

(51) International Patent Classification:
G06Q 30/02 (2012.01)(21) International Application Number:
PCT/GR2013/000065(22) International Filing Date:
18 December 2013 (18.12.2013)

(25) Filing Language: English

(26) Publication Language: English

(71) Applicant: **GOOGLE INC.** [US/US]; 1600 Amphitheatre Parkway, Mountain View, California 94043 (US).

(72) Inventors; and

(71) Applicants: **FILIPPOVA, Ekaterina** [RU/CH]; Hallwylstr. 36, 8004 Zurich (CH). **ALFONSECA, Enrique** [ES/CH]; Einsiedlerstrasse 71, 8810 Horgen (CH). **TSOCHANTARIDIS, Ioannis** [GR/CH]; Zurich, ZH (CH). **ALTUN, Yasemin** [TR/CH]; Schulhausstrasse 15, Zurich 8002 (CH). **CIARAMITA, Massimiliano** [IT/CH]; Wannerstrasse 28/66, 8045 Zurich (CH).(74) Agent: **KILIMIRIS, Constantinos**; Patrinos & Kilimiris, 7, Hatziyianni Mexi Str., GR-11528 Athens (GR).

(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, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, 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, 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,

[Continued on next page]

(54) Title: ANNOTATING VIDEOS WITH ENTITIES BASED ON COMMENT SUMMARIES




FIG. 3

(57) Abstract: Implementations of the present disclosure include actions of receiving content data, the content data including a plurality of comments associated with an item of digital content, processing comments of the plurality of comments to provide a set of relevant sentences, receiving a set of entities including one or more entities, each entity in the set of entities being provided based on the set of relevant sentences and being associated with a respective score, selecting at least one entity of the set of entities based on respective scores, and associating the at least one entity with the item of digital content.

WO 2015/092446 A1



TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, **Published:**
KM, ML, MR, NE, SN, TD, TG).

— *with international search report (Art. 21(3))*

ANNOTATING VIDEOS WITH ENTITIES BASED ON COMMENT SUMMARIES

BACKGROUND

5 Users of web services can search for and view digital content. For example, a content-sharing service can enable users to share and/or view digital videos, e.g., over the Internet. Being aware of what items of digital content, such as videos, are about can be used to improve search, recommendations and discovery of videos. In some examples, videos can be about entities, e.g., a person, a place, a thing, a name, and/or other
10 concepts. In some examples, an item of digital content, such as a video, can be associated with one or more entities. For example, information regarding entities can be determined based on as video frames, title, description, and/or audio of the video.

SUMMARY

 This specification relates to associating entities with digital content.
15 Implementations of the present disclosure are generally directed to associating one or more entities with digital content based on a plurality of comments associated with the digital content. Search, recommendations and/or discovery of digital content can be improved based on the entities associated with the digital content.

 In general, innovative aspects of the subject matter described in this specification
20 can be embodied in methods that include actions of receiving content data, the content data including a plurality of comments associated with an item of digital content, processing comments of the plurality of comments to provide a set of relevant sentences, receiving a set of entities including one or more entities, each entity in the set of entities being provided based on the set of relevant sentences and being associated with a
25 respective score, selecting at least one entity of the set of entities based on respective scores, and associating the at least one entity with the item of digital content. Other implementations of this aspect include corresponding systems, apparatus, and computer programs, configured to perform the actions of the methods, encoded on computer storage devices.

30 These and other implementations can each optionally include one or more of the following features: processing comments of the plurality of comments to provide a set of relevant sentences includes: receiving a set of sentences, sentences in the set of sentences including sentence tokens and word tokens, and filtering sentences from the set of

sentences to provide a sub-set of sentences, the set of relevant sentences being provided based on the sub-set of sentences; sentences are filtered from the set of sentences based on at least one of language, length, punctuation, symbols, and letter case; processing comments of the plurality of comments to provide a set of relevant sentences includes
5 providing the sub-set of sentences and digital content data to a summarizer component, the summarizer component processing the sub-set of sentences to provide the set of relevant sentences; the summarizer component processes the sub-set of sentences to provide at least one word set, the at least one word set including words that are common to sentences of the sub-set of sentences, the set of relevant sentences being provided
10 based on the at least one word set; actions further include providing the set of relevant sentences to an entity identifier component, wherein the entity identifier component processes the set of relevant sentences to provide the set of entities; the respective scores each indicate a frequency, at which a respective entity is identified in the set of relevant sentences; each entity of the set of entities includes at least one of a person, a location, a
15 thing, and a concept; and the item of digital content comprises a video.

