



US 20230186348A1

(19) **United States**

(12) **Patent Application Publication**  
**Sarkar et al.**

(10) **Pub. No.: US 2023/0186348 A1**

(43) **Pub. Date: Jun. 15, 2023**

(54) **IMAGE RECOGNITION BASED CONTENT ITEM SELECTION**

(71) Applicant: **Google LLC**, Mountain View, CA (US)

(72) Inventors: **Subhadip Sarkar**, Sunnyvale, CA (US); **Ravi Jain**, Palo Alto, CA (US)

(21) Appl. No.: **18/168,217**

(22) Filed: **Feb. 13, 2023**

**Publication Classification**

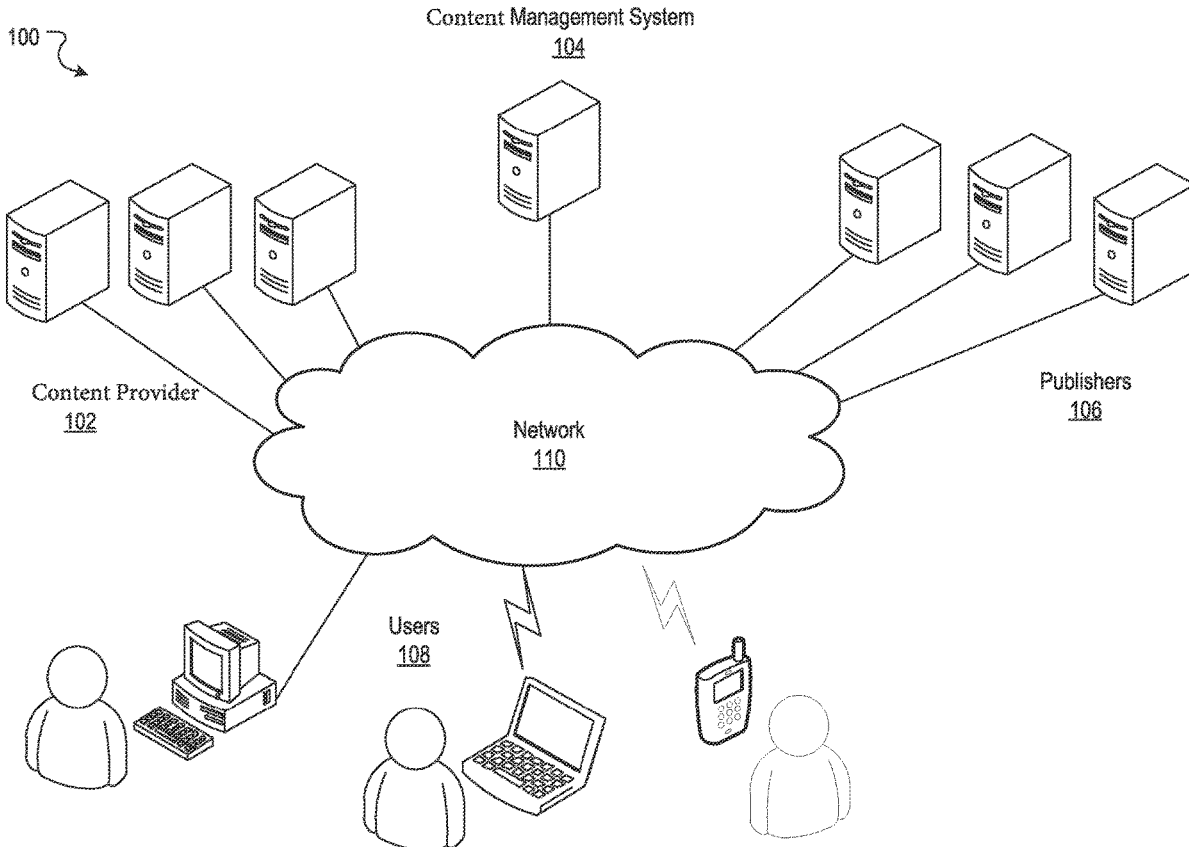
(51) **Int. Cl.**  
**G06F 16/242** (2019.01)  
**G06Q 90/00** (2006.01)  
**G06F 16/9535** (2019.01)  
**G06F 16/58** (2019.01)  
(52) **U.S. Cl.**  
CPC ..... **G06Q 30/0256** (2013.01); **G06F 16/242** (2019.01); **G06Q 90/00** (2013.01); **G06F 16/9535** (2019.01); **G06F 18/24** (2023.01); **G06F 16/5866** (2019.01)

