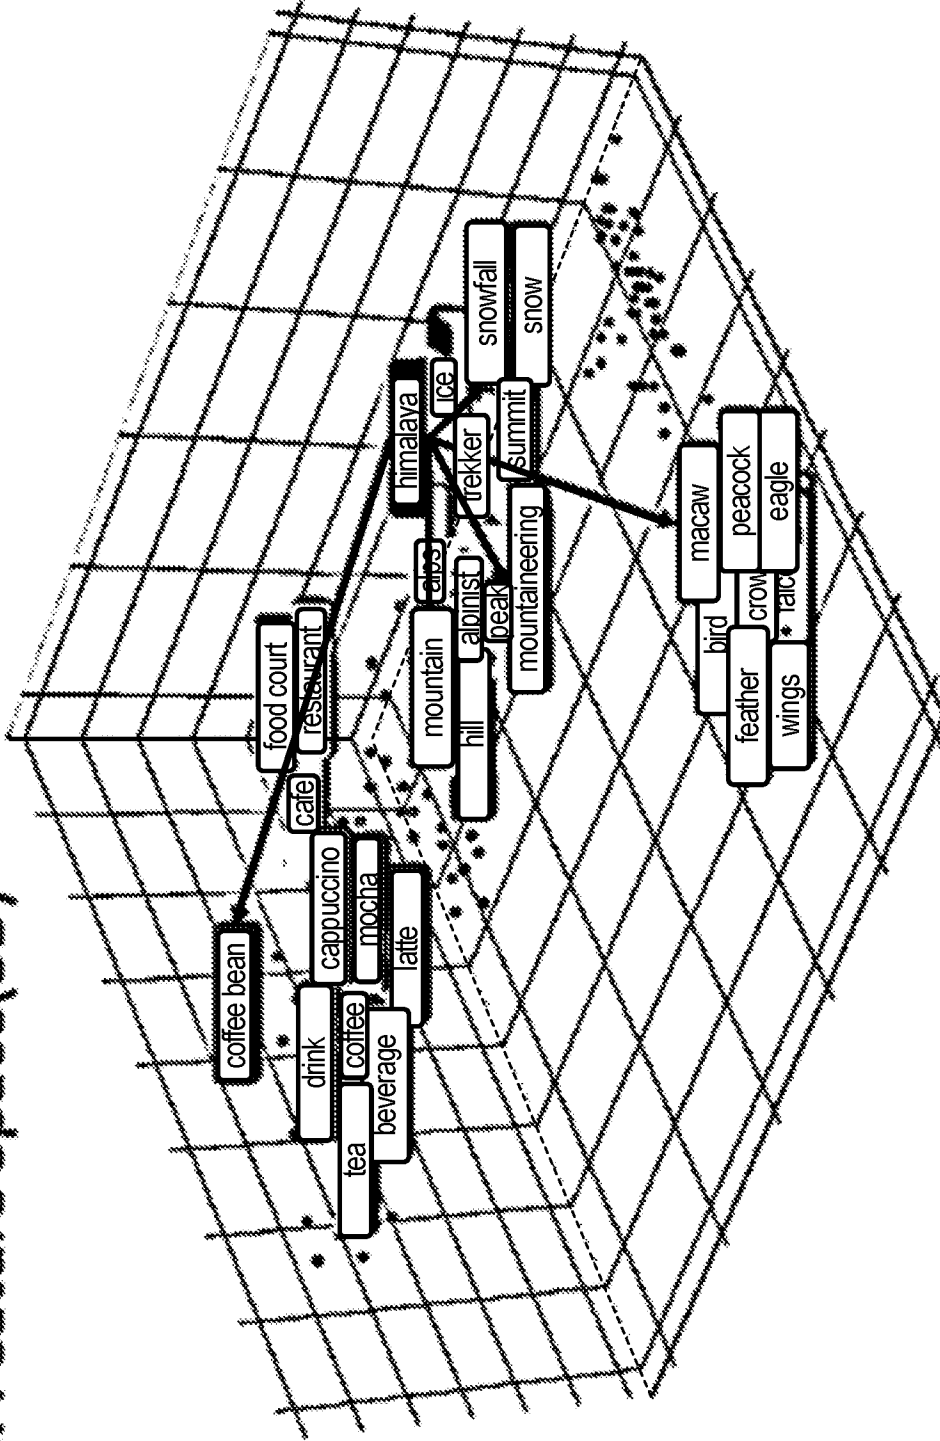


FIG. 20

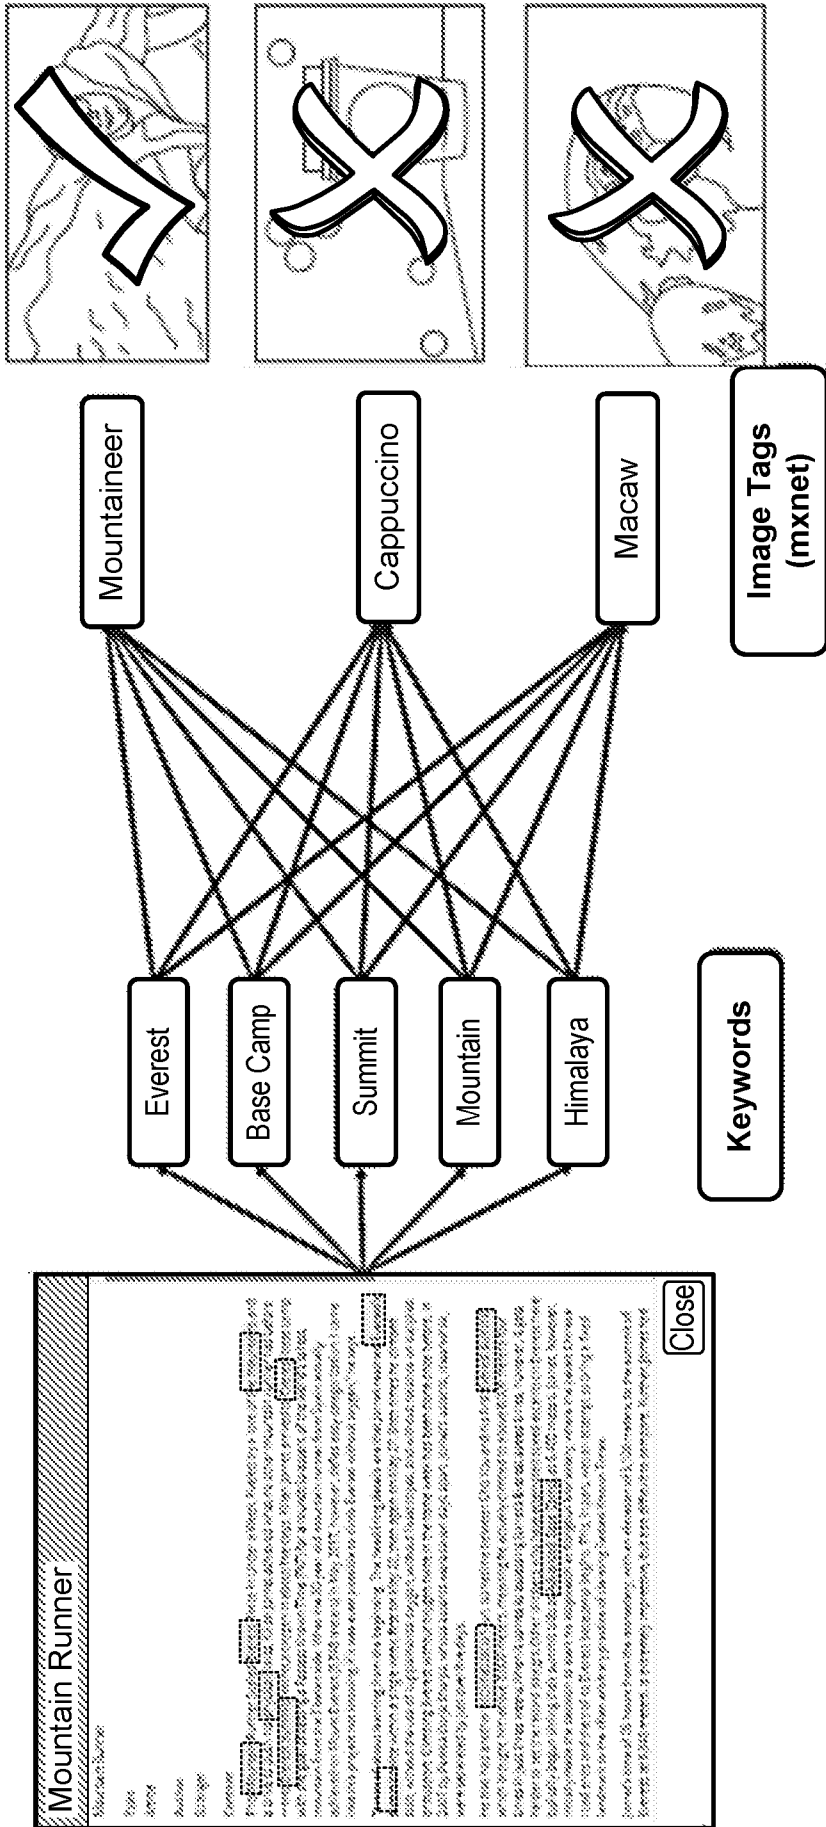


FIG. 21

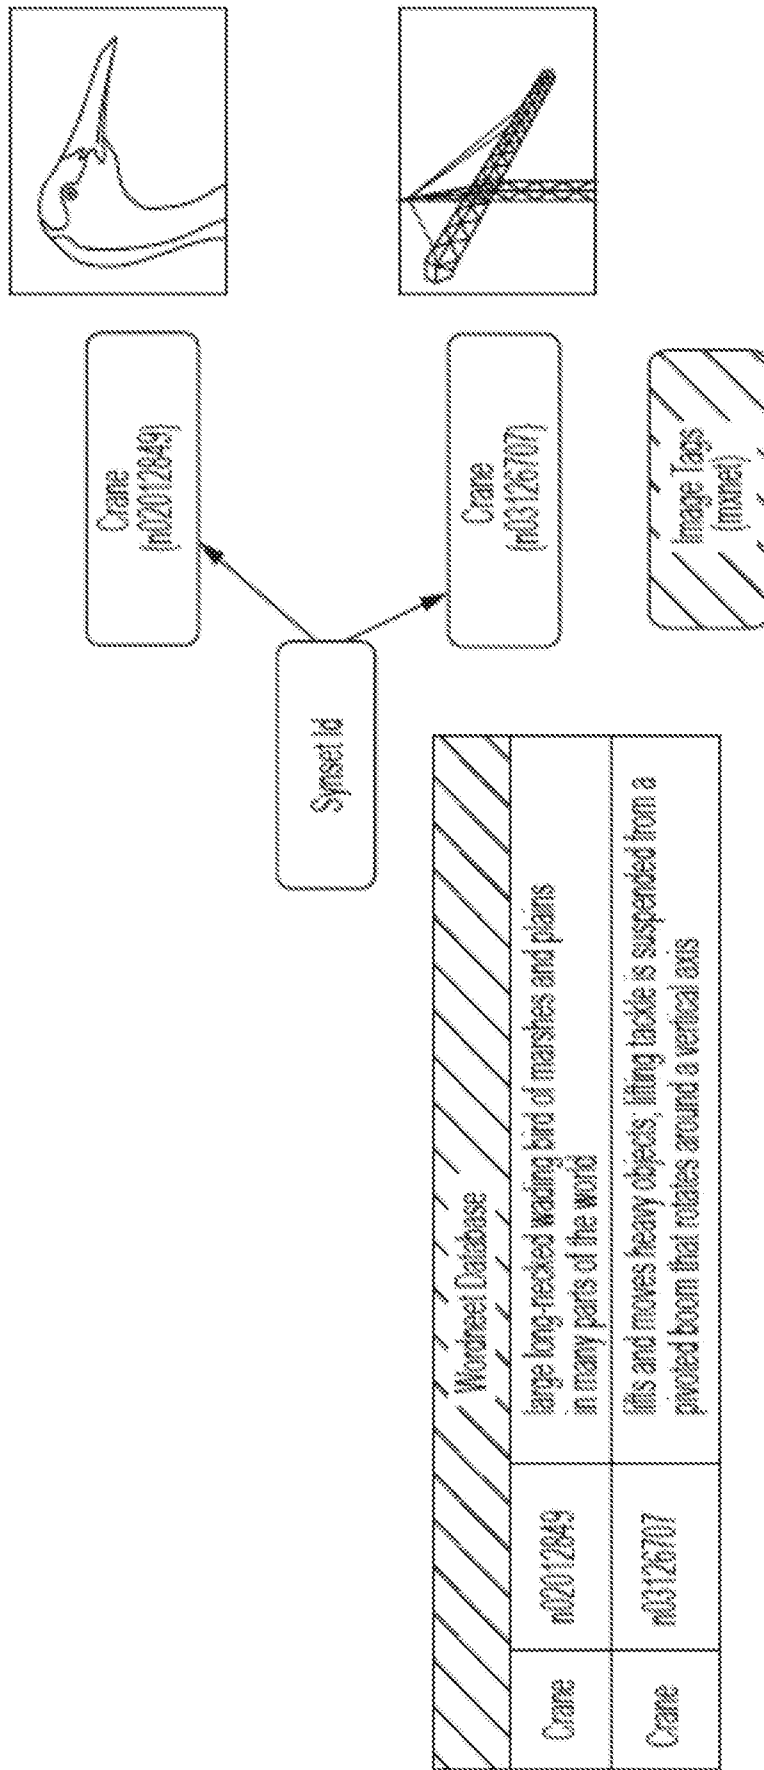


FIG. 22

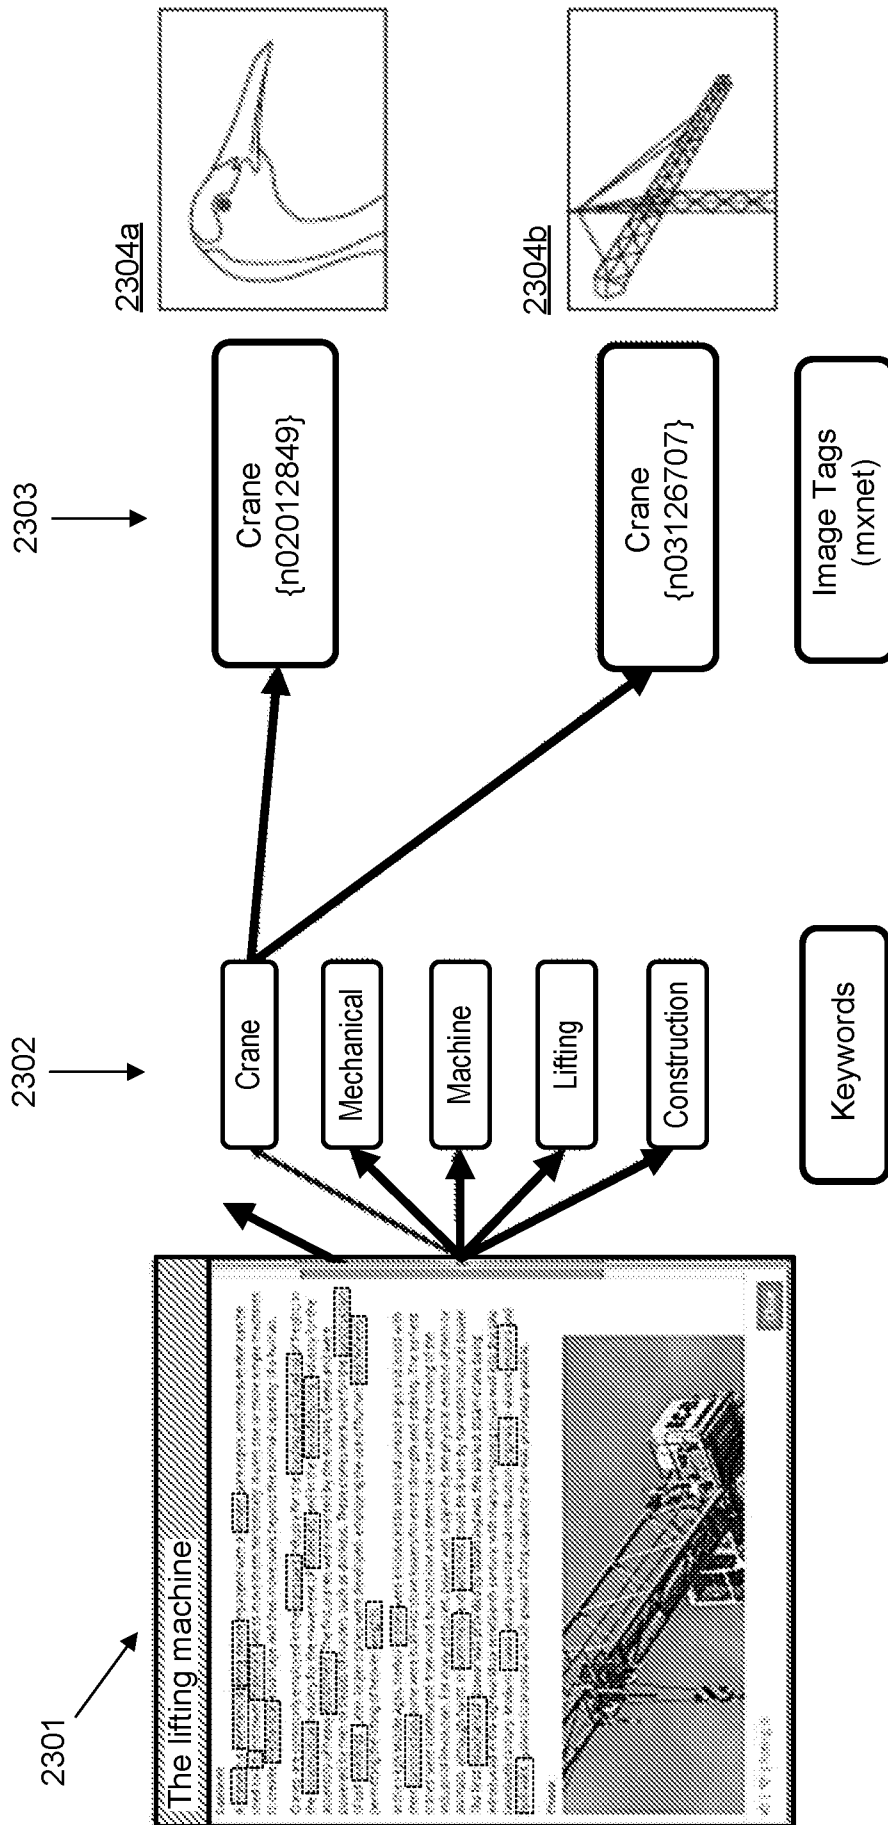


FIG. 23

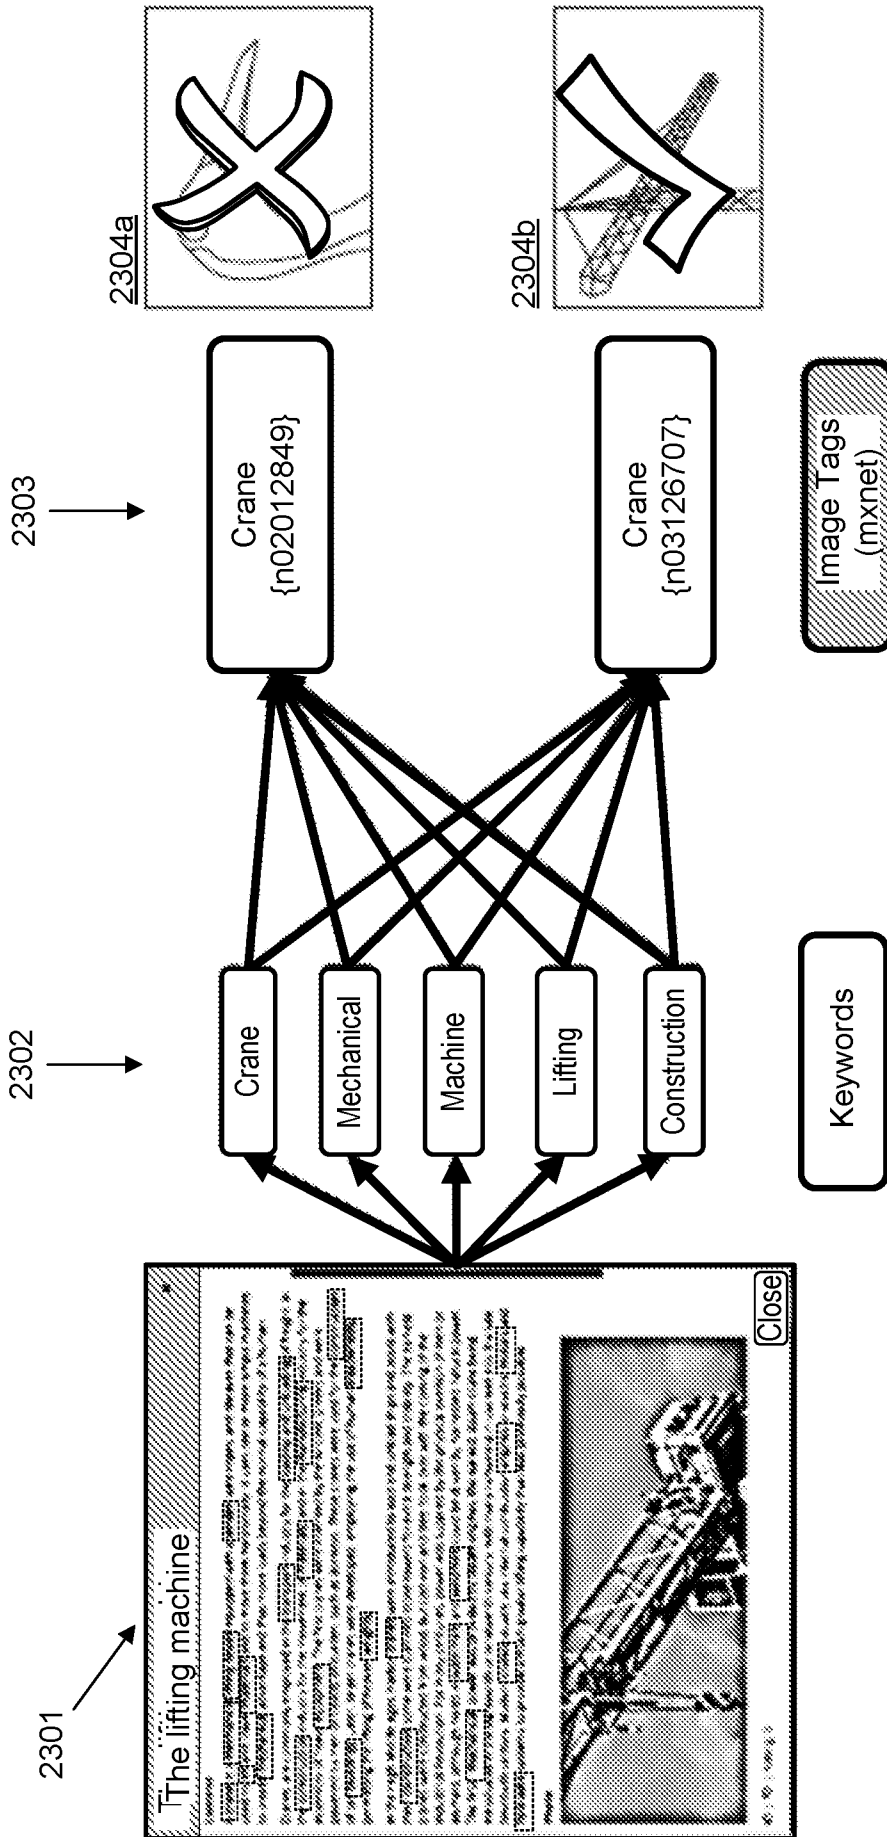


FIG. 24

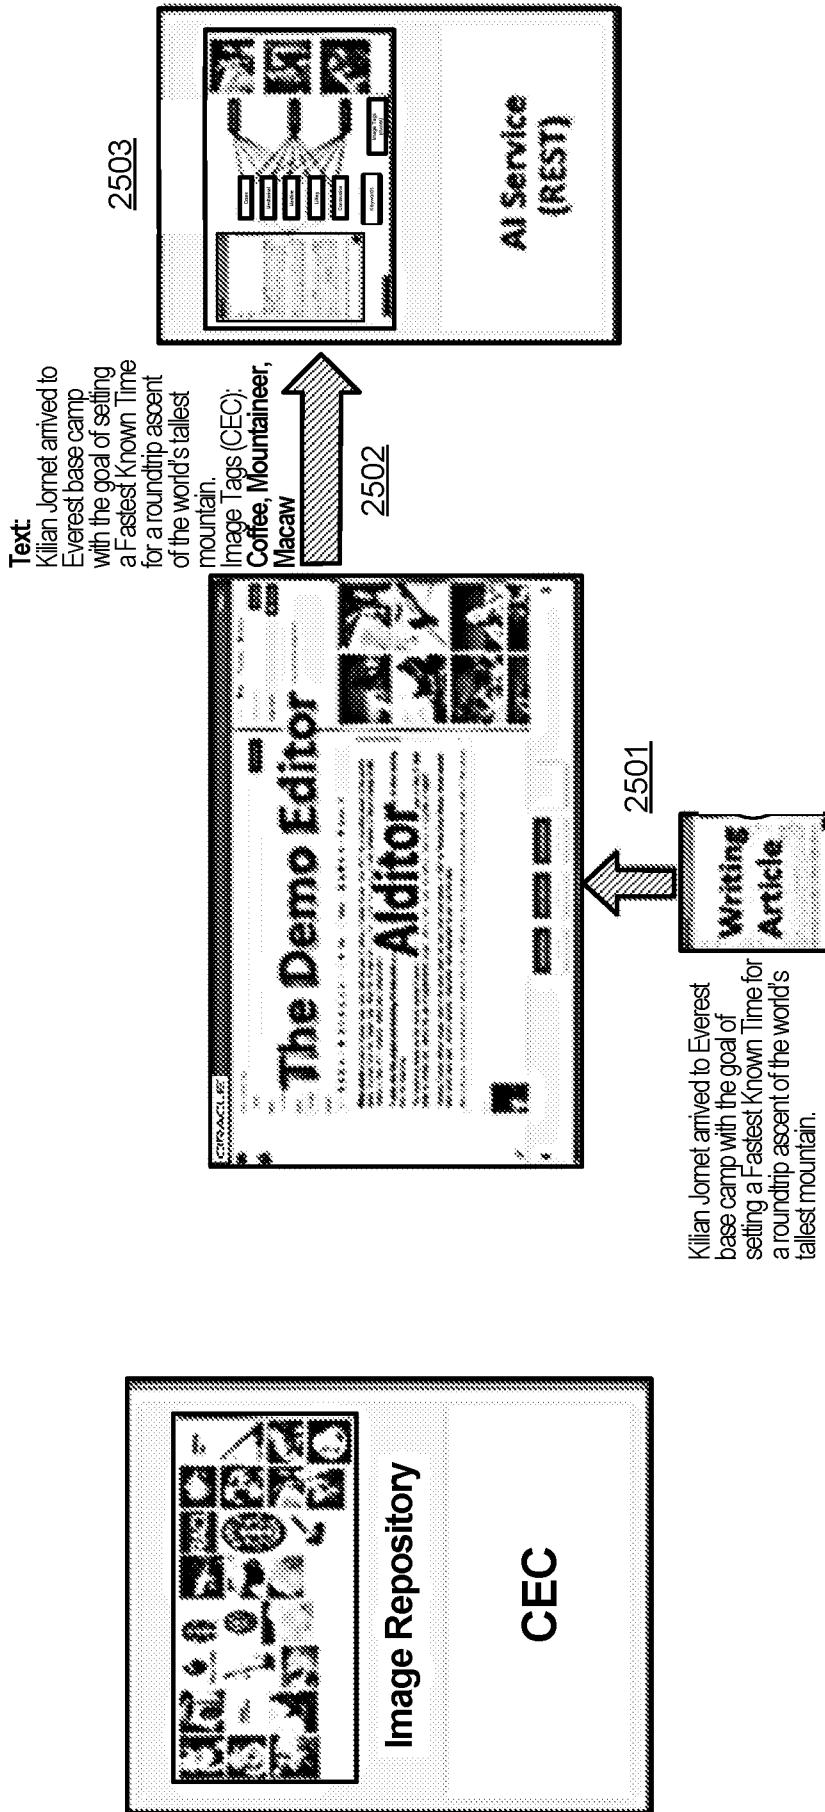


FIG. 25

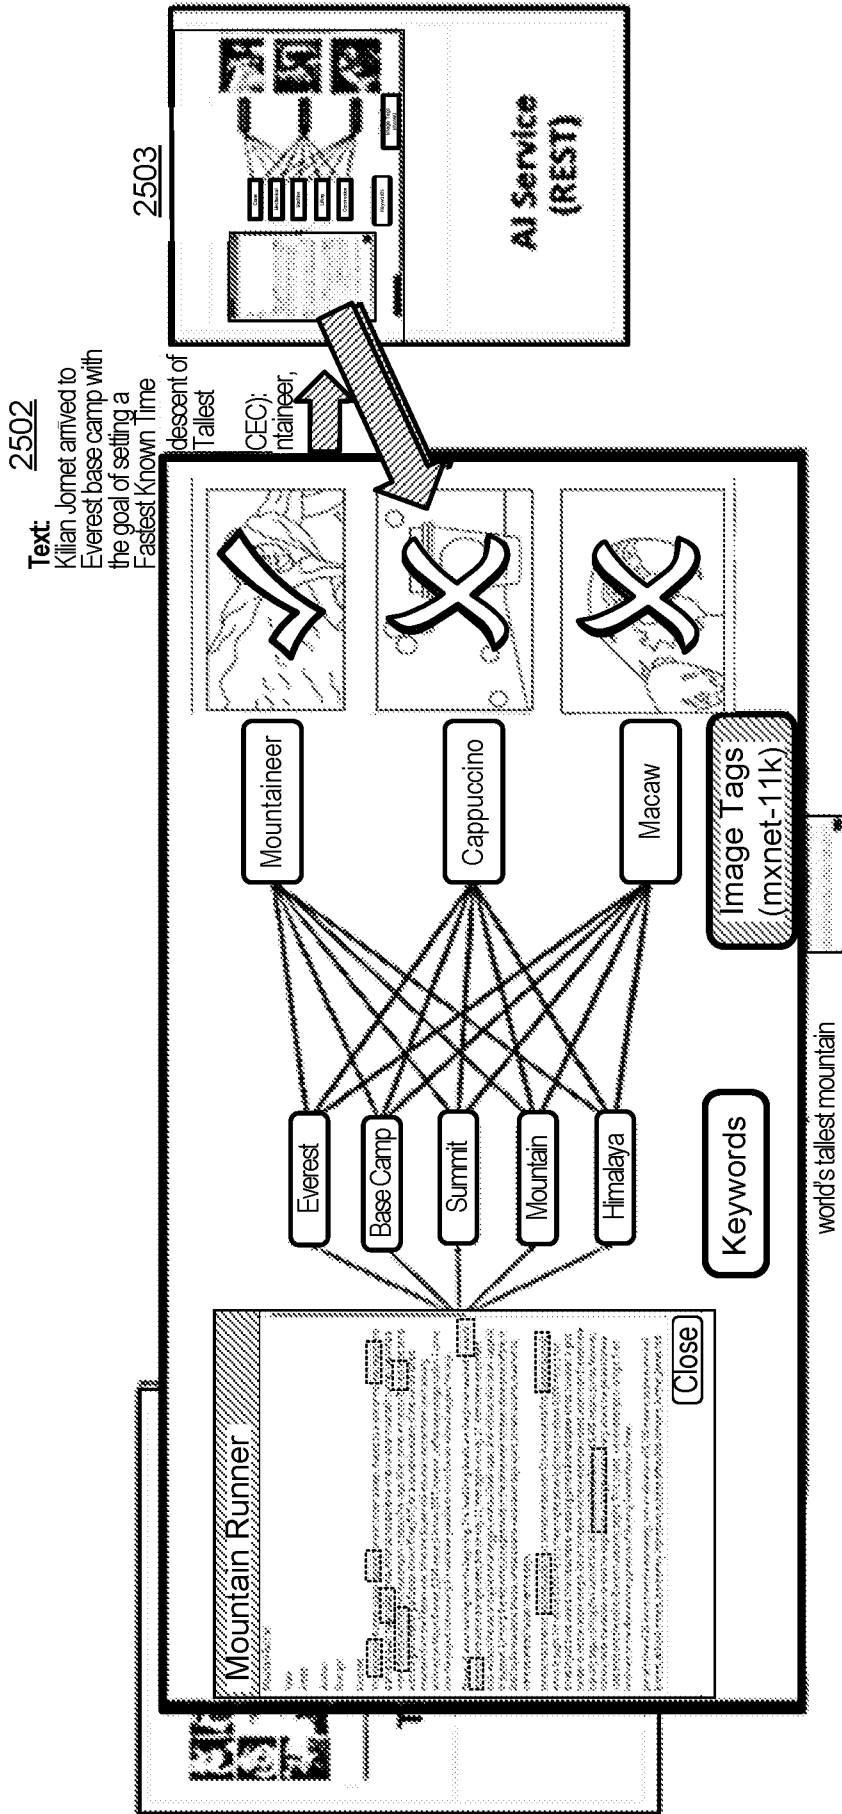


FIG. 26

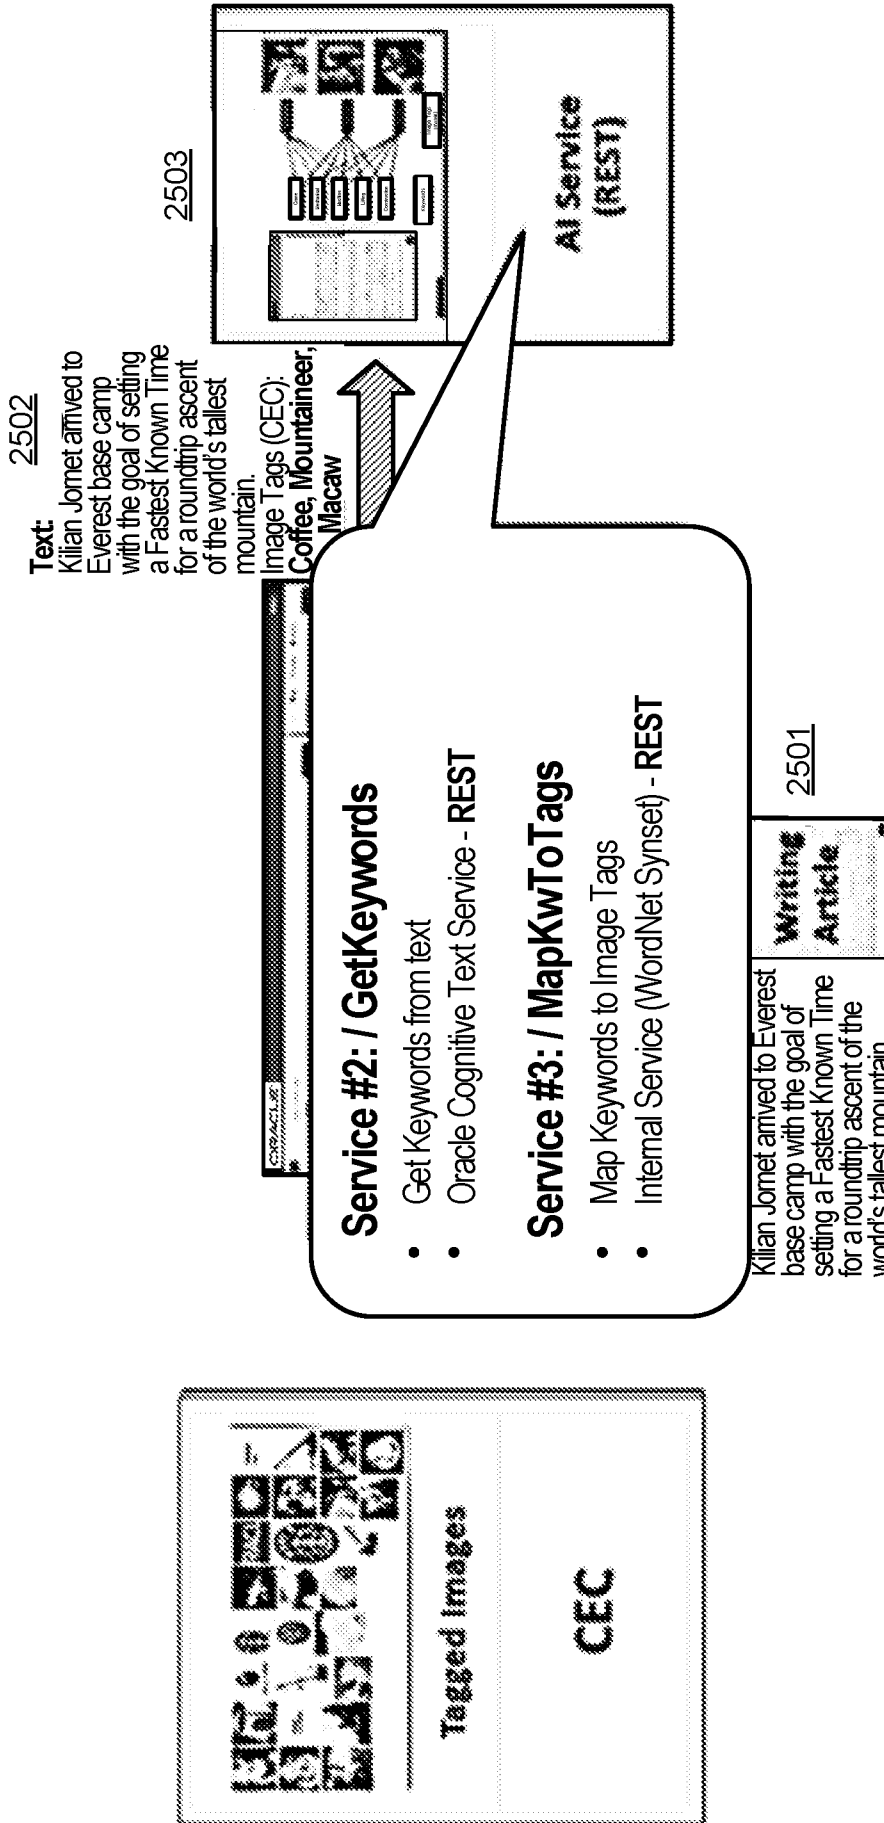


FIG. 27

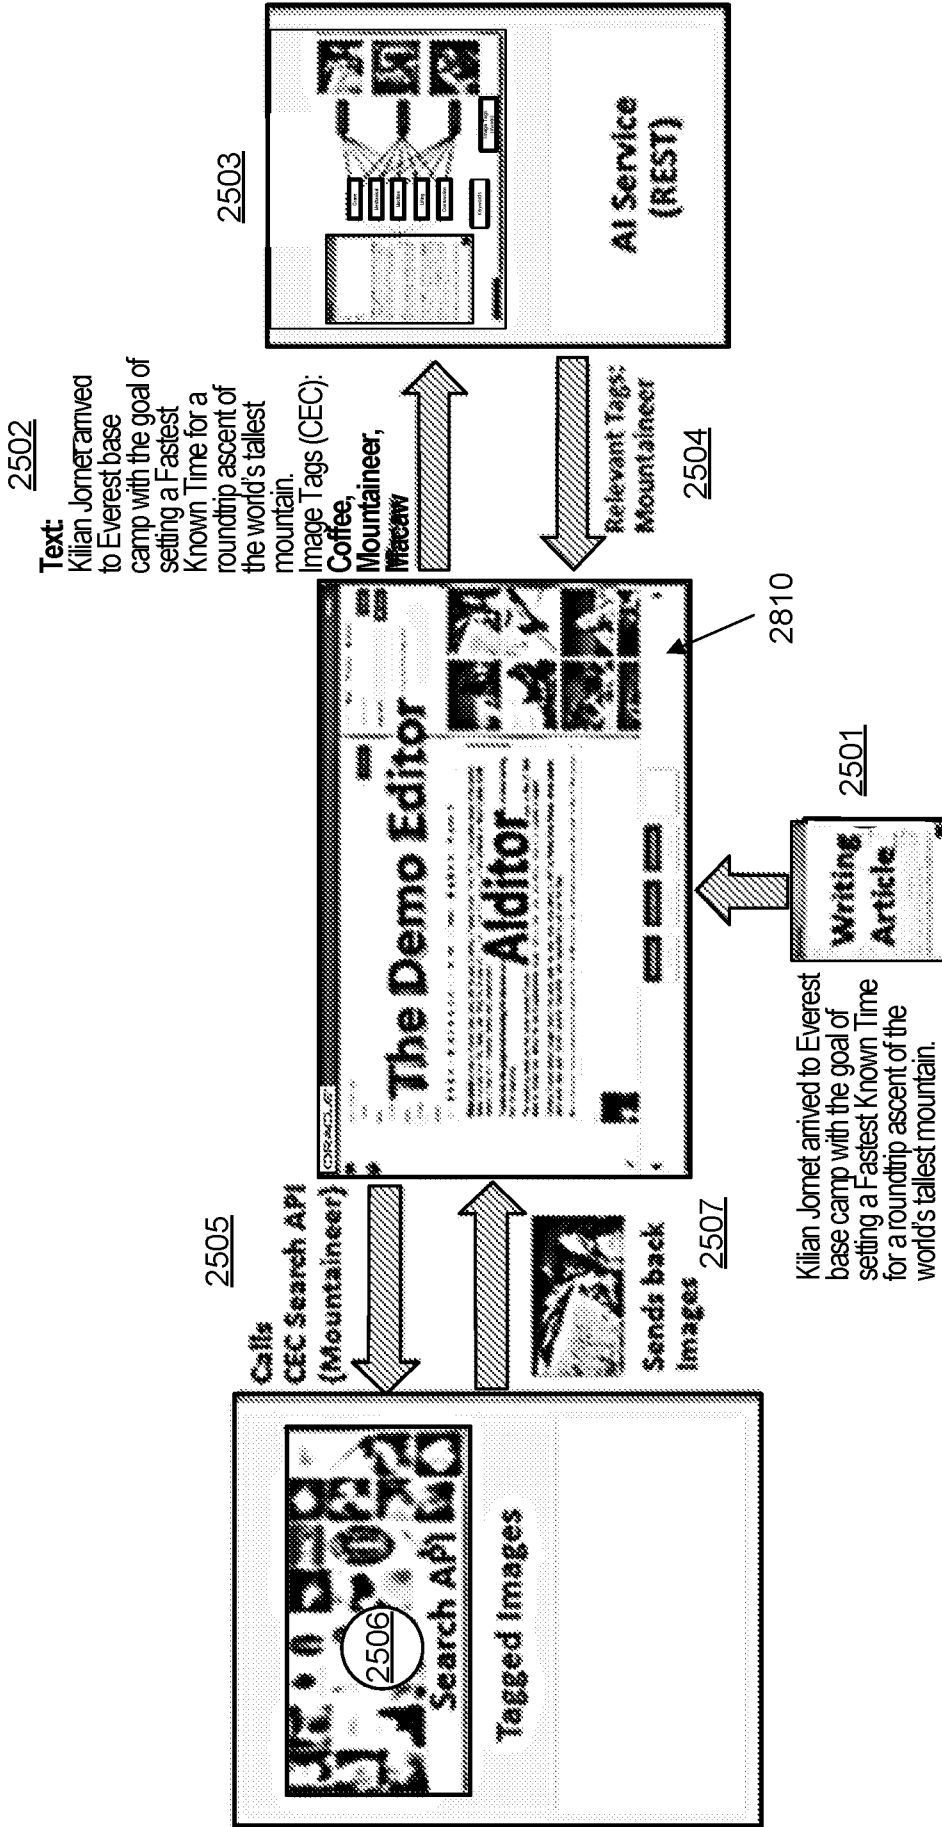