Particular implementations of the subject matter described in this specification can be implemented so as to realize one or more of the following advantages. In some examples, implementations of the present disclosure enhance search, recommendations and discovery of based on comments associated therewith. In some examples,
20 implementations of the present disclosure deal with noise in the comments, e.g., spam, and extract relevant information from comments. In some examples, implementations of the present disclosure enable extraction of entities that are likely to be relevant, if not most central, to the underlying digital content.

The details of one or more implementations of the subject matter described in this
25 specification are set forth in the accompanying drawings and the description below. Other features, aspects, and advantages of the subject matter will become apparent from the description, the drawings, and the claims.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 depicts an example environment in which users can interact with one or
30 more computer-implemented services.

FIG. 2 depicts an example search content-sharing page.

FIG. 3 depicts an example process that can be executed in accordance with implementations of the present disclosure.

Like reference numbers and designations in the various drawings indicate like elements.

DETAILED DESCRIPTION

5 FIG. 1 depicts an example environment 100 in which users can interact with one or more computer-implemented services, e.g., web services. Example computer-implemented services can include a search service, an electronic mail service, a chat service, a document sharing service, a calendar sharing service, a blogging service, a micro-blogging service, a social networking service, a location (location-aware) service, a
10 check-in service, a ratings and review service, and a content-sharing service, e.g., a photo sharing service, a video sharing service. In the example of FIG. 1, a content-sharing service is depicted, which is described in further detail herein. It is appreciated, however, that implementations of the present disclosure can include one or more computer-implemented services, such as the examples described herein.

15 With continued reference to FIG. 1, a content-sharing system 120 provides content-sharing services. Example content can include digital videos and digital images. The example environment 100 includes a network 102, e.g., a local area network (LAN), a wide area network (WAN), the Internet, or a combination thereof, that connects user devices 106, and the content-sharing system 120. In some examples, the network 102
20 can be accessed over a wired and/or a wireless communications link. For example, mobile computing devices, such as smartphones, can utilize a cellular network to access the network. The environment 100 may include millions of user devices 106.

 In some examples, a user device 106 is an electronic device that is capable of requesting and receiving resources, e.g., web pages, over the network 102. Example user
25 devices 106 include personal computers, mobile computing devices, e.g., smartphones and/or tablet computing devices that can send and receive data over the network 102. As used throughout this document, the term mobile computing device ("mobile device") refers to a user device that is configured to communicate over a mobile communications network. A smartphone, e.g., a phone that is enabled to communicate over the Internet, is
30 an example of a mobile device. A user device 106 can include a user application, e.g., a web browser, to facilitate the sending and receiving of data over the network 102.

 In some examples, to facilitate sharing of content, the content-sharing system 120 can receive content from one or more sources. For example, users can use respective

computing devices 106 to upload content to the content-sharing service 120. A content data repository 122, e.g., database, is provided to store content received by the content-sharing system 120. In some examples, content is stored as one or more computer-readable files, which can include content data and meta-data associated with the content.

5 Example meta-data can include an author of the content, a date that the content was created, geo-location data associated with the content, and/or annotations associated with the content, as described in further detail herein. In some examples, content can be indexed to facilitate searching of the content. In the example of FIG. 1, a content index 124 is provided.

10 The user devices 106 can be used to submit queries 109 to the content-sharing system 120. In some examples, a user device 106 can include one or more input modalities. Example modalities can include a keyboard, a touchscreen and/or a microphone. For example, a user can use a keyboard and/or touchscreen to type in a query. As another example, a user can speak a search query, the user speech being
15 captured through a microphone, and being processed through speech recognition to provide the search query.

In response to receiving a query 109, the content-sharing system 120 accesses the content index 122 to identify content that is relevant to, e.g., have at least a minimum specified relevance score for, the query 109. The content-sharing system 120 identifies
20 the content, generates a results display 111 that includes the content 112, and returns the results display 111 to the user devices 106. In an example context, a results display can include one or more web pages, e.g., one or more results pages. In some examples, a web page can be provided based on a web document that can be written in any appropriate machine-readable language. It is contemplated, however, that implementations of the
25 present disclosure can include other appropriate display types. For example, the content can be provided in a display generated by an application that is executed on a computing device, and/or a display generated by an operating system, e.g., mobile operating system. In some examples, content can be provided based on any appropriate form, e.g., Javascript-html, plaintext.