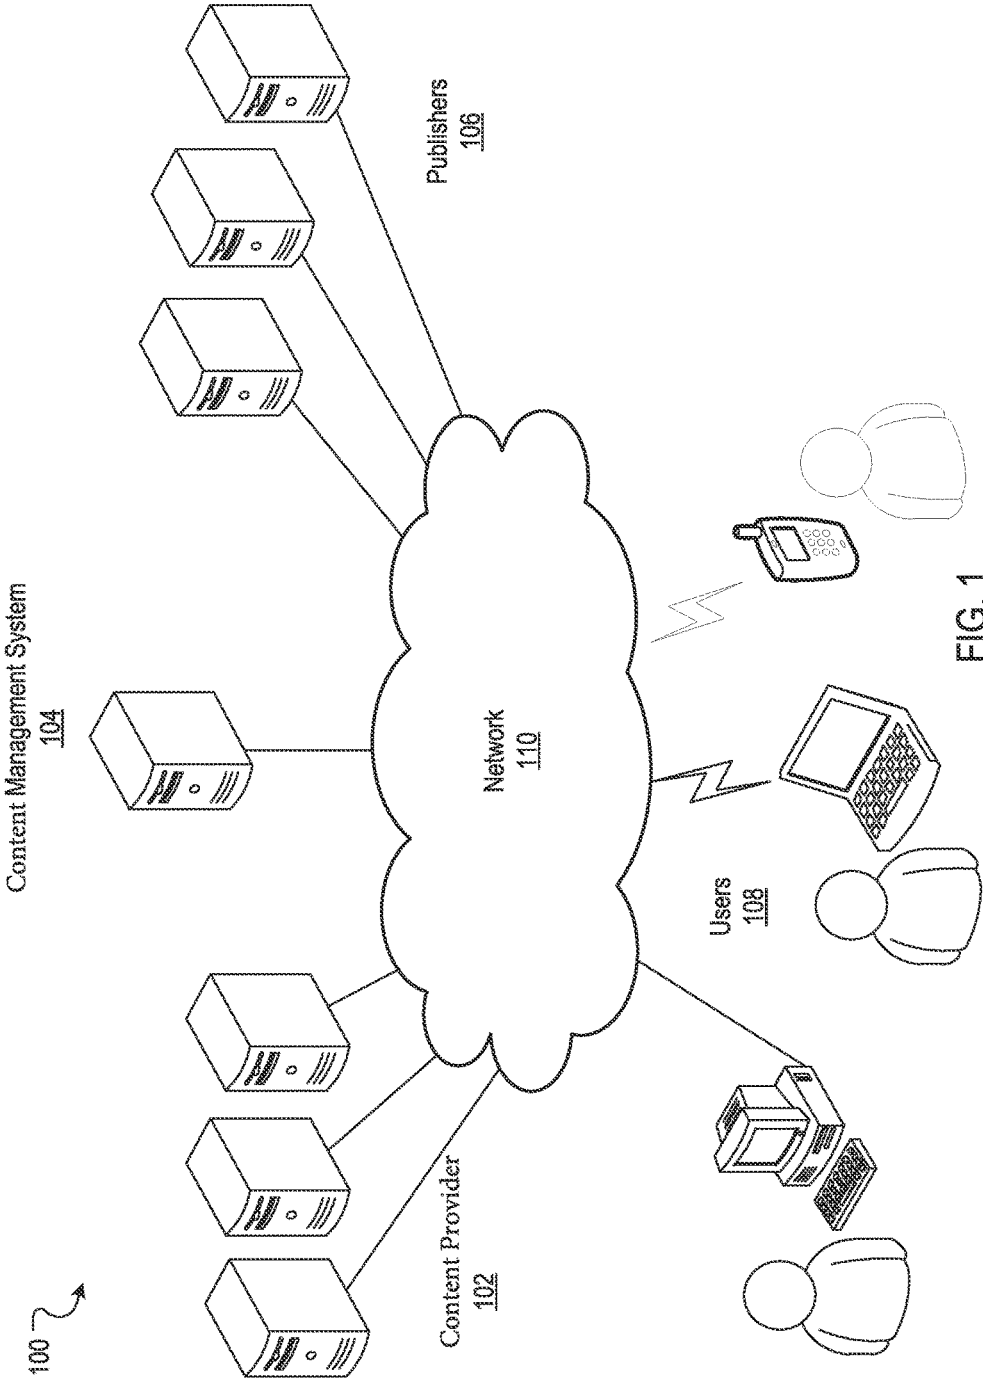
**Related U.S. Application Data**

(63) Continuation of application No. 17/394,123, filed on Aug. 4, 2021, now Pat. No. 11,593,906, which is a continuation of application No. 15/638,333, filed on Jun. 29, 2017, now Pat. No. 11,100,538, which is a continuation of application No. 14/190,523, filed on Feb. 26, 2014, now abandoned, which is a continuation of application No. 13/168,904, filed on Jun. 24, 2011, now Pat. No. 8,688,514.

(57) **ABSTRACT**

Methods, systems, and apparatus, including computer programs encoded on a computer storage medium, for content presentation. In one aspect, a method includes receiving a request for a content item to be presented to a user; for a collection of content items, identifying one or more candidate content items including using image keywords derived from images associated with the user to identify the one or more candidate content items; selecting a content item responsive to the content request from the candidate content items; and sending the selected content item for presentation to the user.