FIG. 28

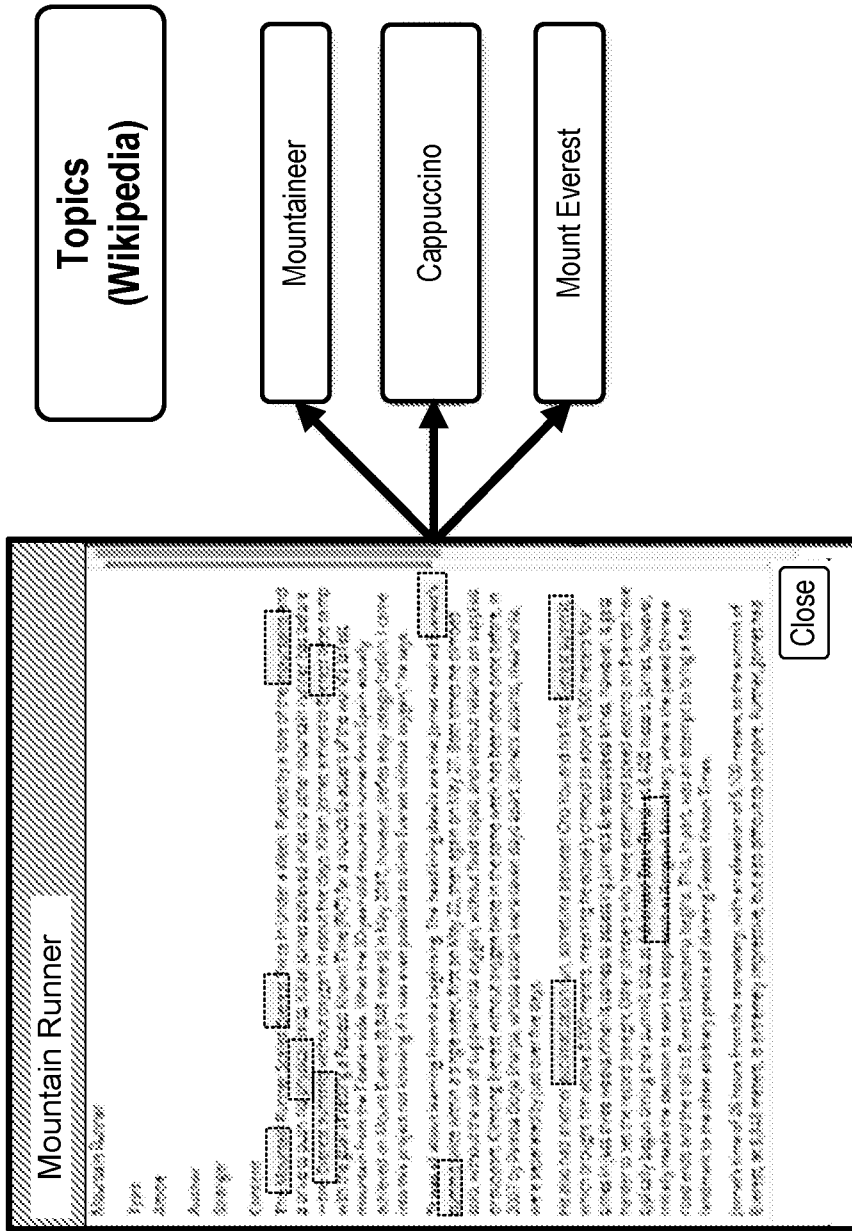


FIG. 29

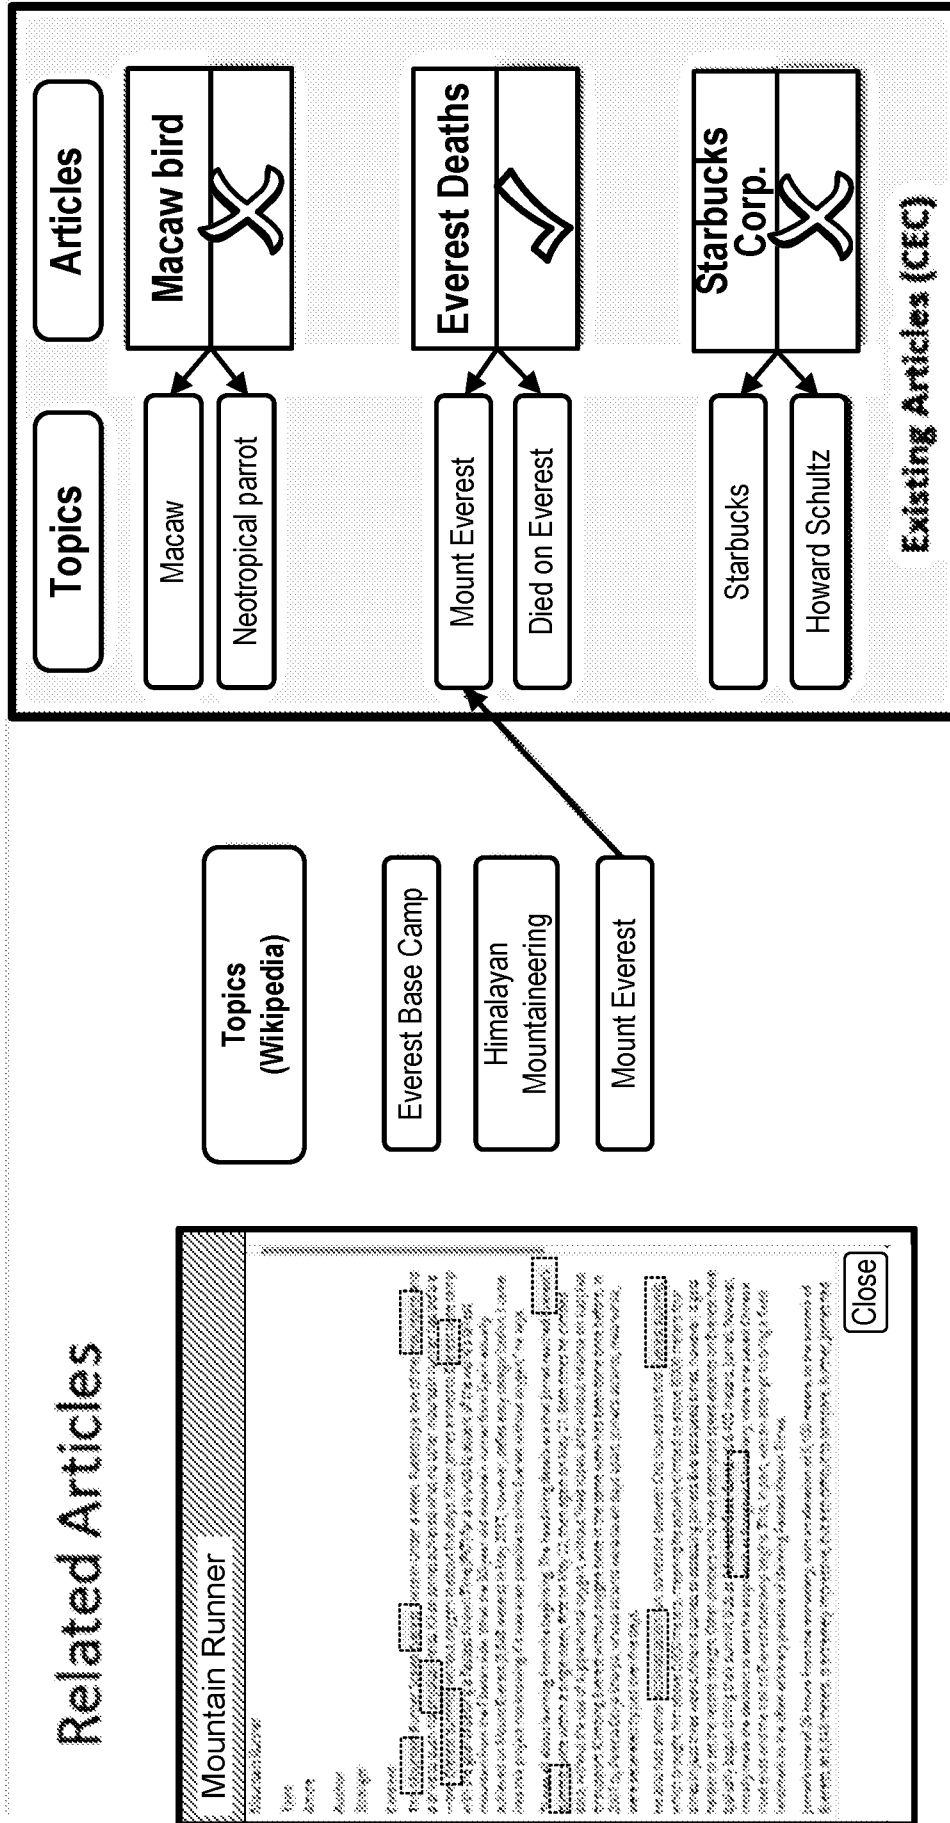


FIG. 30

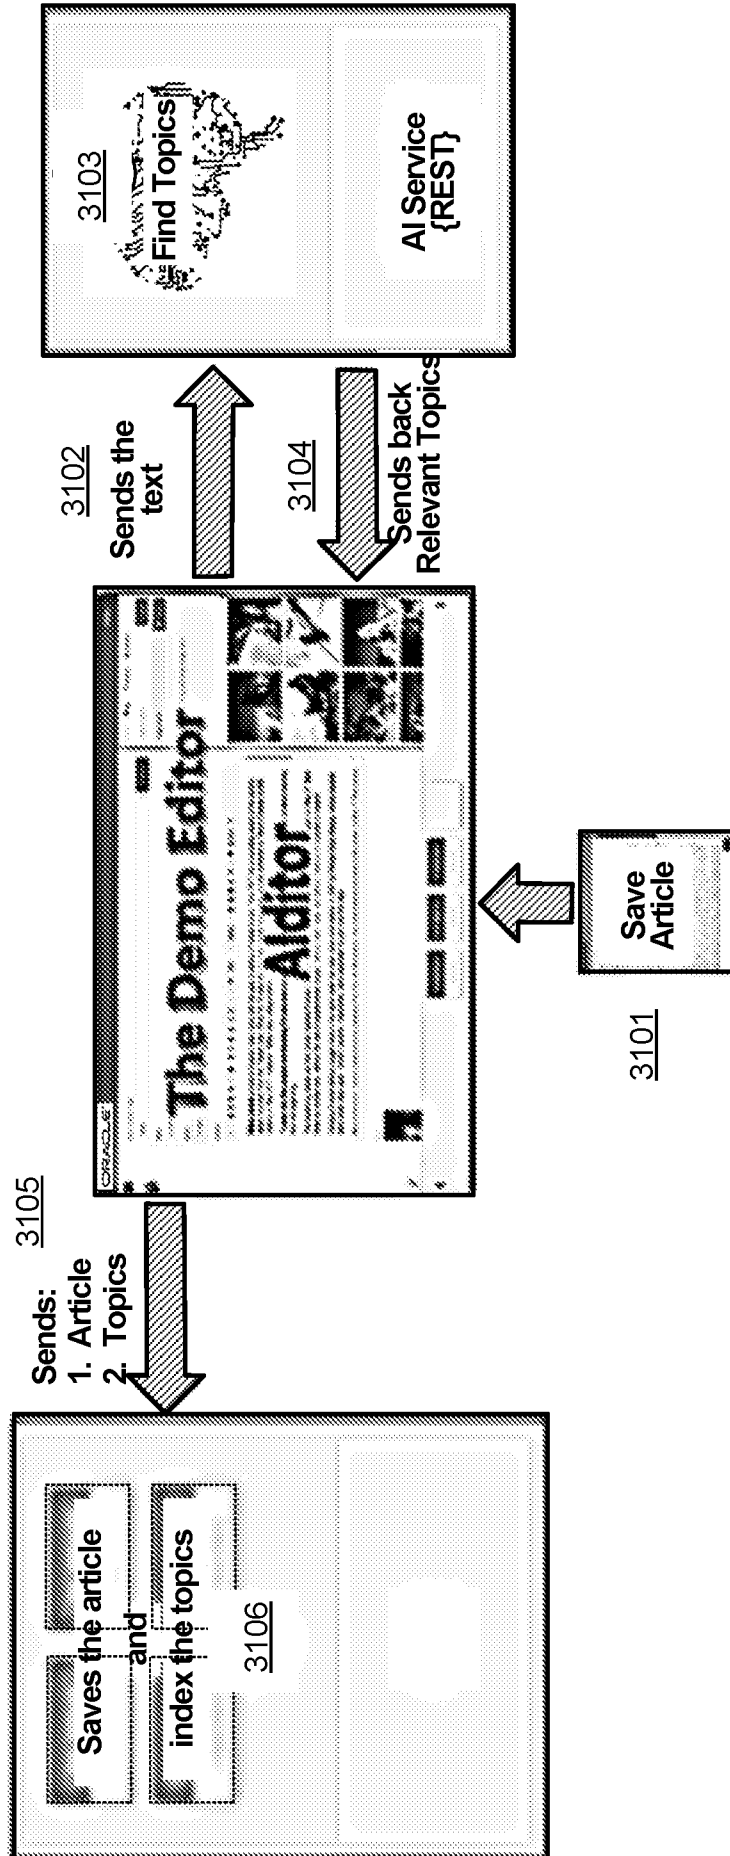


FIG. 31

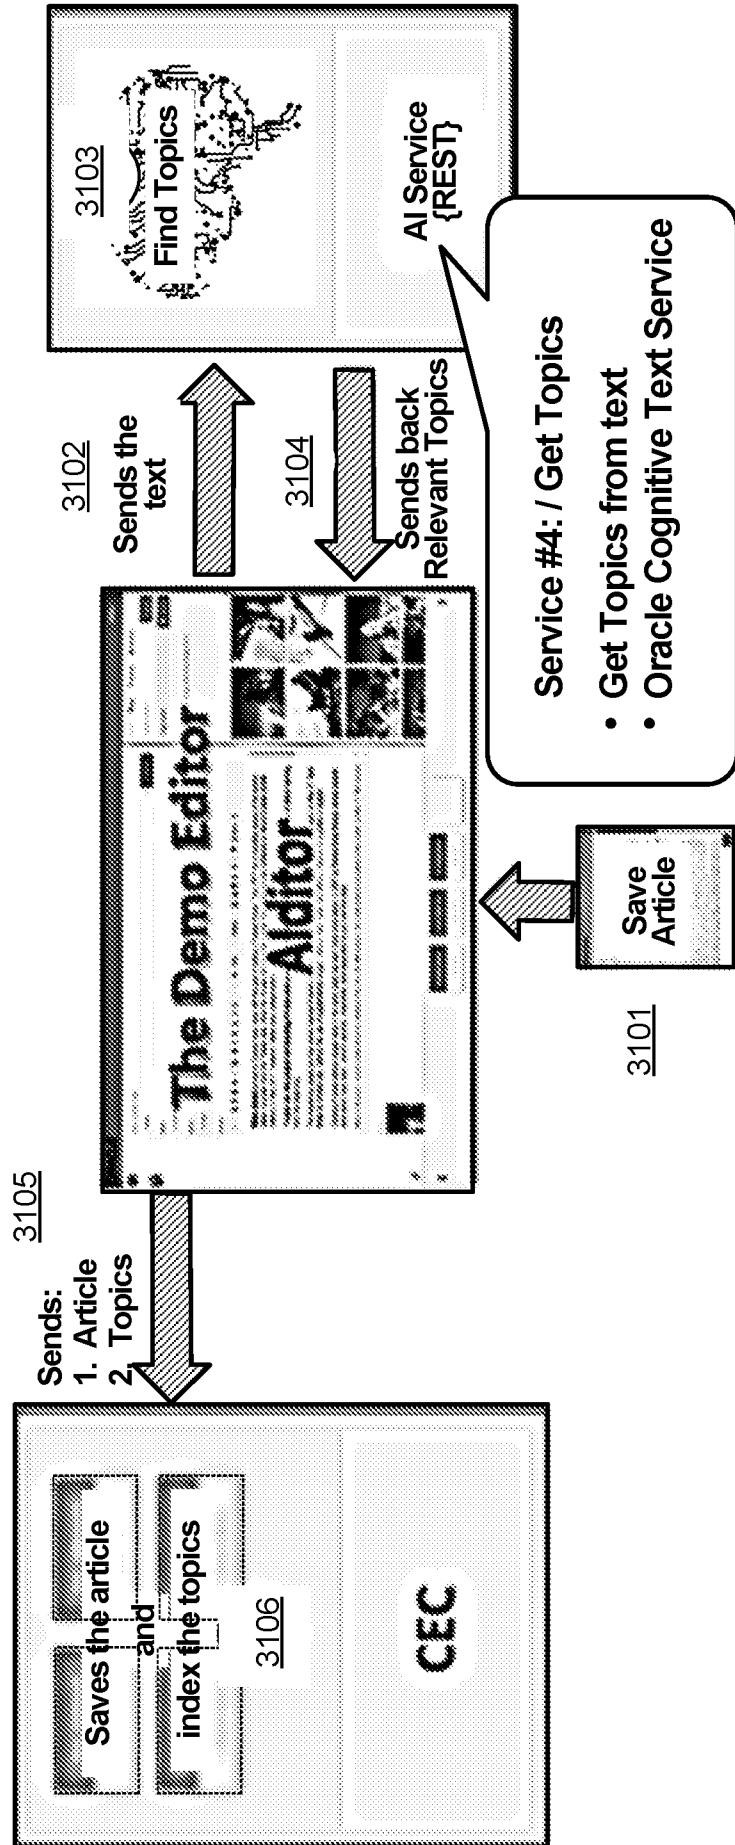


FIG. 32

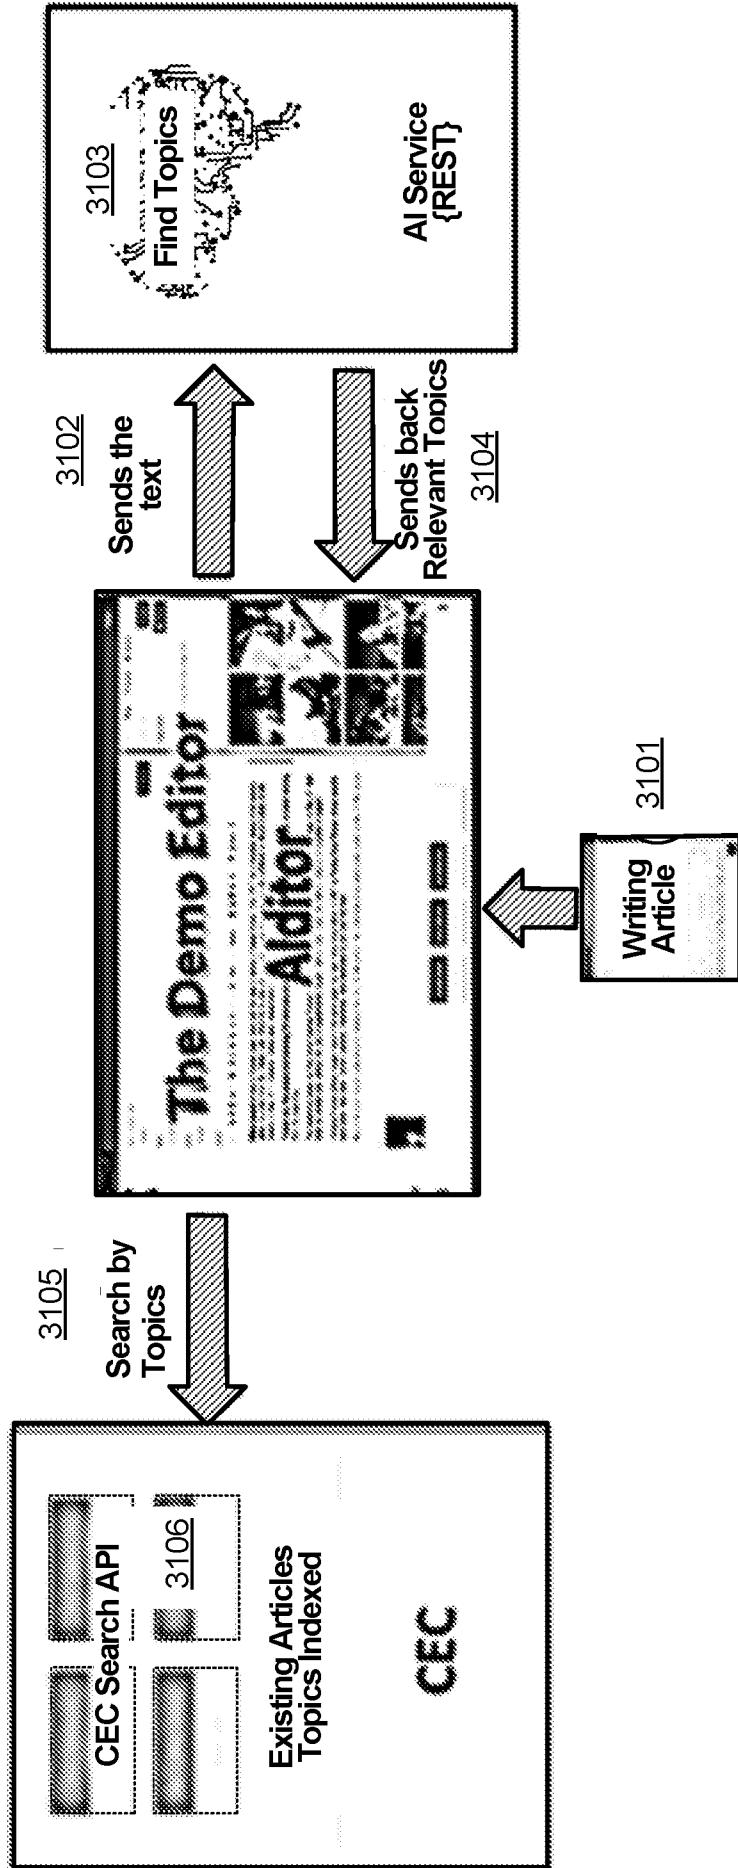


FIG. 33

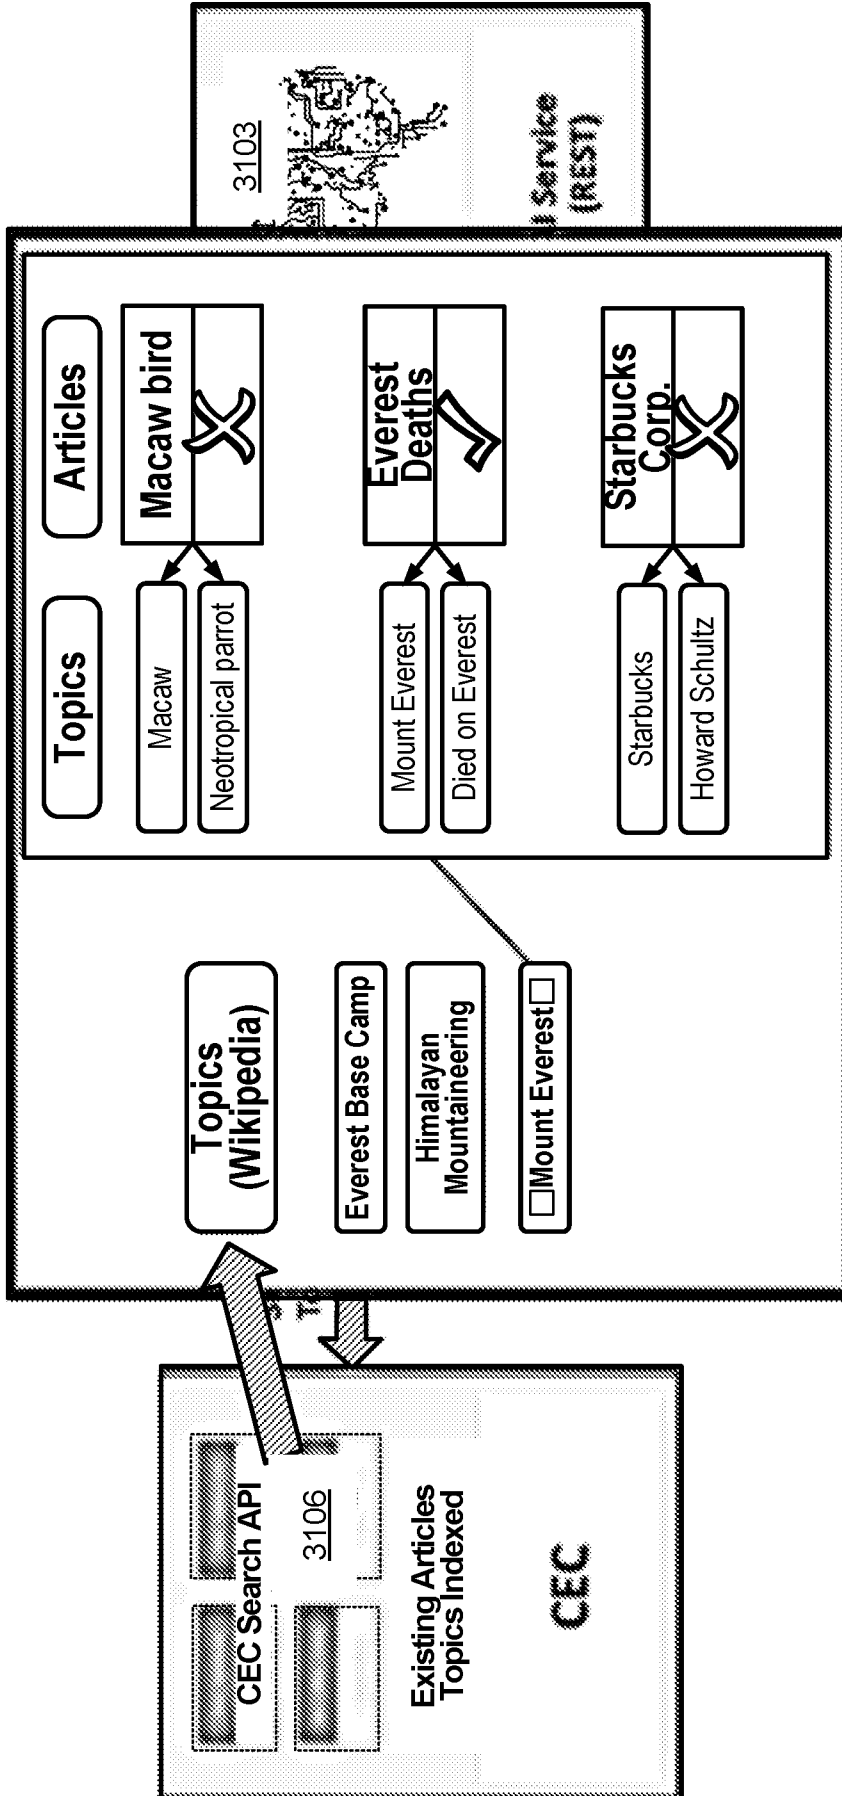


FIG. 34

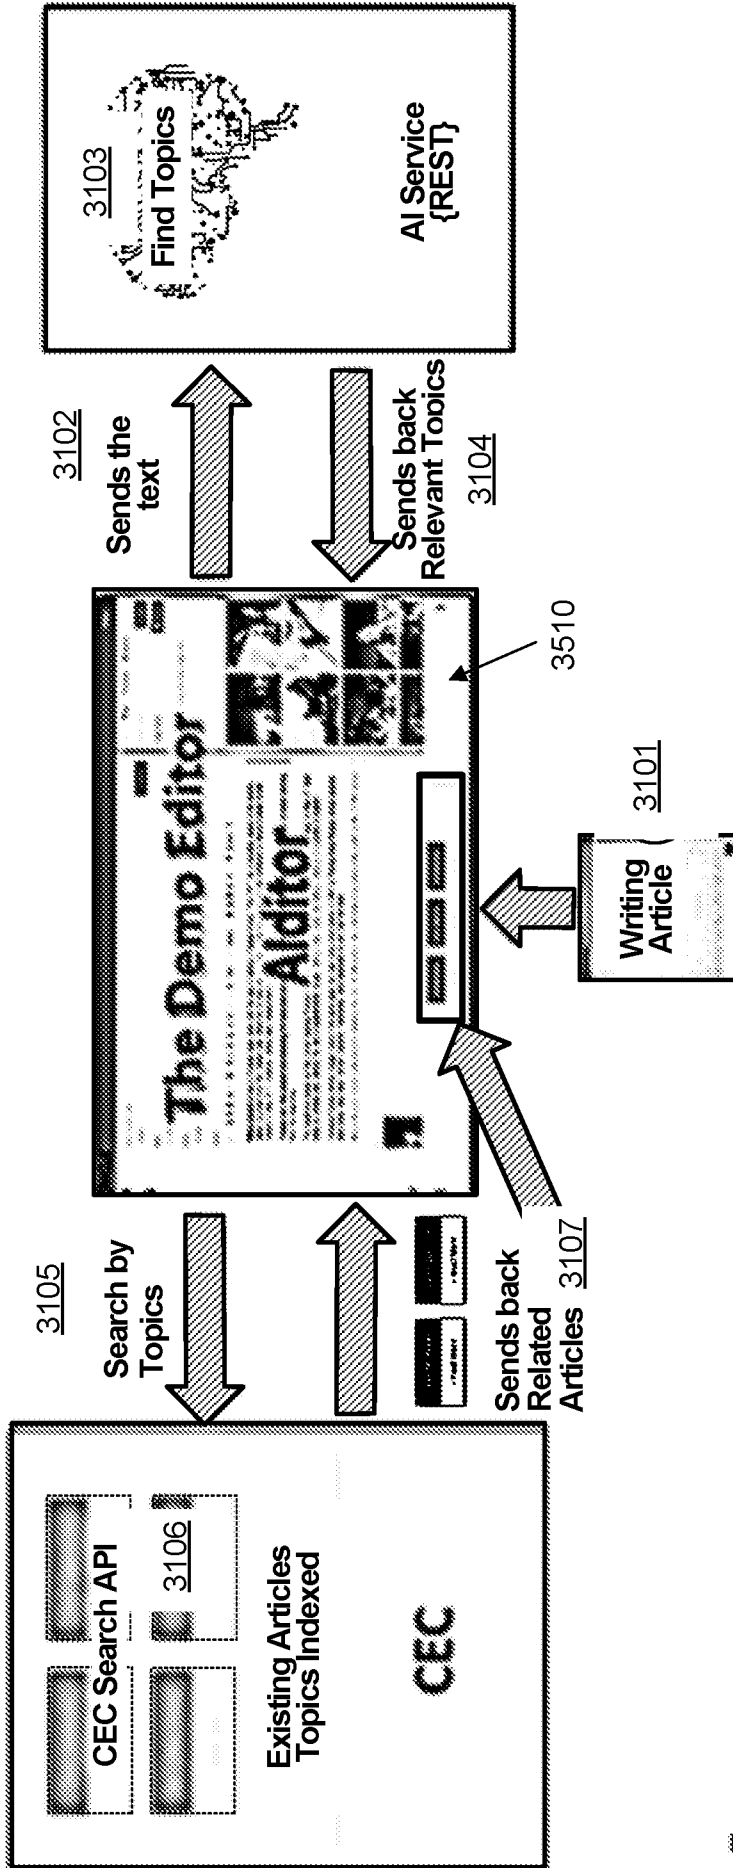


FIG. 35

3600

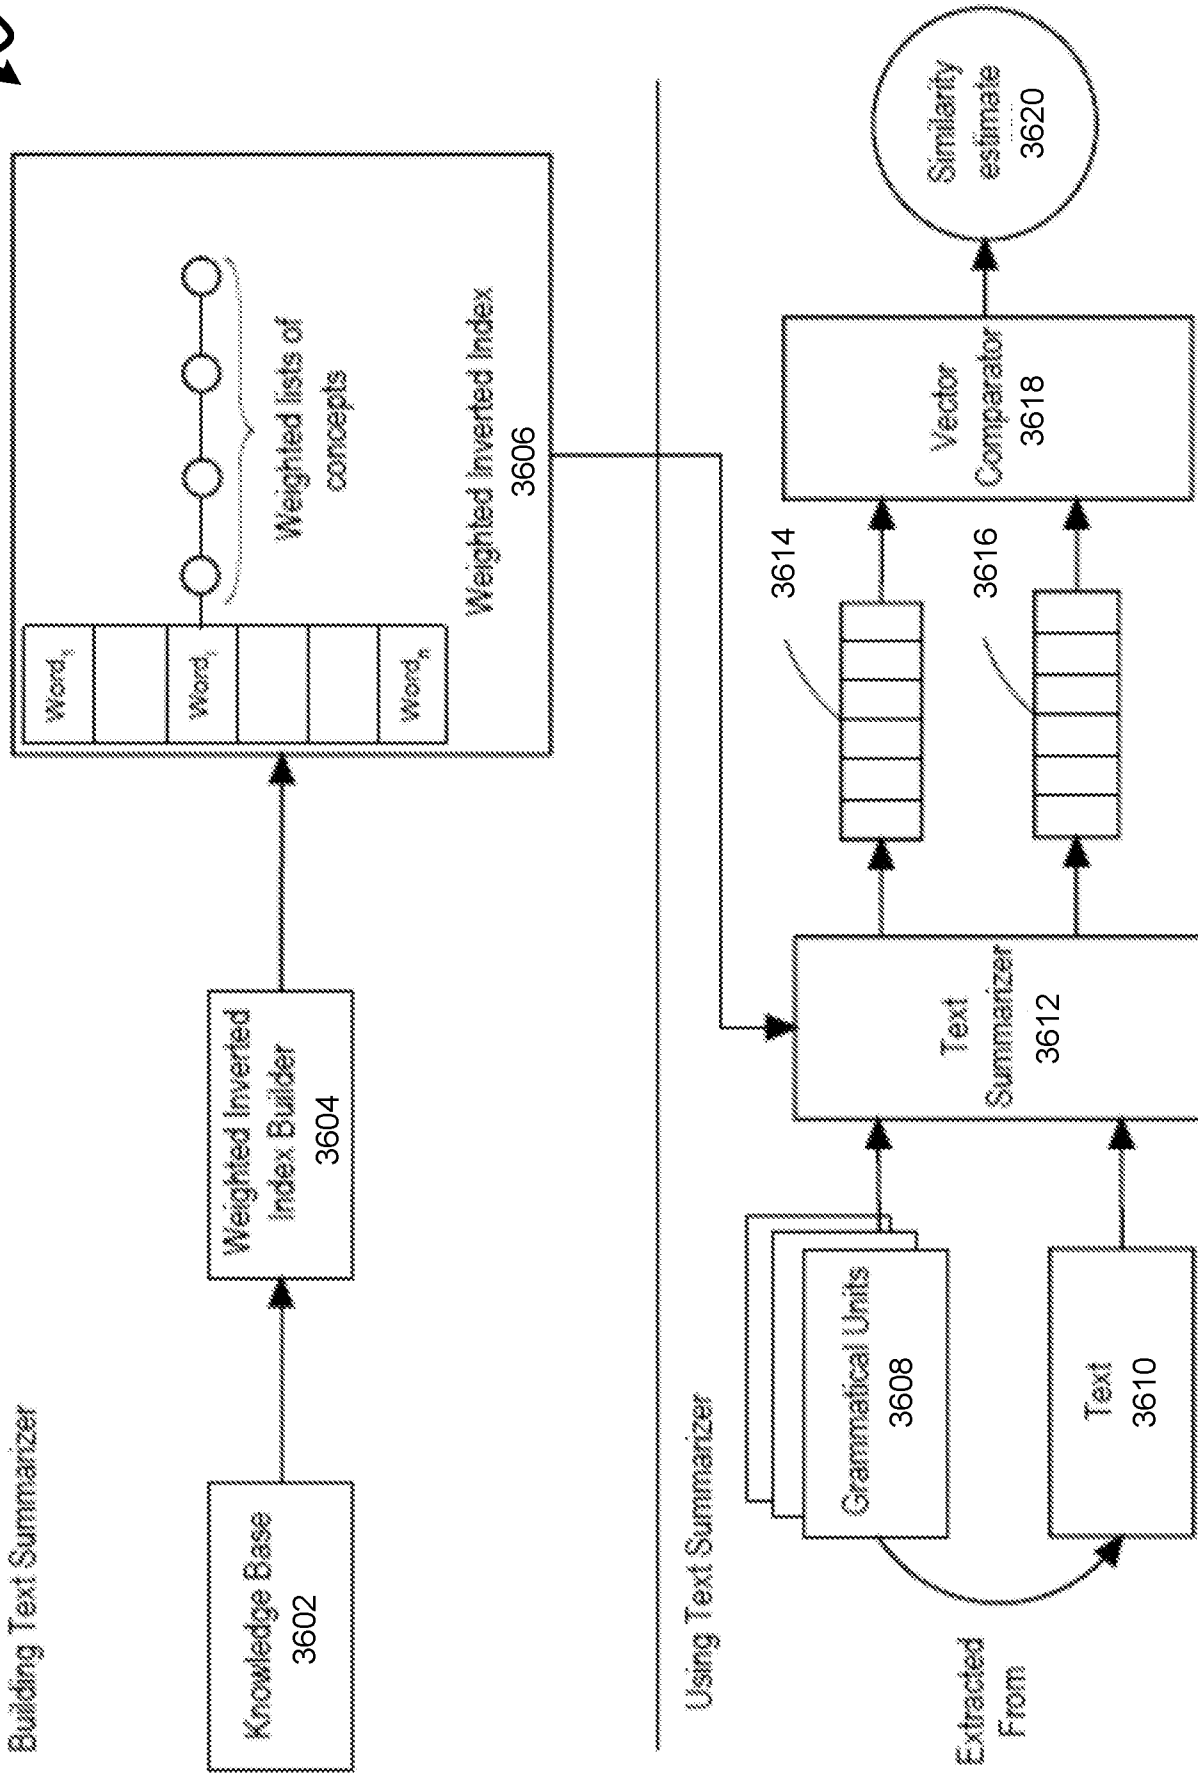


FIG. 36

ORACLE Content and Experience Cloud