30 In some examples, the content 112 includes one or more items of content, e.g., videos, images, that were determined to be relevant to the query 109. The results display 111 can also include meta-data associated with the content 112. In the example of FIG. 1, meta-data can include one or more comments 113 that are associated with the content 112. For example, the content-sharing system 120 can facilitate user comments to shared

content. In this manner, users that are able to view the content can submit comments, which comments are associated with the content, e.g., stored as meta-data in a computer-readable file containing the content data.

Implementations of the present disclosure are directed to annotating content with entities that are determined based on comments associated with the content.

Implementations will be described in further detail herein with the content including digital videos. In accordance with implementations of the present disclosure, video data is received and includes comments. In some implementations, comments are processed and one or more entities relevant to the video are determined based on the comments.

Example entities can include a person, a place, a thing, a location, and a name. In general, an entity can include any defined concept.

Referring again to FIG. 1, the environment 100 includes a comment processing system 130. In some examples, the comment processing system 130 processes a set of comments associated with each item of content, e.g., video, to provide one or more entities that can be associated with a respective item of content. In some examples, the one or more entities can be provided as meta-data associated with the item of content, e.g., stored in the computer-readable file of the content.

In some implementations, an initial set of comments are provided from the content data. In some examples, the initial set of comments includes all comments associated with the content. In some examples, if there are an insufficient number of comments (C), the content is not processed, and no entities are associated with the content. For example, the number of comments (C) can be compared to a threshold number of comments (C_{THRLO}), and, if C is less than C_{THRLO} , the content is not processed and no entities are associated with the content. In some examples, if there are too many comments, comments can be filtered from the set of comments. For example, the number of comments (C) can be compared to a threshold number of comments (C_{THRHI}), and, if C exceeds C_{THRHI} , comments are filtered from the set of comments. In some examples, comments can be filtered based on time. For example, the oldest comments can be removed from the set of comments. As another example, the most recent comments can be removed from the set of comments. In some examples, comments can be randomly filtered from the set of comments.

In some implementations, comments in the set of comments are processed to provide a set of sentences. In some examples, the comments of the set of comments are provided to a linguistics component, which processes the comments to provide the set of

sentences. In some examples, the linguistics component can be provided as one or more computer-executable programs executed by one or more computing devices. In some examples, processing of the comments can include annotating the comments with sentence boundaries and/or token boundaries. In some examples, sentence boundaries can include an annotation indicating the beginning of a sentence, and an annotation indicating the end of the sentence. In some examples, token boundaries indicate individual words, symbols and/or punctuations in the comments.

In some implementations, one or more sentences are filtered from the set of sentences. In some examples, sentences can be filtered based on language, length, punctuation, symbols, and/or letter case. For example, sentences that are determined not to be in a specified language, e.g., English, can be removed. As another example, sentences that are determined to be too short, e.g., the number of words in the sentence is less than a minimum number of words, are removed. As another example, sentences that are determined to be too long, e.g., the number of words in the sentence exceeds a maximum number of words, are removed. As another example, sentences having too many uppercase letters relative to the letters of the sentence, e.g., a ratio of uppercase letters to all letters in the sentence exceeds a threshold ratio, can be removed. As another example, sentences having too many punctuations relative to the overall number of characters, e.g., a ratio of punctuations to characters of a sentence exceeds a threshold ratio, can be removed. As another example, sentences having a number of symbols that exceeds a threshold number of symbols can be removed.

In some implementations, if there are too many sentences in the set of sentences, e.g., after the set of sentences has been filtered, sentences can be further filtered from the set of sentences, e.g., to provide a sub-set of sentences. For example, the number of sentences (S) can be compared to a threshold number of sentences (S_{THR}), and, if S exceeds S_{THR} , sentences are filtered from the set of sentences. In some examples, sentences can be filtered based on time. For example, sentences provided from the oldest comments can be removed from the set of sentences. As another example, sentences provided from the most recent comments can be removed from the set of sentences. In some examples, sentences can be randomly filtered from the set of sentences. In some examples, above-described thresholds, e.g., threshold ratios, threshold number of symbols, minimum/maximum number of words, and the like, can be tightened to further filter sentences based on the respective parameters, e.g., language, length, punctuation,

symbols, and/or letter case. In some examples, if no sentences are filtered from the set of sentences, the sub-set of sentences includes the sentences of the set of sentences.

In some implementations, the sub-set of sentences is provided to a summarizer component, which provides a set of relevant sentences. In some examples, the summarizer component can be provided as one or more computer-executable programs executed by one or more computing devices. In some examples, additional content data for the respective item of content, e.g., video, can also be provided to the summarizer component. Example additional content data can include a title associated with the content, and a description associated with the content. For example, a user that provided the item of content to the content-sharing service 120 can provide a title for the content and/or a description of the content.