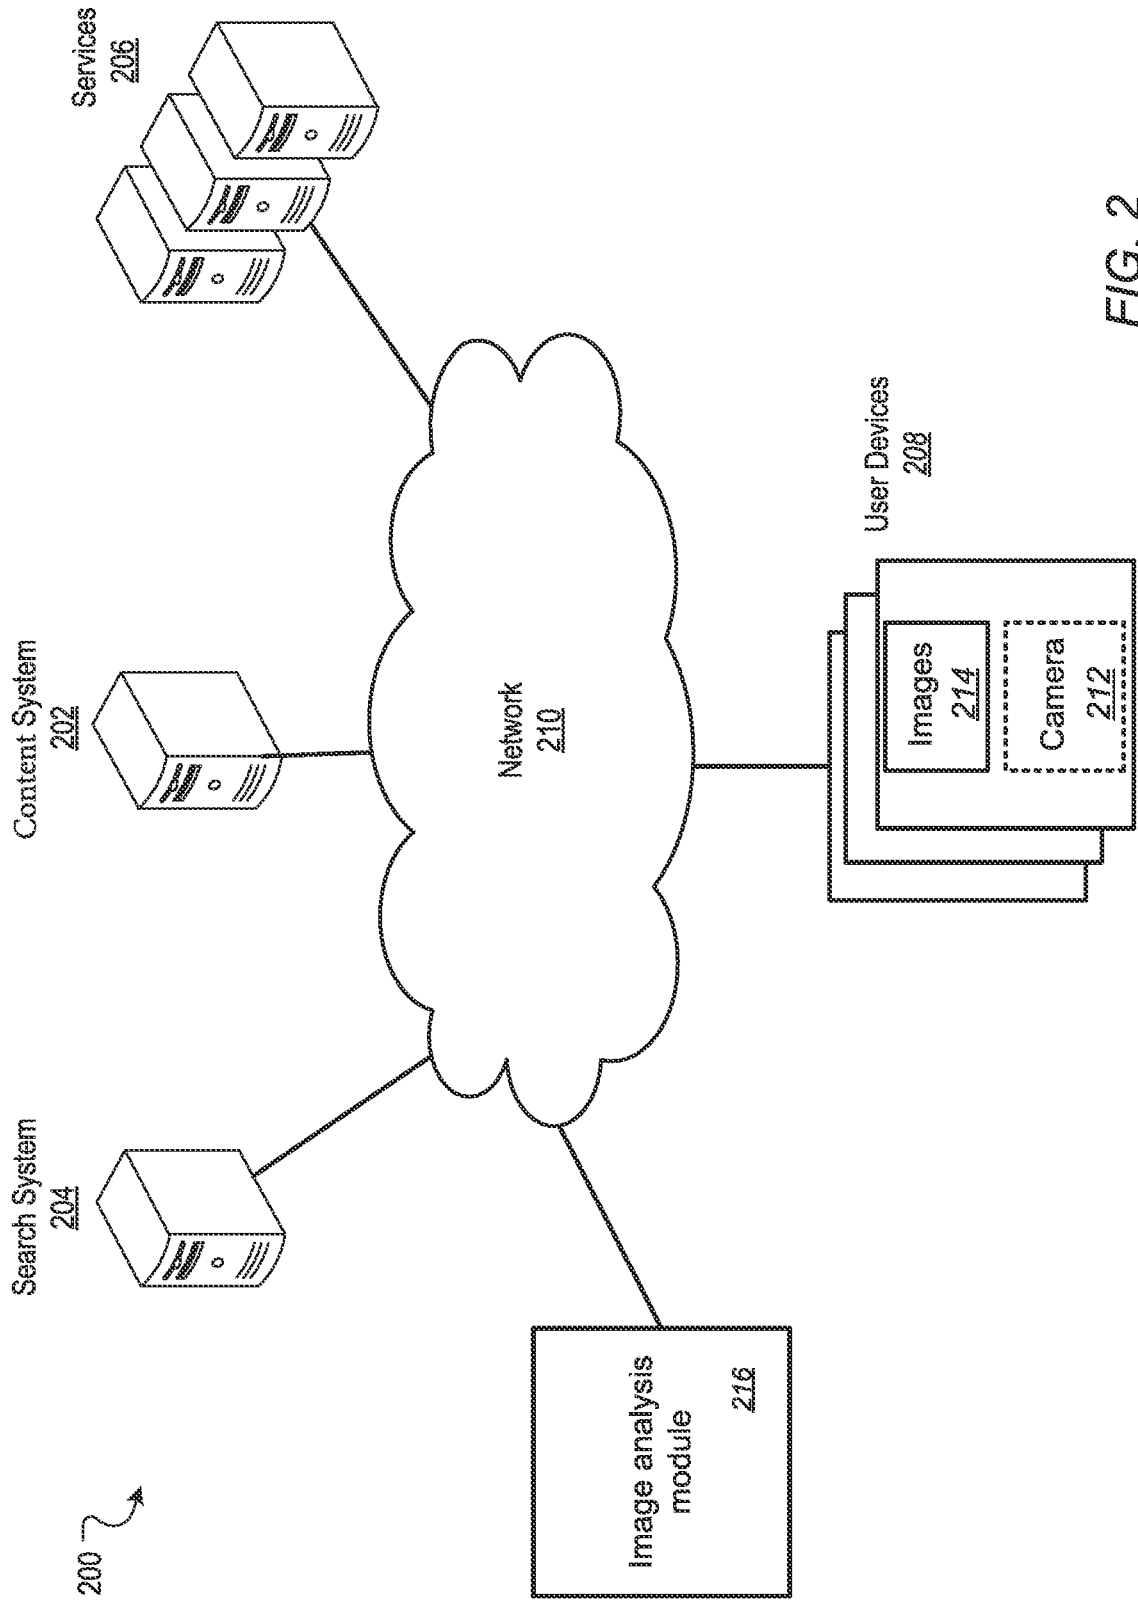


FIG. 2

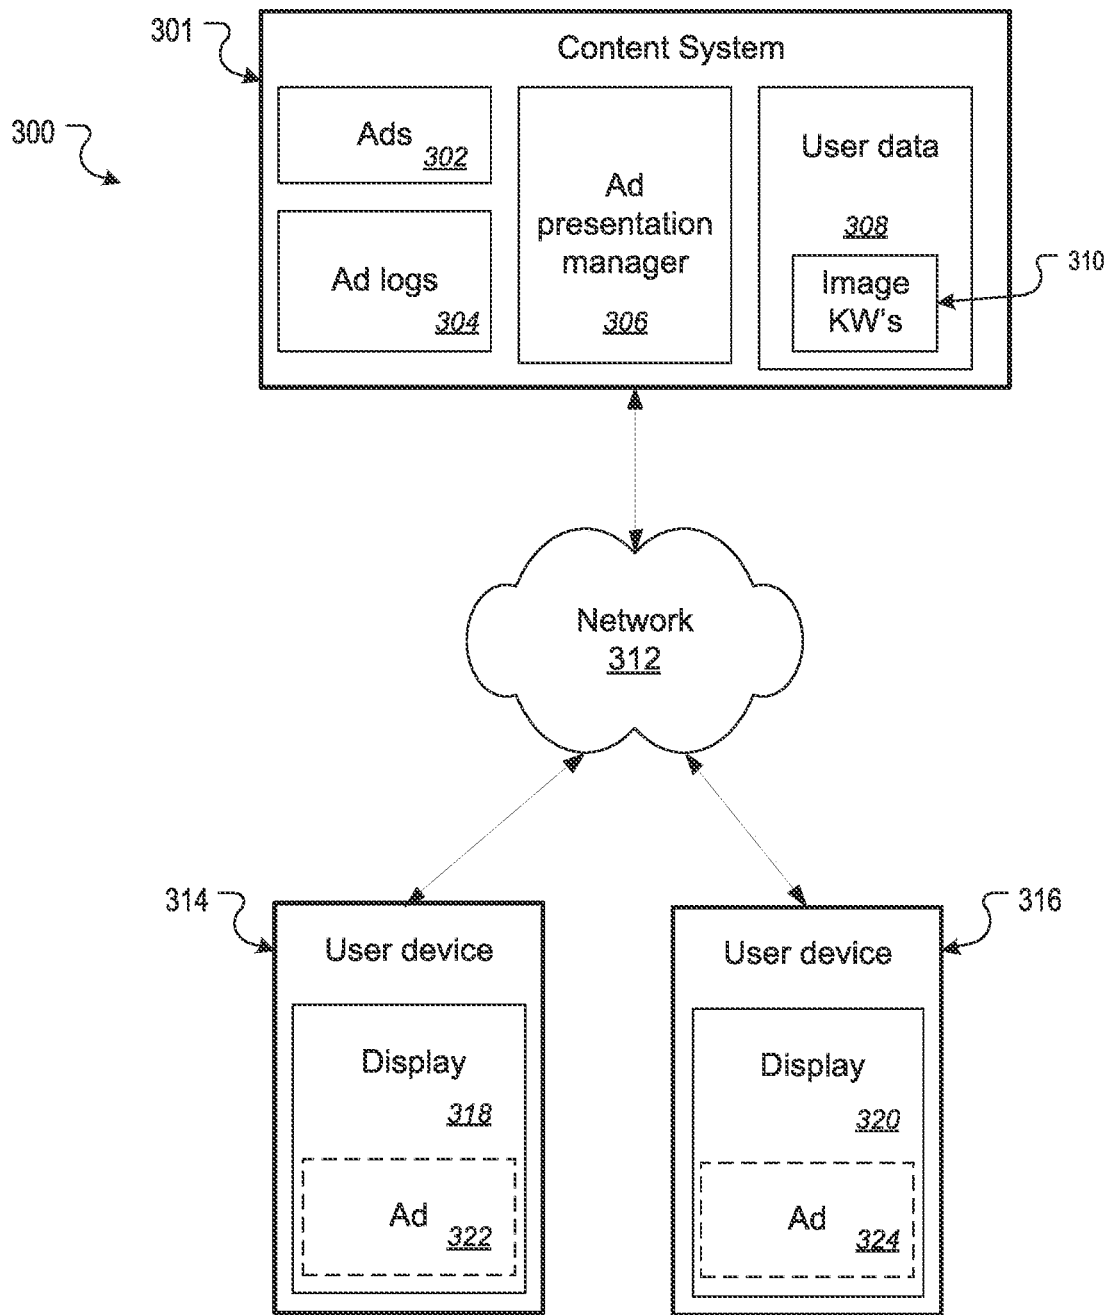


FIG. 3

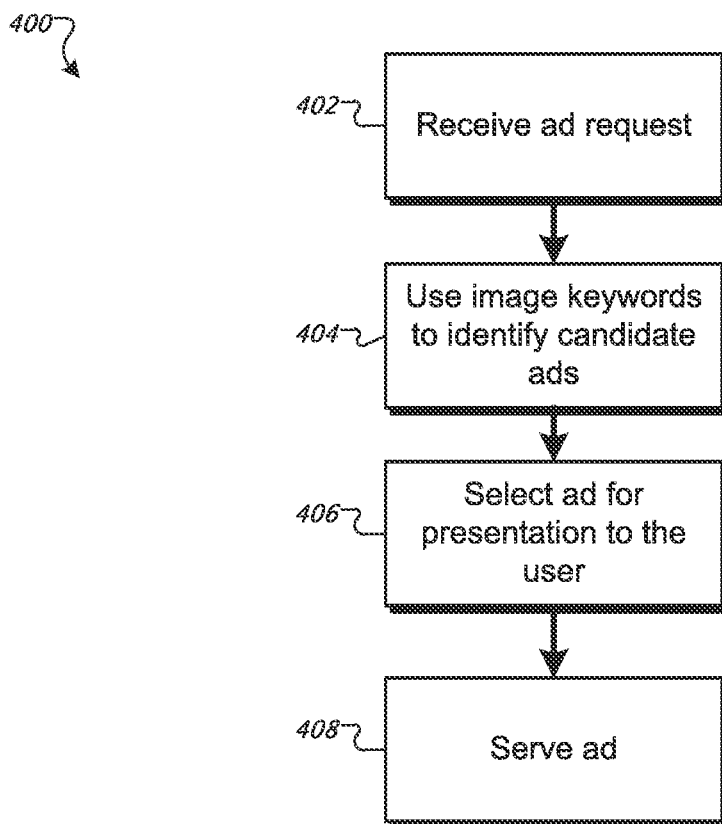


FIG. 4

## IMAGE RECOGNITION BASED CONTENT ITEM SELECTION

### CROSS-REFERENCE TO RELATED APPLICATIONS

**[0001]** This application claims the benefit of priority under 35 U.S.C. § 120 as a continuation of U.S. patent application Ser. No. 15/638,333, filed Jun. 29, 2017, which claims the benefit of priority under 35 U.S.C. § 120 as a continuation of U.S. patent application Ser. No. 14/190,523, filed Feb. 26, 2014, which claims the benefit of priority under 35 U.S.C. § 120 as a continuation of U.S. patent application Ser. No. 13/168,904, filed Jun. 24, 2011, each of which are hereby incorporated herein by reference in their entirety.

### BACKGROUND

**[0002]** Computing systems can present content items. However, it may be challenging for a computing system to select content when there is minimal context.

### SUMMARY

**[0003]** This specification describes technologies relating to content presentation.

**[0004]** In general, one aspect of the subject matter described in this specification can be embodied in methods that include the actions of receiving a request for an ad to be presented to a user; for a collection of ads, identifying one or more candidate ads including using image keywords derived from images associated with the user to identify the one or more candidate ads; selecting an ad responsive to the ad request from the candidate ads; and sending the selected ad for presentation to the user. Other embodiments of this aspect include corresponding systems, apparatus, and computer programs recorded on computer storage devices, each configured to perform the operations of the methods.

**[0005]** These and other embodiments can each optionally include one or more of the following features. The image keywords are derived from images uploaded by the user. The images include images captured by the user using a camera of a mobile device. The image keywords are derived from images selected by the user. Selecting an ad responsive to the ad request includes determining a similarity between one or more keywords associated with each ad of the collection of ads and the image keywords. Determining the similarity include calculating a similarity score for each ad and where selecting the ad responsive to the ad request includes using a combination of the similarity score and one or more other scores for each. Each image keyword is associated with a scoring weight for scoring the respective image keywords relative to the ads in the collection of ads and where the weight decreases according to age of the image keyword. The ad request is the first ad request following user image activity. The image content is used to directly infer the user's intent.