Coffee Article 1 **Submit for Review** **View** **Create** **Save** **Close** **Full Screen**

Assets > Coffee Article 1 | English (United States) (en-US)

Content Item Properties

Name

 The item name should not contain the following characters: ',":?<%>|/*

Description (Optional)

Content Item Data Fields

Text *
 26 of 200 characters

Large Text *
 26 of 200 characters

Date *

Number *

Media *

Media Recommendations

Collection: All
 Drag-n-drop items onto sections within form
 Relevance: Size

	Cafe Created 11/29/17 12:00PM ET JPG
	Content item title Created 11/29/17 12:00PM ET Advertisement
	Asset Title Created 11/29/17 12:00PM ET JPG

ADMINISTRATION
 Documents
 Conversations
 Sites
Assets
 Developer
 Analytics
 Settings
 Integrations
 Assets

FIG. 37

ORACLE: Content and Experience Cloud 924 ?

Assets Upload Add Create

Find filename.jpg Name v

Similar to filename.jpg

Café Supreme
 Clear All
 Collection: All
 Channel: All
 Language: All

Digital Assets
 Images
 Videos
 Documents
 HTML
 Other

Content Items
 Article
 Videos

Status
 Published
 Approved
 Rejected
 HCE In Review
 Draft
 Translated

	Image Title Created 01/02/18 by John Smith [CAPSULE TEXT]		Image Title Created 01/02/18 by John Smith [CAPSULE TEXT]		Image Title Created 01/02/18 by John Smith [CAPSULE TEXT]
	Image Title Created 01/02/18 by John Smith [CAPSULE TEXT]		Image Title Created 01/02/18 by John Smith [CAPSULE TEXT]		Image Title Created 01/02/18 by John Smith [CAPSULE TEXT]
	Image Title Created 01/02/18 by John Smith [CAPSULE TEXT]		Image Title Created 01/02/18 by John Smith [CAPSULE TEXT]		Image Title Created 01/02/18 by John Smith [CAPSULE TEXT]
	Image Title Created 01/02/18 by John Smith [CAPSULE TEXT]		Image Title Created 01/02/18 by John Smith [CAPSULE TEXT]		Image Title Created 01/02/18 by John Smith [CAPSULE TEXT]