In some implementations, the summarizer component processes the sentences to provide the set of relevant sentences. In some examples, the summarizer processes the sentences to provide one or more word sets. An example first word set can include words that are common for the particular language, e.g., English. In some example, common words can include words, e.g., the, a, an, and, but, or, that do not reflect a topic of a particular sentence, e.g., words that are frequently seen in sentences regardless of topic(s). An example second word set can include words that are common to the particular sentences of the sub-set of sentences. In some examples, a frequency of each word across the sentences of the sub-set of sentences can be determined. In some examples, if the frequency of a word exceeds a threshold frequency, and the word is not included in the first word set, e.g., words that are common to the language, then the word is included in the second word set. In some examples, if the frequency of a word does not exceed a threshold frequency, and/or the word is included in the first word set, e.g., words that are common to the language, then the word is not included in the second word set. In some examples, sentences that include one or more words of the second word set are included in the set of relevant sentences.

In some examples, relevant sentences can be identified based on one or more words that are included in the title for the content and/or the description of the content. In some examples, if a word in a sentence corresponds to a word in the title and/or description, and the word is not included in the first word set, the sentence can be included in the set of relevant sentences.

In some implementations, relevant sentences of the set of relevant sentences are processed to identify one or more entities. In some examples, the set of relevant sentences

is provided to an entity identifier component, which processes the relevant sentences to identify one or more entities that can be associated with the underlying item of digital content, e.g., the digital content, from which the set of relevant sentences was provided.

In some examples, the entity identifier component can be provided as one or more

5 computer-executable programs executed by one or more computing devices. In some examples, one or more entities can be identified based on words provided in the relevant sentences. In some examples, one or more entities can be identified as related entities based on entities identified from the words. In some examples, each entity has an associated topicality score that reflects the frequency, at which the entity is identified
10 from the relevant sentences. In some examples, an entity that is frequently identified from the relevant sentences has a higher topicality score than an entity that is less frequently identified from the relevant sentences. In some examples, the entity identifier component provides a set of entities, the set of entities including one or more entities, each entity in the set of entities being associated with a respective topicality score.

15 In some implementations, one or more entities are selected from the set of entities to be associated with the underlying digital content. In some examples, entities can be selected based on respective topicality scores. For example, entities can be provided in ranked order, where the top X entities are provided as primary entities, the next top Y entities are provided as secondary entities, and the next Z entities are provided as tertiary
20 entities. In some examples, one or more primary entities, e.g., $X \geq 1$, one or more secondary entities, e.g., $Y \geq 1$, and one or more tertiary entities, e.g., $Z \geq 1$. In some examples, a primary entity indicates a topic that the underlying digital content is determined to be primarily about. In some examples, a secondary entity indicates a topic that the underlying digital content is determined to be secondarily about. In some
25 examples, a tertiary entity indicates a topic that the underlying digital content is determined to be relevant and/or related to.

In some implementations, each entity of one or more entities that have been selected is associated with the underlying item of digital content. In some examples, metadata indicating each of the one or more entities can be included in computer-readable
30 file of the content.

In some examples, one or more recommendations can be provided to users based on entities associated with an item of digital content. In some examples, recommendations can be provided based on an intersection of entities associated with a user, e.g., provided from a user profile, and entities associated with digital content. For

example, and referring again to FIG. 1, the environment 100 includes a recommendation system 140. In some examples, the recommendation system 140 can access an index of users, which index provides one or more entities associated with respective users. In some examples, the recommendation system 140 can access an index of digital content, which index provides one or more entities associated with respective items of digital content. In some examples, digital content, for which entities match with entities associated with a particular user, can be recommended to the user. In some examples, the recommendation system 140 can match digital content to other digital content based on entities that the digital content have in common, and can be recommended relative to one another, as described in further detail below.

Although, in the example of FIG. 1, the content-sharing system 120, the comment processing system 130, and the recommendation system 140 are depicted as communicating directly with one another, it is contemplated that the content-sharing system 120, the comment processing system 130, and the recommendation system 140 can communicate with one another through the network 102.

FIG. 2 depicts an example content-sharing page 200. In the example of FIG. 2, the content-sharing page 200 includes a content display region 202, and a recommended content region 204. In the depicted example, a video 206 is displayed. The video 206 includes a title 208, a description 210, and comments 212, 214, 216, 218, 220 associated therewith. In some examples, the video 206 can be displayed based on a query submitted by a user. In the depicted example, videos 230, 232 are displayed in the recommended content region 204.