**[0006]** Particular embodiments of the subject matter described in this specification can be implemented so as to realize one or more of the following advantages. Using keywords derived from images uploaded or selected by a user provides information for presenting ads directed to the user's interests. Images can provide context for identifying ads in situations where there may be little context available, for example, for mobile applications. Images can also pro-

vide user intent directly. For example, if the input image is for a digital camera, the system can identify suitable ads based on the image content.

**[0007]** The details of one or more embodiments of the subject matter described in this 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

**[0008]** FIG. 1 is a diagram of an example content presentation system.

**[0009]** FIG. 2 is a diagram of an example system.

**[0010]** FIG. 3 is an example system for serving ads.

**[0011]** FIG. 4 is an example process for presenting ads to users.

**[0012]** Like reference numbers and designations in the various drawings indicate like elements.

### DETAILED DESCRIPTION

**[0013]** Images can be used to provide context for selecting ads. In particular, images associated with a user can provide context in situations where there is little context available, for example, when serving ads to mobile applications. The image context can include one or more keywords derived from images associated with particular users. For example, users can upload images (for example, to a photo album, a search system, or mail system as an attachment), which can be processed to identify keywords. Alternatively, images selected by users (e.g., in response to presented image search results) can be processed to identify keywords. These keywords can be used to identify particular ads for presentation to the respective users in response to an ad request.

**[0014]** While reference will be made below to content selection systems and methods, other forms of content including other forms of content items can be managed, presented, and logged in accordance with the description below.

**[0015]** FIG. 1 is a block diagram of an example content presentation system 100. In some implementations, one or more content providers 102 can directly, or indirectly, enter, maintain, and log ad information in a content management system 104 ("system 104"). Though reference is made to content selection, other forms of content, including other forms of sponsored content, can be delivered by the system 100. The ads can be in the form of graphical ads, for example, banner ads, text only ads, image ads, barcode ads (e.g., ads that include one or more barcodes, for example, for use in redeeming the particular ad), audio ads, video ads, animated ads, ads combining one or more of any of such components, etc. The ads can also include embedded information, such as links, meta-information, and/or machine executable instructions. One or more users 108 may submit requests for ads to the system 104. The system 104 responds by sending ads to the requesting users 108 for placement on or association with one or more of a publisher's 106 content items (e.g., web properties, mobile applications, or other third party content). Example web properties can include web pages, television and radio content item slots, or print media space.

**[0016]** Users 108 and the content providers 102, can provide usage information to the system 104, for example,

whether or not a conversion (e.g., a purchase or other interaction) or a click-through related to an ad (e.g., a user has selected an ad) has occurred. This usage information can include measured or observed user behavior related to ads that have been served. The system **104** may perform financial transactions, for example, crediting the publishers **106** and charging the content providers **102** based on the usage information.

**[0017]** A network **110**, such as a local area network (LAN), wide area network (WAN), the Internet, one or more telephony networks or a combination thereof, connects the content providers **102**, the system **104**, the publishers **106**, and the users **108**.

**[0018]** One example publisher **106** is a general content server that receives requests for content (e.g., articles, discussion threads, music, video, graphics, search results, web page listings, information feeds, etc.), and retrieves the requested content in response to the request.

**[0019]** The user **108** can submit a request for ads to a content server in the system **104**. The ad request can include a number of ads desired. The ad request can also include content request information, including information associated with the user **108** or the user's device (e.g., a mobile device), as described in greater detail below.

**[0020]** In some implementations, a client browser, application, or other content provider, combines the requested content with one or more of the ads provided by the system **104**. The combined content and ads can be rendered to the users **108** that requested the content for presentation in a viewer (e.g., a browser, application interface, or other content display system) of a user device. The user device can transmit information about the ads back to the content server, including information describing how, when, and/or where the ads are to be rendered (e.g., in HTML or JavaScript™). In some implementations, the content and ads are combined separately. For example, content can be rendered while waiting for ads to be received and incorporated into the content.

**[0021]** Another example publisher **106** is a search service. A search service can receive queries for search results. In response, the search service can retrieve relevant search results from an index of documents (e.g., from an index of web pages). Search results can include, for example, lists of web page titles, snippets of text extracted from those web pages, and hypertext links to those web pages, and may be grouped into a predetermined number of (e.g., ten) search results. In some implementations, the search results can be presented to users with one or more ads.

**[0022]** A user **108**, or the user's device, can submit a request for ads to the system **104**. The request may include a number of ads desired. The system **104** can use various information to identify a group of ads to provide to the user **108**. Ads from this group of ads can be provided by the user's device for presentation in response to a request from a user application, or associated with content to be displayed (e.g., a web page or search results). Thus, ads can be provided directly to the user device for presentation or indirectly through content provided to the user device (e.g., ads presented with search results where the ads are provided to the user with the search results).

**[0023]** In some implementations, the content management system **104** can use an auction process to select ads from the content providers **102**, which can then be provided to the users **108** in response to the receive requests for ads. For

example, the content providers **102** may be permitted to select, or bid, an amount the content providers are willing to pay for each presentation of or interaction with (e.g., click) of an ad, e.g., a cost-per-click amount an content provider pays when, for example, a user clicks on an ad. The cost-per-click can include a maximum cost-per-click, e.g., the maximum amount the content provider is willing to pay for each click of an ad based on a keyword, e.g., a word or words in a query. Other bid types, however, can also be used. Based on these bids, ads can be selected and ranked for presentation.