FIG. 38

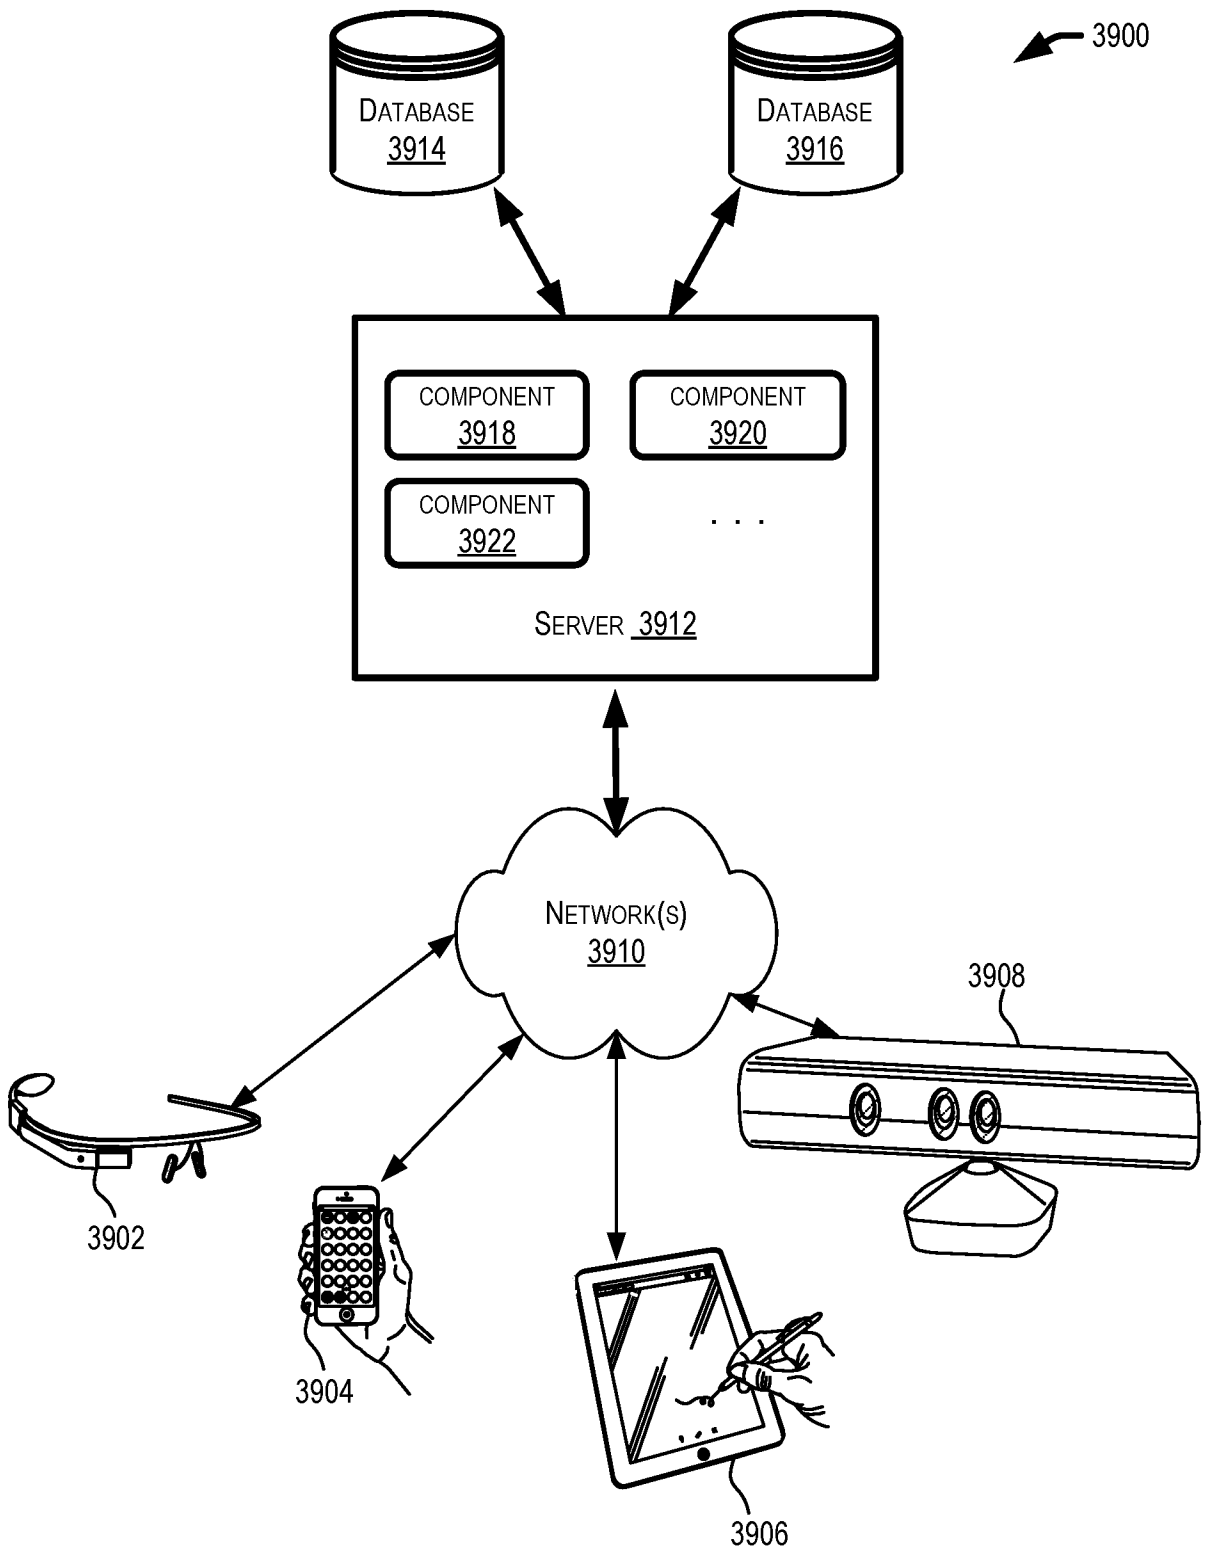


FIG. 39

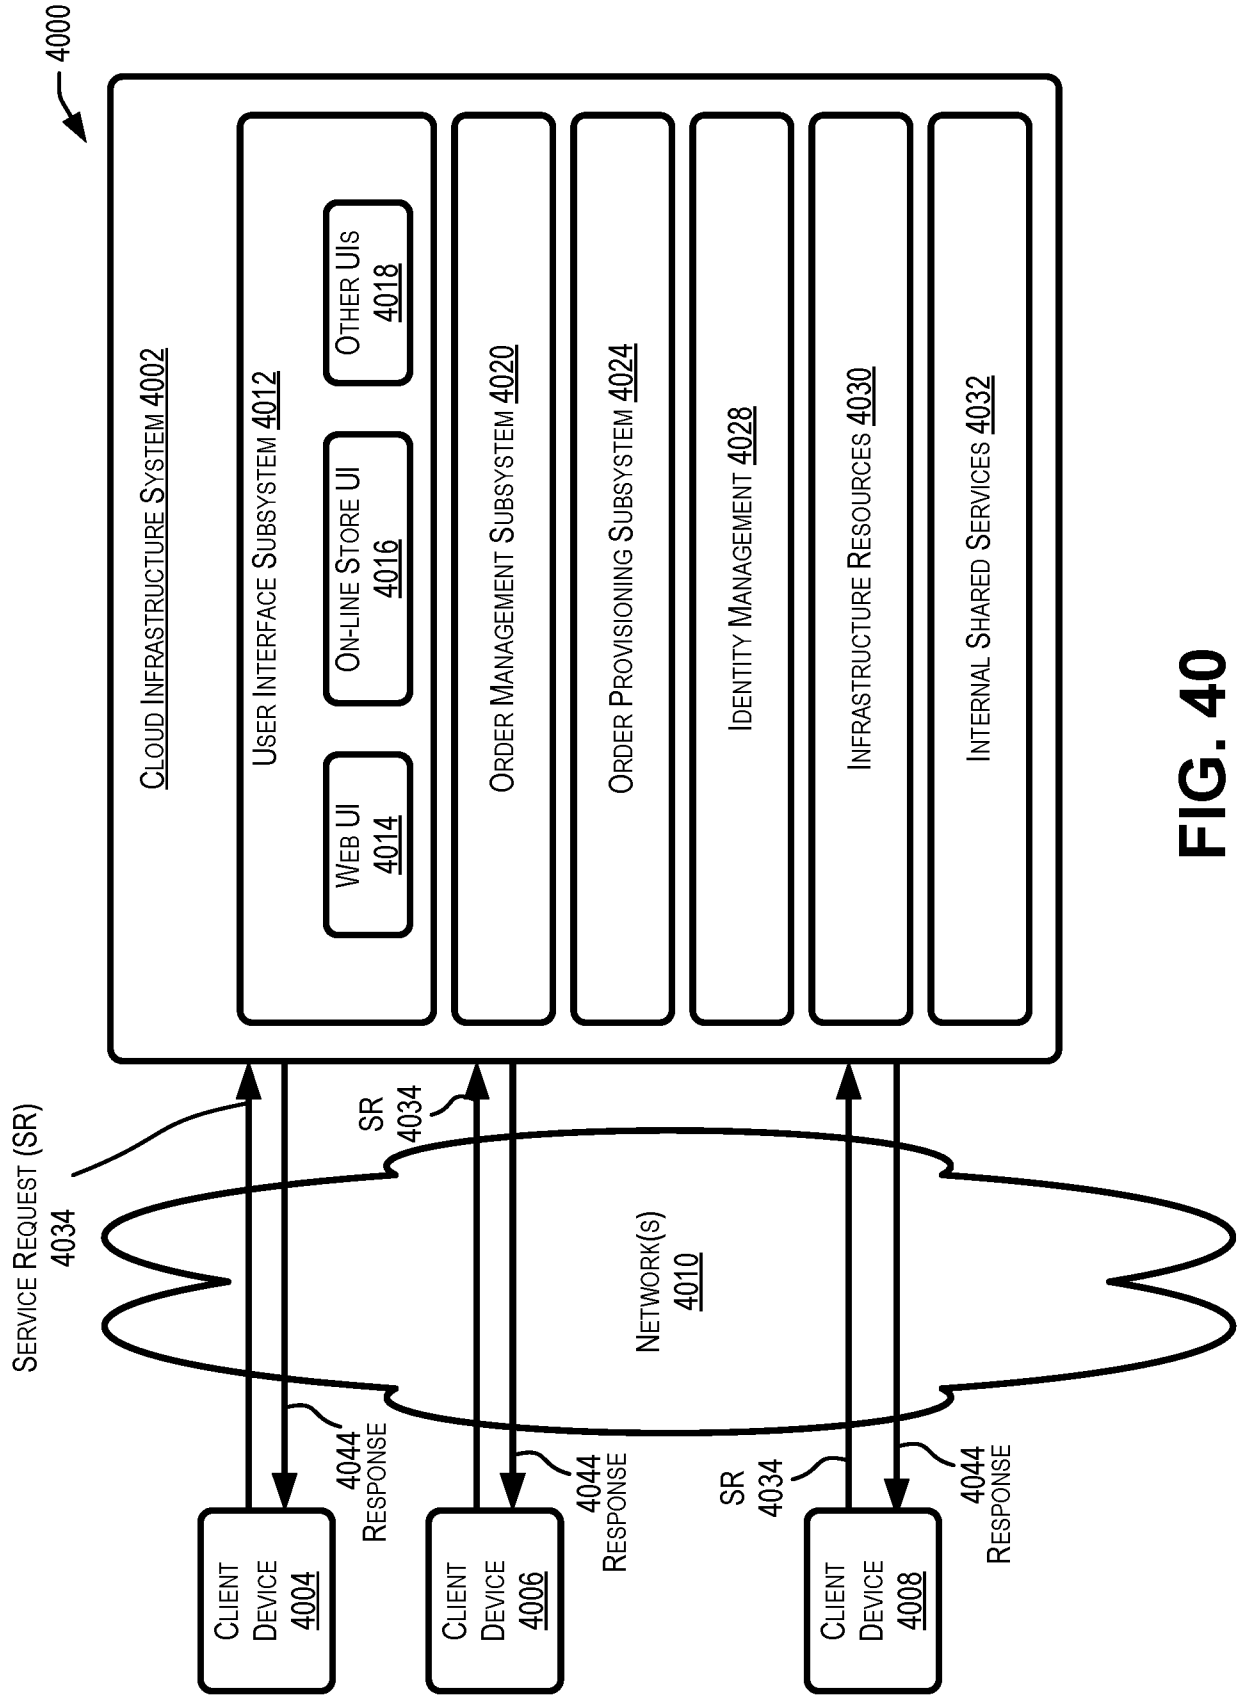


FIG. 40

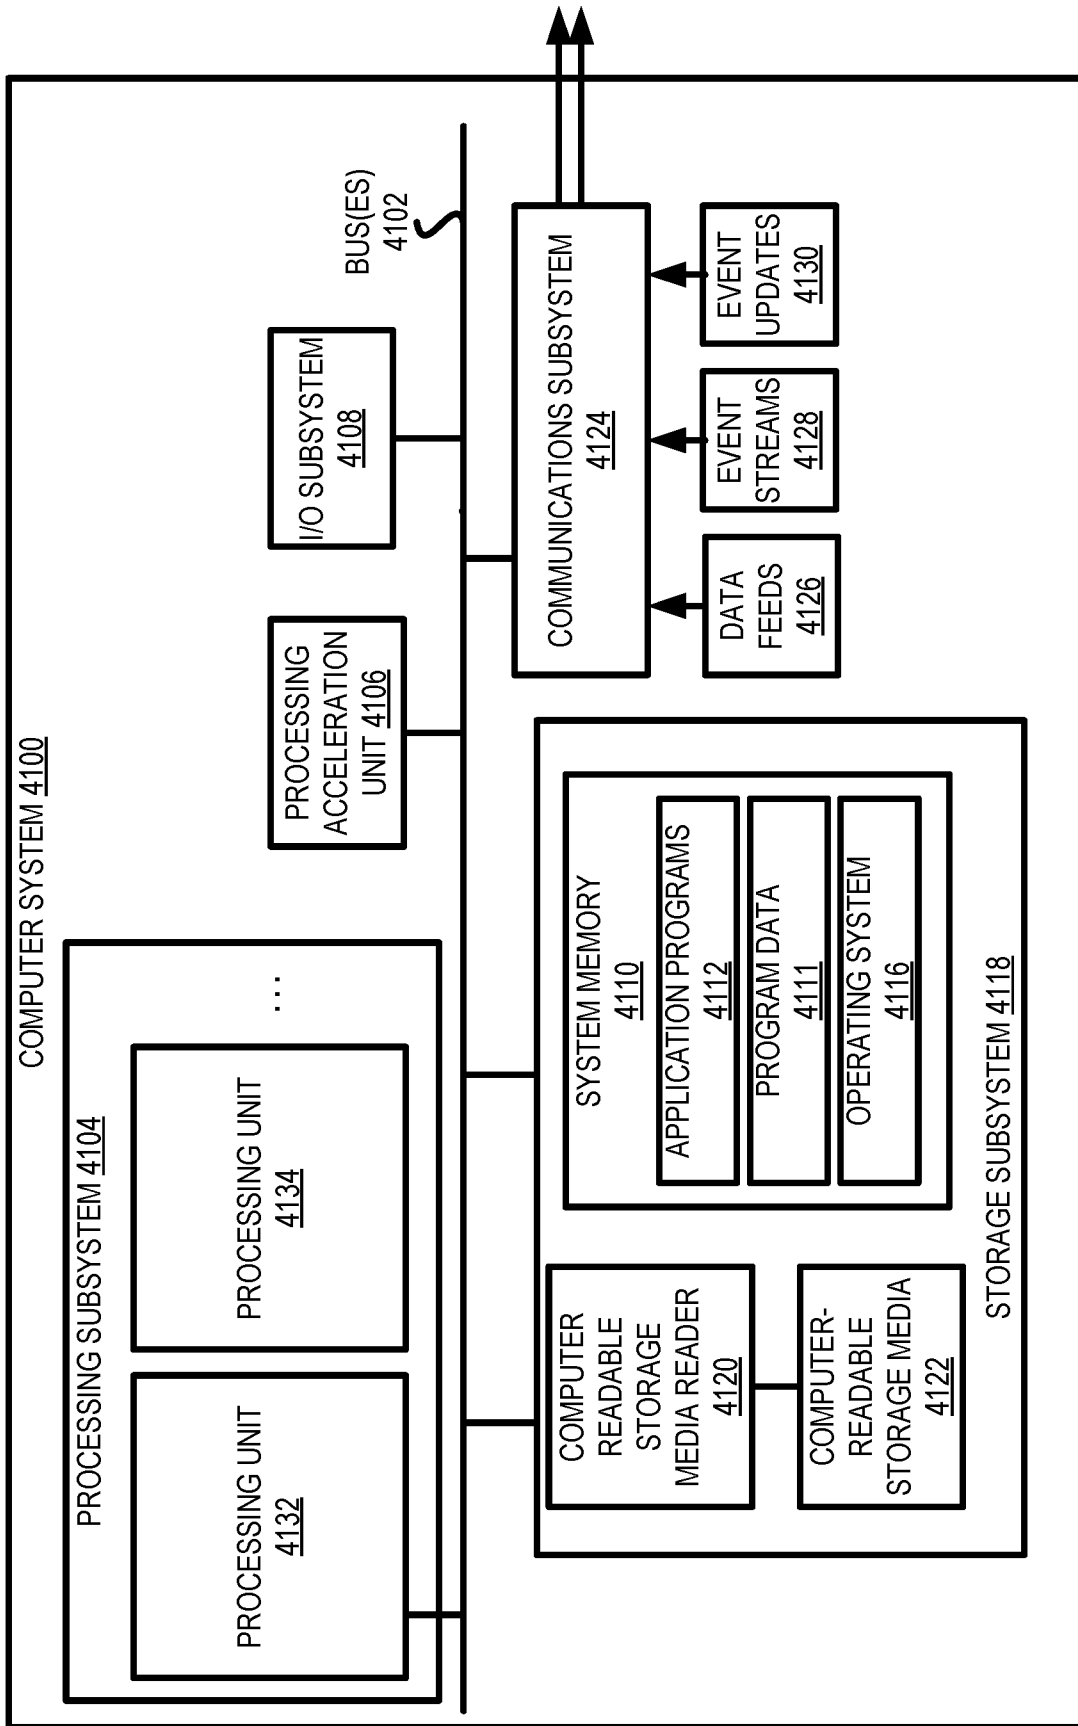


FIG. 41

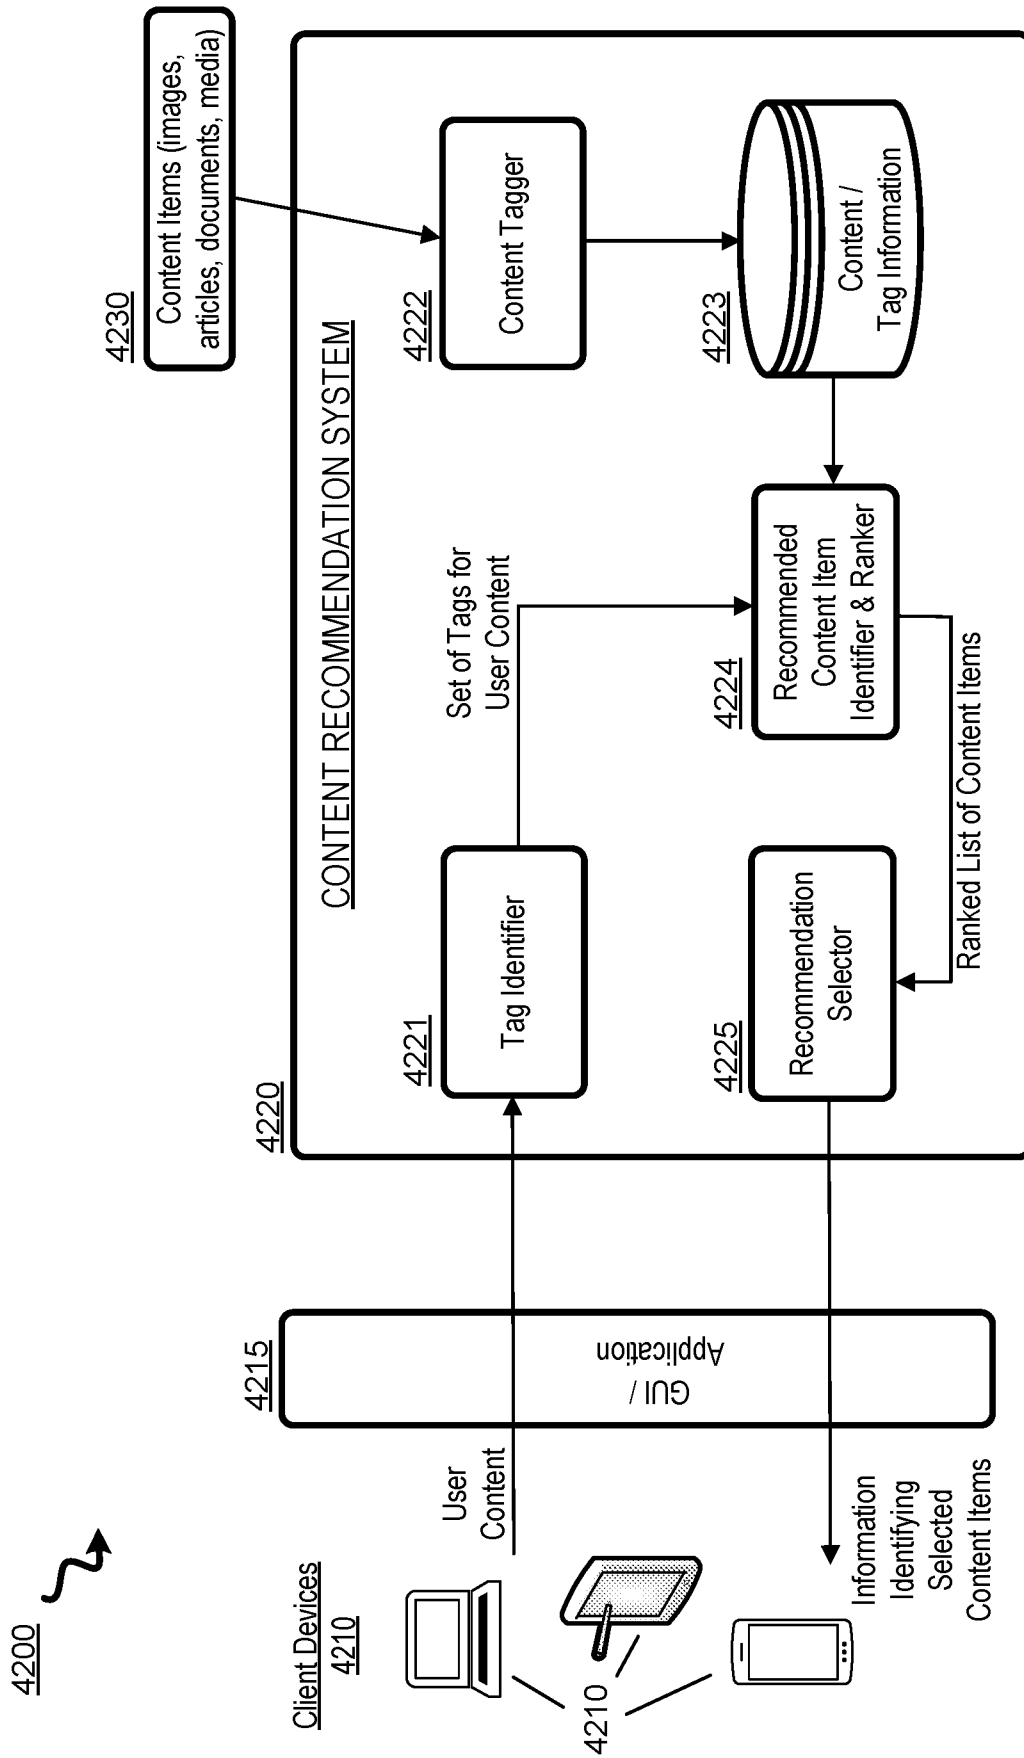


FIG. 42

4300

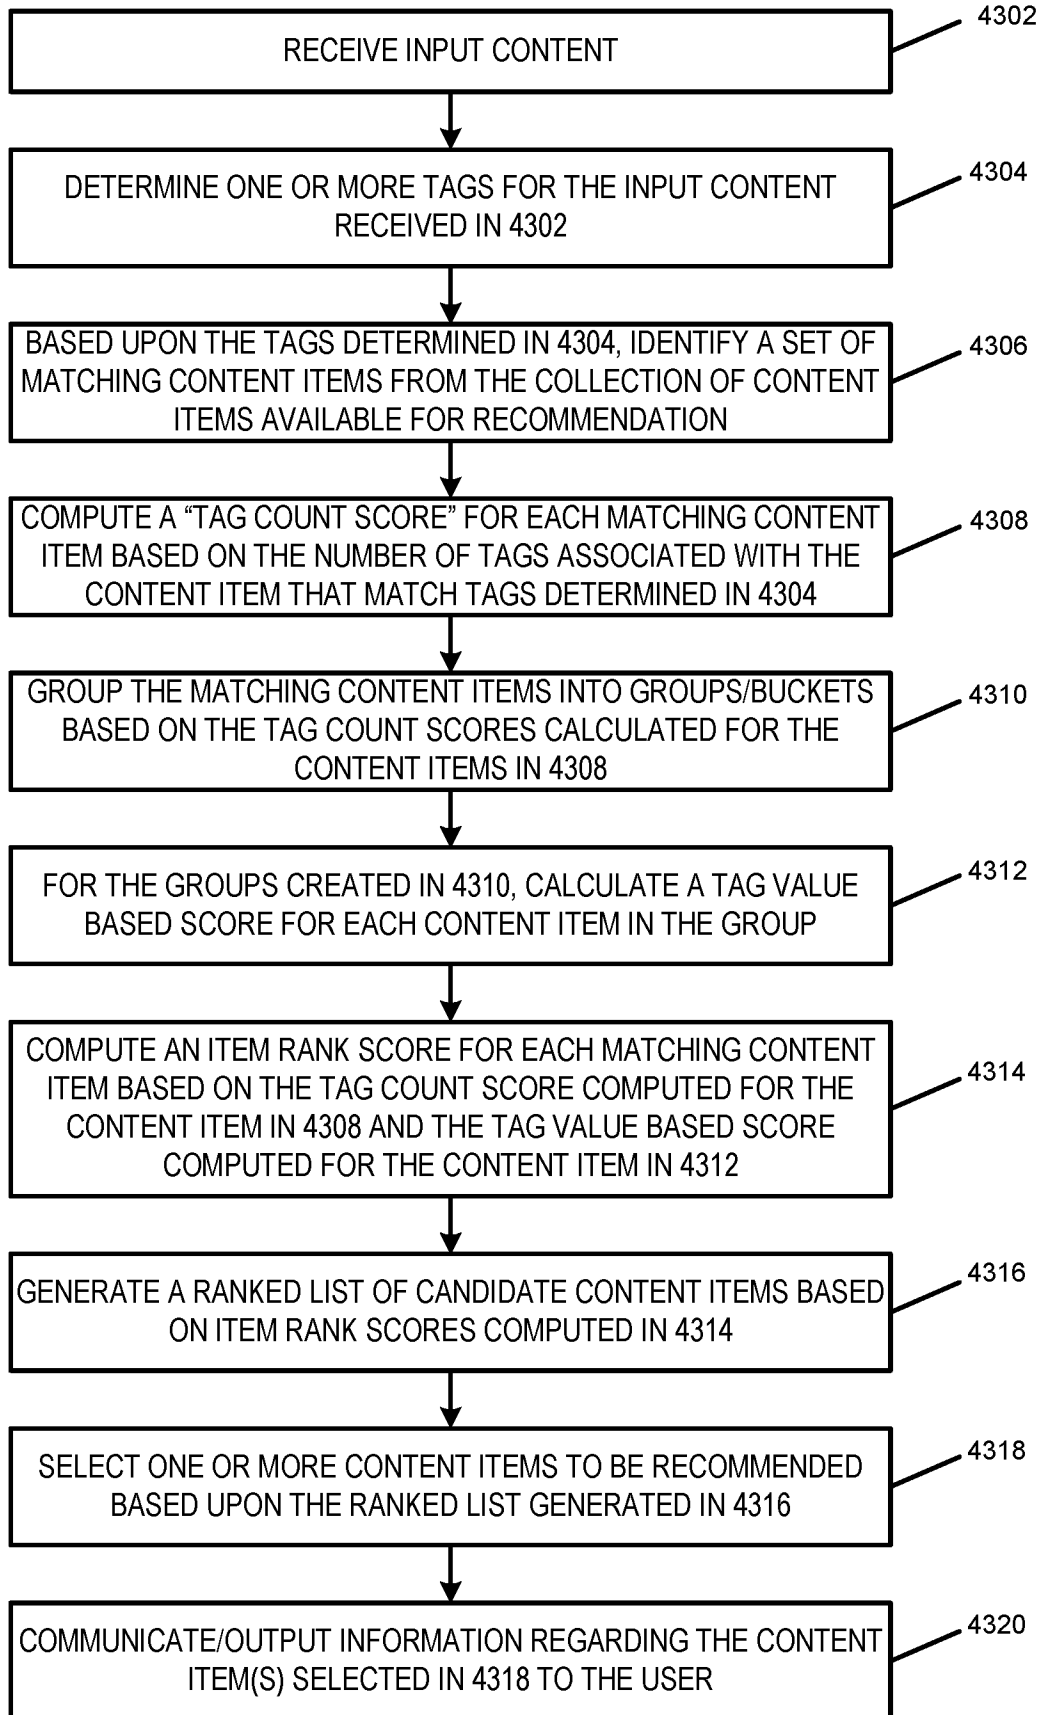


FIG. 43

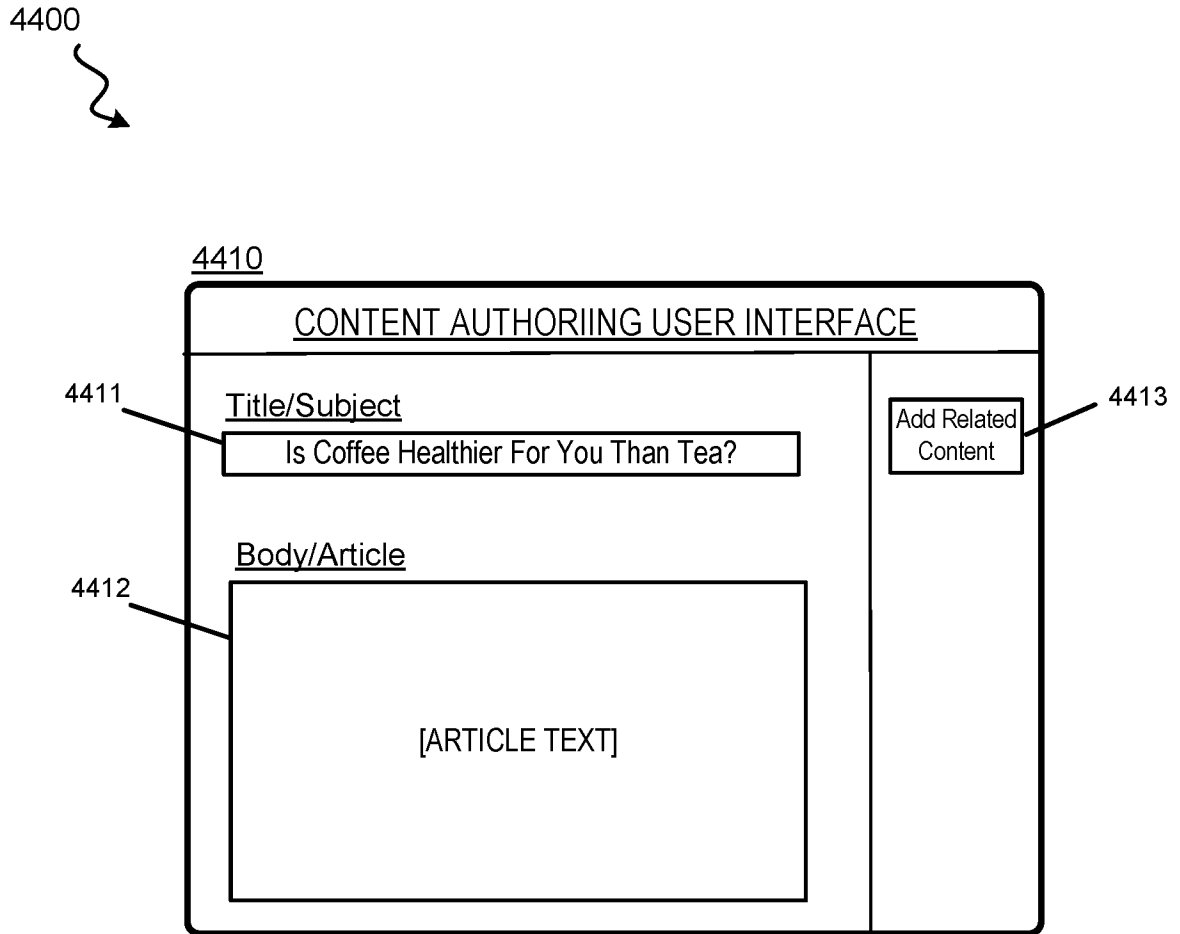


FIG. 44

4500

4501

4502

4503

4504

<u>Image ID</u>	<u>Image Description</u>	<u>Content Tag</u>	<u>Tag Value</u>
Image_1	Man sitting alone on park bench	Human	0.93
		Park	0.07
Image_2	Woman holding cup of coffee	Coffee	0.5
		Human	0.5
Image_3	Tea leaves drying in greenhouse	Tea	0.65
		Greenhouse	0.35
Image_4	Coffee shop employee behind counter smiling and holding cup of coffee.	Coffee	0.35