In accordance with implementations of the present disclosure, comments associated with the video 206, e.g., including comments 212, 214, 216, 218, 220, can be processed to associate one or more entities with the video 206. For the example of FIG. 2, example entities can include cats, cars and motorcycles. In some examples, the comments 216, 220 can be filtered from determining entities to be associated with the video 206. For example, the comment 216 can be determined to be spam, and consequently, can be filtered from consideration. As another example, the comment 220 can be considered to be too short, contain too many uppercase letters, and/or contain too many punctuations, and consequently, can be filtered from consideration. In some examples, the comments 212, 214, 218 can be considered in determining entities to be associated with the video 206. In some examples, each of the comments 212, 214, 218 can include one or more sentences having appropriate parameters, e.g., language, length, punctuation, symbols,

and/or letter case, and that are determined to be relevant sentences, e.g., based on the terms “cats,” “cars,” “motorcycle,” and “motorcycles,” and the title 208 and description 210 of the video 206. In some examples, the comments 212, 214, 218 can be processed to associate the example entities “cat,” “car,” “motorcycle,” and “vehicle” with the video 206.

In some implementations, the videos 230, 232 can be recommended at least partially based on the entities associated with the video 206. For example, it can be determined that the video 230 is associated with the entity “cat.” Consequently, the video 230 can be identified as a recommended video, e.g., the entity “cat” is associated with both the video 206 and the video 230, and can be displayed in the recommended content region 204. As another example, it can be determined that the video 232 is associated with the entity “motorcycle.” Consequently, the video 232 can be identified as a recommended video; e.g., the entity “motorcycle” is associated with both the video 206 and the video 232, and can be displayed in the recommended content region 204.

FIG. 3 depicts an example process 300 that can be executed in accordance with implementations of the present disclosure. The example process 300 can be implemented, for example, by the example environment 100 of FIG. 1. In some examples, the example process 300 can be provided by one or more computer-executable programs executed using one or more computing devices. In some examples, the process 300 can be performed for each item of content, e.g., video, provided, e.g., by the content processing system 130 of FIG. 1.

Video data is received (302). For example, the video data can be received by the comment processing system 130. In some examples, the video data includes comments, e.g., C comments. In some examples, the video data includes a title and/or description associated with the underlying video. It is determined whether C exceeds C_{THRLO} (304). For example, the content processing system can determine whether C exceeds C_{THRLO} . If C does not exceed C_{THRLO} , the comments are not processed (306). If C exceeds C_{THRLO} , it is determined whether C exceeds C_{THRHI} (308). If C exceeds C_{THRHI} , the comments can be filtered (310). In some examples, comments can be filtered as described herein. For example, comments can be removed, such that less than or equal to C_{THRHI} comments are included in the set of comments that are to be processed.

Comments are processed to provide sentences (312). For example, the comments can be provided to a linguistics component, which processes the comments to provide the set of sentences, as described herein. Sentences are filtered (314). For example, sentences

can be filtered from the set of sentences to provide a sub-set of sentences, as described herein. The number of sentences is capped to provide a sub-set of sentences (316). The sub-set of sentences and video data, e.g., title, description, are provided to a summarizer component (318). In some examples, and as described herein, the summarizer component
5 processes the sentences to provide a set of relevant sentences. The set of relevant sentences is provided to an identifier component (322). In some examples, and as described herein, the set of relevant sentences is provided to an entity identifier component, which processes the relevant sentences to identify one or more entities that can be associated with the video. For example, the entity identifier component can
10 provide a list of entities, each entity being associated with a respective topicality score.

Entities are selected to be associated with the video (326). In some examples, and as described herein, entities can be selected from the list of entities based on the respective topicality scores. For example, primary, secondary and/or tertiary entities can be provided. A video data set is provided (328). For example, the video data set can
15 include the one or more entities, e.g., primary, secondary and/or tertiary entities that are associated with the underlying video. In some examples, the content-sharing system 120 and/or the comment processing system 130 can provide the video data set.

Implementations of the subject matter and the operations described in this specification can be realized in digital electronic circuitry, or in computer software,
20 firmware, or hardware, including the structures disclosed in this specification and their structural equivalents, or in combinations of one or more of them. Implementations of the subject matter described in this specification can be realized using one or more computer programs, i.e., one or more modules of computer program instructions, encoded on computer storage medium for execution by, or to control the operation of, data processing
25 apparatus. Alternatively or in addition, the program instructions can be encoded on an artificially-generated propagated signal, e.g., a machine-generated electrical, optical, or electromagnetic signal that is generated to encode information for transmission to suitable receiver apparatus for execution by a data processing apparatus. A computer storage medium can be, or be included in, a computer-readable storage device, a computer-
30 readable storage substrate, a random or serial access memory array or device, or a combination of one or more of them. Moreover, while a computer storage medium is not a propagated signal, a computer storage medium can be a source or destination of computer program instructions encoded in an artificially-generated propagated signal.