**[0024]** The user **108** can combine the search results with one or more of the ads provided by the system **104**. This combined information can then be displayed by the devices of the users **108** that requested the content. Alternatively, portions can be received at different times and integrated in the display as received. For example, the ads can be received and rendered separately by a browser or application. The search results can be maintained as distinct from the ads, so as not to confuse the user between paid ads and presumably neutral search results.

**[0025]** In some implementations, one or more publishers **106** provide web properties to the users **108**. The users **108** submit requests for ads to the content management system **104**. The system **104** responds by sending ads to the requesting users **108** for placement on one or more web properties (e.g., websites and other network-distributed content) that are relevant to the web property. In some implementations, the user's device identifies a particular ad from the received ads for presentation with the provided web properties. In some implementations, the requests are executed by devices associated with the user **108**, e.g., by the execution of a particular script (e.g., JavaScript™) when the publisher's web page is loading on a client device.

**[0026]** Another example publisher **106** is a mobile application developer. A mobile application is an application specifically designed for operation on a mobile device (e.g., a smart phone). The mobile application can also include ads positioned within the content of the mobile application. Similar to publishers **106** described above, the ads can be received from the content management system **104** for placement in the mobile application when accessed by a user (e.g., when a particular page of a mobile application is loaded on the mobile device).

**[0027]** FIG. 2 is a diagram of an example system **200**. The example system **200** includes a content system **202**, a search system **204**, services **206**, and an image analysis module **216**, each of which can communicate through one or more network **210**.

**[0028]** The content system **202** provides ads for presentation to users (e.g., of user devices **208**) in response to ad requests. The ad requests can be received from the user devices **208** or from other sources, for example, the search system **204** or content publishers (not shown). The content system **202** can use image keywords associated with particular users to identify candidate ads to provide to the user devices **208**.

**[0029]** The search system **204** can provide search results to the user devices **208** responsive to received queries. In some implementations, the search results are image search results, for example, thumbnail images corresponding to image resources responsive to the received query. Users presented with image search results can select particular images from the image search results in order to view the

corresponding image resource. Image search results can be provided with one or more ads, for example, in a particular portion or portions of the search user interface. In some implementations, the search system 204 includes an image search in which users can submit images as a query (e.g., an image taken from a camera on a mobile device).

[0030] The services 206 can include network photo albums or other photo repositories to which individual users can upload images. The services 206 can also include a mail service where, for example, users can receive images in messages or can attach images in messages to be sent to others.

[0031] The user devices 208 can include mobile devices, desktop computing devices, notebook computing devices, personal data assistants, tablet devices, etc. One or more of the user devices 208 can include a camera 212. The camera 212 can be a still camera or a video camera. The camera 212 can be, for example, a camera integrated into the device (e.g., a mobile device with an integrated camera) or a camera attached to a user device (e.g., a webcam attached to a desktop computer).

[0032] The user devices 208 also includes images 214. The images 214 can include images captured using camera 212 or from various other sources. For example, a user can transfer images from a digital camera to the images 214. In another example, a user can receive images from other users or network locations (e.g., by downloading the images or receiving images by e-mail). One or more images from the images 214 can also be uploaded to network locations using the network 210. For example, a user can upload one or more images to a web album provided by the services 206 or as a query to an image search provided by the search system 204. In some implementations, the images are encrypted prior to being uploaded.

[0033] The image analysis module 216 can analyze images uploaded, for example, to search system 204 or to services 206 in order to identify image keywords. The image analysis module 216 can also analyze images selected by users, for example, from presented image search results. The image analysis module 216 can then send the image keywords to the ad system 202. Image analysis is described in greater detail below. In some implementations, the image analysis module 216 is part of a system affiliated with each of the search system 204, ad system 202, and image services 206. In some other implementations, the image analysis module 216 is a component of one or more of the search system 204, ad system 202, and image services 206. In some implementations, the images are encrypted prior to being uploaded to the image analysis module 216. Additionally, uploaded images can be stored in a protected database and/or not logged into the system.

[0034] FIG. 3 is an example system 300 for serving ads. The system 300 includes an ad system 301, user device 314, and user device 316. The ad system 301 can communicate with each of the user devices 314 and 316 through a network 312.

[0035] The ad system 301 includes ads 302, ad logs 304, ad presentation manager 306, and user data 308.

[0036] The ads 302 are a collection of ads available for presentation to users. Each ad can include associated information used to determine whether the ad should be presented to a user in response to a request, for example, keywords associated with the ad and demographic goals for the ad (e.g., particular groups to which the ad is directed).

The ad presentation manager 306 can select ads for presentation on user devices. The ad logs 304 can store presentation and conversion events (e.g., users an ad is presented to, time stamps, ad identifiers, device identifiers, wallet status of ads, user ad selections, publisher conversion events, etc.).