		Human	0.60
		Cafe	0.05
Image_5	Group of men sitting outside drinking tea	Tea	0.45
		Human	0.55
Image_6	Cup of coffee sitting on desk	Coffee	1.00
Image_7	Man sipping tea while listening to music	Human	0.6
		Tea	0.3
		Music	0.1
Image_8	Piles of coffee beans in a sack	Coffee	0.85
		Farming	0.15

FIG. 45

4600

4601

4602

4603

4604

<u>Image ID</u>	<u>Tag Count Score</u>	<u>Tag Value Score</u>	<u>Ranking Score</u>
Image_1	1	0.93	1.93
Image_2	2	0.25	2.25
Image_3	1	0.65	1.65
Image_4	2	0.21	2.21
Image_5	2	0.248	2.248
Image_6	1	1.0	2.00
Image_7	2	0.18	2.18
Image_8	1	0.85	1.85

FIG. 46

4700
↘

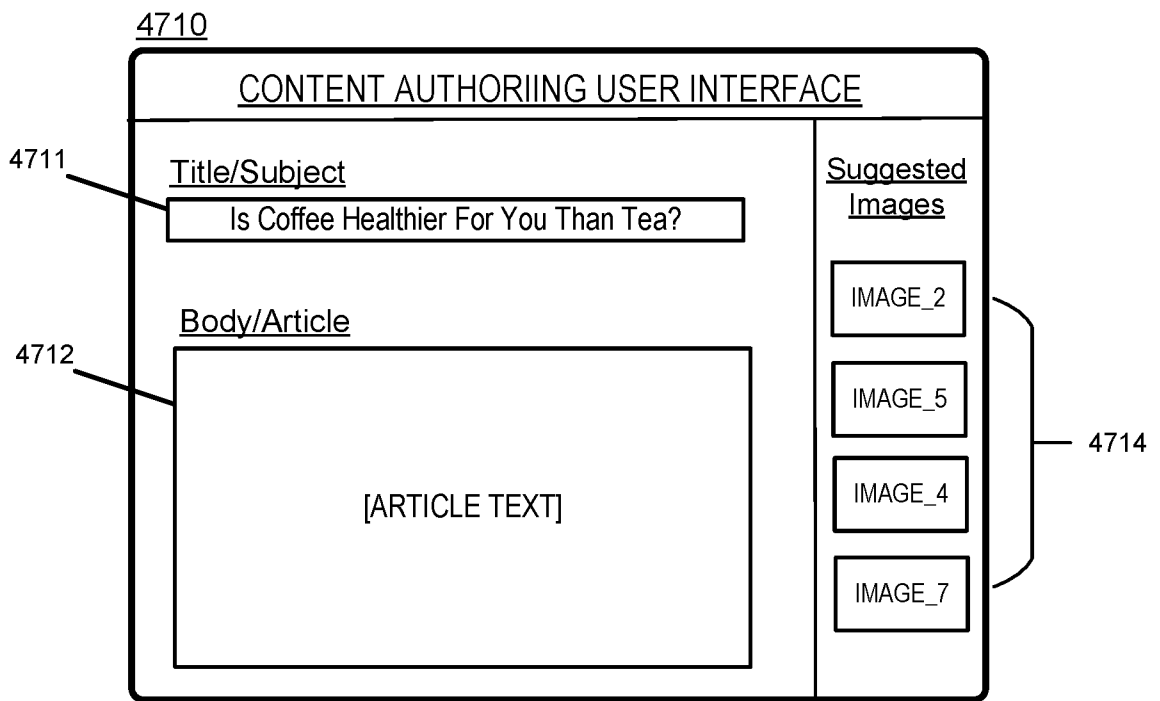


FIG. 47