The computer storage medium can also be, or be included in, one or more separate physical components or media (e.g., multiple CDs, disks, or other storage devices).

The operations described in this specification can be implemented as operations performed by a data processing apparatus on data stored on one or more computer-readable storage devices or received from other sources.

The term “data processing apparatus” encompasses all kinds of apparatus, devices, and machines for processing data, including by way of example a programmable processor, a computer, a system on a chip, or multiple ones, or combinations, of the foregoing. The apparatus can include special purpose logic circuitry, e.g., an FPGA (field programmable gate array) or an ASIC (application-specific integrated circuit). The apparatus can also include, in addition to hardware, code that creates an execution environment for the computer program in question, e.g., code that constitutes processor firmware, a protocol stack, a database management system, an operating system, a cross-platform runtime environment, a virtual machine, or a combination of one or more of them. The apparatus and execution environment can realize various different computing model infrastructures, such as web services, distributed computing and grid computing infrastructures.

A computer program (also known as a program, software, software application, script, or code) can be written in any form of programming language, including compiled or interpreted languages, declarative or procedural languages, and it can be deployed in any form, including as a stand-alone program or as a module, component, subroutine, object, or other unit suitable for use in a computing environment. A computer program may, but need not, correspond to a file in a file system. A program can be stored in a portion of a file that holds other programs or data (e.g., one or more scripts stored in a markup language document), in a single file dedicated to the program in question, or in multiple coordinated files (e.g., files that store one or more modules, sub-programs, or portions of code). A computer program can be deployed to be executed on one computer or on multiple computers that are located at one site or distributed across multiple sites and interconnected by a communication network.

The processes and logic flows described in this specification can be performed by one or more programmable processors executing one or more computer programs to perform actions by operating on input data and generating output. The processes and logic flows can also be performed by, and apparatus can also be implemented as, special

purpose logic circuitry, e.g., an FPGA (field programmable gate array) or an ASIC (application-specific integrated circuit).

Processors suitable for the execution of a computer program include, by way of example, both general and special purpose microprocessors, and any one or more

5 processors of any kind of digital computer. Generally, a processor will receive instructions and data from a read-only memory or a random access memory or both.

Elements of a computer can include a processor for performing actions in accordance with instructions and one or more memory devices for storing instructions and data.

Generally, a computer will also include, or be operatively coupled to receive data from or

10 transfer data to, or both, one or more mass storage devices for storing data, e.g., magnetic, magneto-optical disks, or optical disks. However, a computer need not have such

devices. Moreover, a computer can be embedded in another device, e.g., a mobile

telephone, a personal digital assistant (PDA), a mobile audio or video player, a game console, a Global Positioning System (GPS) receiver, or a portable storage device (e.g., a

15 universal serial bus (USB) flash drive), to name just a few. Devices suitable for storing computer program instructions and data include all forms of non-volatile memory, media and memory devices, including by way of example semiconductor memory devices, e.g.,

EPROM, EEPROM, and flash memory devices; magnetic disks, e.g., internal hard disks or removable disks; magneto-optical disks; and CD-ROM and DVD-ROM disks. The

20 processor and the memory can be supplemented by, or incorporated in, special purpose logic circuitry.

To provide for interaction with a user, implementations of the subject matter described in this specification can be implemented on a computer having a display device, e.g., a CRT (cathode ray tube) or LCD (liquid crystal display) monitor, for displaying

25 information to the user and a keyboard and a pointing device, e.g., a mouse or a trackball, by which the user can provide input to the computer. Other kinds of devices can be used

to provide for interaction with a user as well; for example, feedback provided to the user can be any form of sensory feedback, e.g., visual feedback, auditory feedback, or tactile

30 feedback; and input from the user can be received in any form, including acoustic, speech, or tactile input. In addition, a computer can interact with a user by sending documents to

and receiving documents from a device that is used by the user; for example, by sending web pages to a web browser on a user's client device in response to requests received from the web browser.

Implementations of the subject matter described in this specification can be implemented in a computing system that includes a back-end component, e.g., as a data server, or that includes a middleware component, e.g., an application server, or that includes a front-end component, e.g., a client computer having a graphical user interface or a Web browser through which a user can interact with an implementation of the subject matter described in this specification, or any combination of one or more such back-end, middleware, or front-end components. The components of the system can be interconnected by any form or medium of digital data communication, e.g., a communication network. Examples of communication networks include a local area network (“LAN”) and a wide area network (“WAN”), an inter-network (e.g., the Internet), and peer-to-peer networks (e.g., ad hoc peer-to-peer networks).

The computing system can include clients and servers. A client and server are generally remote from each other and typically interact through a communication network. The relationship of client and server arises by virtue of computer programs running on the respective computers and having a client-server relationship to each other. In some implementations, a server transmits data (e.g., an HTML page) to a client device (e.g., for purposes of displaying data to and receiving user input from a user interacting with the client device). Data generated at the client device (e.g., a result of the user interaction) can be received from the client device at the server.

While this specification contains many specific implementation details, these should not be construed as limitations on the scope of any implementation of the present disclosure or of what may be claimed, but rather as descriptions of features specific to example implementations. Certain features that are described in this specification in the context of separate implementations can also be implemented in combination in a single implementation. Conversely, various features that are described in the context of a single implementation can also be implemented in multiple implementations separately or in any suitable sub-combination. Moreover, although features may be described above as acting in certain combinations and even initially claimed as such, one or more features from a claimed combination can in some cases be excised from the combination, and the claimed combination may be directed to a sub-combination or variation of a sub-combination.

Similarly, while operations are depicted in the drawings in a particular order, this should not be understood as requiring that such operations be performed in the particular order shown or in sequential order, or that all illustrated operations be performed, to achieve desirable results. In certain circumstances, multitasking and parallel processing

may be advantageous. Moreover, the separation of various system components in the implementations described above should not be understood as requiring such separation in all implementations, and it should be understood that the described program components and systems can generally be integrated together in a single software product

5 or packaged into multiple software products.

Thus, particular implementations of the subject matter have been described. Other implementations are within the scope of the following claims. In some cases, the actions recited in the claims can be performed in a different order and still achieve desirable results. In addition, the processes depicted in the accompanying figures do not

10 necessarily require the particular order shown, or sequential order, to achieve desirable results. In certain implementations, multitasking and parallel processing may be advantageous.

CLAIMS

1. A computer-implemented method executed using one or more processors, the method comprising:

receiving, by the one or more processors, content data, the content data comprising a plurality of comments associated with an item of digital content;

processing, by the one or more processors, comments of the plurality of comments to provide a set of relevant sentences;

receiving, by the one or more processors, a set of entities comprising one or more entities, each entity in the set of entities being provided based on the set of relevant sentences and being associated with a respective score;

selecting, by the one or more processors, at least one entity of the set of entities based on respective scores; and

associating, by the one or more processors, the at least one entity with the item of digital content.

2. The method of claim 1, wherein processing comments of the plurality of comments to provide a set of relevant sentences comprises:

5 receiving a set of sentences, sentences in the set of sentences comprising sentence tokens and word tokens; and

filtering sentences from the set of sentences to provide a sub-set of sentences, the set of relevant sentences being provided based on the sub-set of sentences.

10 3. The method of claim 2, wherein sentences are filtered from the set of sentences based on at least one of language, length, punctuation, symbols, and letter case.

4. The method of claim 1, wherein processing comments of the plurality of comments to provide a set of relevant sentences comprises providing the sub-set of
15 sentences and digital content data to a summarizer component, the summarizer component processing the sub-set of sentences to provide the set of relevant sentences.

5. The method of claim 4, wherein the summarizer component processes the sub-set of sentences to provide at least one word set, the at least one word set comprising words that are common to sentences of the sub-set of sentences, the set of relevant sentences being provided based on the at least one word set.

5

6. The method of claim 1, further comprising providing the set of relevant sentences to an entity identifier component, wherein the entity identifier component processes the set of relevant sentences to provide the set of entities.

10 7. The method of claim 1, wherein the respective scores each indicate a frequency, at which a respective entity is identified in the set of relevant sentences.

8. The method of claim 1, wherein each entity of the set of entities comprises at least one of a person, a location, a thing, and a concept.

15

9. The method of claim 1, wherein the item of digital content comprises a video.

10. A system comprising:

a data store for storing data; and

20 one or more processors configured to interact with the data store, the one or more processors being further configured to perform operations comprising:

receiving content data, the content data comprising a plurality of comments associated with an item of digital content;

processing comments of the plurality of comments to provide a set of relevant sentences;

receiving a set of entities comprising one or more entities, each entity in the set of entities being provided based on the set of relevant sentences and being associated with a respective score;

selecting at least one entity of the set of entities based on respective scores;

and

associating the at least one entity with the item of digital content.

11. The system of claim 10, wherein processing comments of the plurality of comments to provide a set of relevant sentences comprises:

25

receiving a set of sentences, sentences in the set of sentences comprising sentence tokens and word tokens; and

filtering sentences from the set of sentences to provide a sub-set of sentences, the set of relevant sentences being provided based on the sub-set of sentences.

5

12. The system of claim 11, wherein sentences are filtered from the set of sentences based on at least one of language, length, punctuation, symbols, and letter case.

10

13. The system of claim 10, wherein processing comments of the plurality of comments to provide a set of relevant sentences comprises providing the sub-set of sentences and digital content data to a summarizer component, the summarizer component processing the sub-set of sentences to provide the set of relevant sentences.

15

14. The system of claim 13, wherein the summarizer component processes the sub-set of sentences to provide at least one word set, the at least one word set comprising words that are common to sentences of the sub-set of sentences, the set of relevant sentences being provided based on the at least one word set.

20

15. The system of claim 10, wherein operations further comprise providing the set of relevant sentences to an entity identifier component, wherein the entity identifier component processes the set of relevant sentences to provide the set of entities.

25

16. The system of claim 10, wherein the respective scores each indicate a frequency, at which a respective entity is identified in the set of relevant sentences.

17. The system of claim 10, wherein each entity of the set of entities comprises at least one of a person, a location, a thing, and a concept.

30

18. The system of claim 10, wherein the item of digital content comprises a video.

19. A computer readable medium storing instructions that, when executed by one or more processors, cause the one or more processors to perform operations comprising:

receiving content data, the content data comprising a plurality of comments associated with an item of digital content;

processing comments of the plurality of comments to provide a set of relevant sentences;

receiving a set of entities comprising one or more entities, each entity in the set of entities being provided based on the set of relevant sentences and being associated with a respective score;

selecting at least one entity of the set of entities based on respective scores; and associating the at least one entity with the item of digital content.

20. The computer readable medium of claim 19, wherein processing comments of the plurality of comments to provide a set of relevant sentences comprises:

receiving a set of sentences, sentences in the set of sentences comprising sentence tokens and word tokens; and

5 filtering sentences from the set of sentences to provide a sub-set of sentences, the set of relevant sentences being provided based on the sub-set of sentences.

21. The computer readable medium of claim 20, wherein sentences are filtered from the set of sentences based on at least one of language, length, punctuation, symbols, and
10 letter case.

22. The computer readable medium of claim 19, wherein processing comments of the plurality of comments to provide a set of relevant sentences comprises providing the sub-set of sentences and digital content data to a summarizer component, the summarizer
15 component processing the sub-set of sentences to provide the set of relevant sentences.

23. The computer readable medium of claim 22, wherein the summarizer component processes the sub-set of sentences to provide at least one word set, the at least one word set comprising words that are common to sentences of the sub-set of sentences, the set of
20 relevant sentences being provided based on the at least one word set.

24. The computer readable medium of claim 19, wherein operations further comprise providing the set of relevant sentences to an entity identifier component, wherein the entity identifier component processes the set of relevant sentences to provide the set of
25 entities.

25. The computer readable medium of claim 19, wherein the respective scores each indicate a frequency, at which a respective entity is identified in the set of relevant sentences.

5 26. The computer readable medium of claim 19, wherein each entity of the set of entities comprises at least one of a person, a location, a thing, and a concept.

27. The computer readable medium of claim 19, wherein the item of digital content comprises a video.

10





FIG. 1






FIG. 3

INTERNATIONAL SEARCH REPORT

International application No
PCT/GR2013/000065

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06Q30/02
ADD.

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
G06Q

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

EPO-Internal

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2012/303713 A1 (HARBICK ANDREW V [US] ET AL) 29 November 2012 (2012-11-29) abstract	1-27
X	US 2010/071005 A1 (KUSUNOKI YOSHIAKI [JP]) 18 March 2010 (2010-03-18) abstract	1-27



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

24 February 2014

Date of mailing of the international search report

03/03/2014

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040,
Fax: (+31-70) 340-3016

Authorized officer

Weidmann, Matthias

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/GR2013/000065

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2012303713 A1	29-11-2012	US 8260656 B1	04-09-2012
		US 2012303713 A1	29-11-2012
		US 2013254065 A1	26-09-2013

US 2010071005 A1	18-03-2010	CN 101677376 A	24-03-2010
		JP 5322550 B2	23-10-2013
		JP 2010074518 A	02-04-2010
		US 2010071005 A1	18-03-2010