[0037] The ad presentation manager 306 can determine which ad or ads of the ads 302 to select for presentation (e.g., on user device 314 and/or user device 316) in response to an ad request. For example, the ad presentation manager 306 can identify and rank candidate ads from the ads 302 based on particular criteria in order to match ads with requested content, a requesting API for an application, and/or a requesting user.

[0038] The ad presentation manager 306 can identify candidate ads for selection based on keywords associated with the ads in the collection that correspond with the ad request. In some implementations, the ad keywords can be matched with the requested content in which the ads are to be presented (e.g., matching ad keywords to keywords in requested content or query terms). For example, each ad of the ads 302 can include one or more keywords for identifying ads as directed to particular subject matter that can be matched to the ad request. Thus, particular ads can be identified as relevant to particular users. Candidate ads for selection can be identified based on how well the keywords match the request.

[0039] In another example, the ad presentation manager 306 uses information about a user (e.g., user profile information, other information) to identify candidate ads for selection. For example, particular ads in the ads 302 can include group information (e.g., intended for women between the ages of 30 and 35). This information can be compared with the user information to determine whether or not an ad is a candidate ad for selection.

[0040] In yet another example, the ad presentation module 306 can send ads to an application based on the specific application requesting the ads (e.g., the ads are for products or services that appeal to a typical user of the particular application) or based on user interactions with the application.

[0041] The ad presentation manager 306 can further use image keywords associated with the user in identifying ads for presentation, for example, by matching image keywords with ad keywords to identify ad candidates. In some implementations, image keywords are only used to identify ads in particular circumstances. For example, when other sources of information are insufficient or unavailable (e.g., lack of other user information or lack of information associated with the content with which the ad is to be presented). In another example, image keywords are only used within a particular time window of the user image activity (e.g., when the associated image is also associated with the ad request or the ad request follows particular image activity, for example, an image upload or selection). In some implementations, the image is used to identify a primary user intent, e.g., if the user submits an image to the image search system in order to obtain information regarding the image. For example, if the user uploads a particular image, the content of that image can be used to identify relevant ads to present to the user.

[0042] In some implementations, an auction system is used to determine which ad from candidate ads matching criteria for the content will be selected for presentation or to adjust scores for candidate ads. In some other implementations, the candidate ad having a highest score is selected for



presentation. The selected ad can then be sent for presentation to the user (e.g., sent directly to the user device or sent to a publisher or other content provider for integration prior to providing the content to the user device).

**[0043]** The user data **308** can include user information such as user profile information that can be used by the ad presentation manager **306** in identifying matching ads. The user data **308** also includes image keywords **310**. The image keywords are keywords derived from particular images associated with each particular user. In some implementations, the keywords are derived in response to a user uploading an image (e.g., an image taken using a camera on a user device) to a search system. For example, the user can upload an image to be used as a search query. In some other implementations, the keywords are derived in response to a user uploading an image to an image service, for example, to an online photo album. Another image source associated with the user includes images received by or sent to the user using e-mail. A further image source can be images selected by the user, e.g., based on presented image search results.

**[0044]** In some implementations, one or more of the search system and services are affiliated with the ad system in order to facilitate the sharing of derived keyword information.

**[0045]** The image keywords can be derived from various techniques. These keywords can then be provided to the ad system **301** for storage within the image keywords **310**. In some implementations, images associated with the user are directly analyzed. For example, an image can be analyzed to identify text within the image (e.g., using optical character recognition techniques). Identified words or phrases from the image can then be considered keywords. For example an image including a visible sign reading “Tour Eiffel” can yield as keywords “Eiffel” and “Tour”. Additionally, the image can be identified for particular patterns matching known objects. For example a picture of a car can be matched with an image in an image database. Images in the database can have a particular descriptor or caption which can also be used as keywords (e.g., “Peugeot 504”).

**[0046]** In some implementations, images associated with the user are indirectly analyzed. For example, an image (or a discernable object within the image) is matched with existing objects in an image database and metadata associated with those existing objects in the image database are analyzed to obtain keywords. For example, the photo may be similar to a photo from an online guidebook of Paris, and any metadata about the latter (URL, anchor text, etc.) can be analyzed to obtain keywords, e.g. [www.paris.example.com/eiffel](http://www.paris.example.com/eiffel) may yield “paris” and “eiffel” as keywords. Additionally, when a URL is identified, the target location (e.g., a web page) can be crawled to obtain keywords.

**[0047]** In some other implementations, referring search terms can be used to identify keywords. For example, once a matching image is found, search history logs (for images search or any other property) can be analyzed to identify which search queries led to (a) that image being served as a search result and (b) resulted in the image being selected. The search terms can be considered keywords, and their importance can be weighted based on the frequency of occurrence as well as whether they led to a selection of the image. For example, the image from the online guidebook of Paris may be served when the user types “Paris” or “Eiffel”, but clicked more often in the latter case.

**[0048]** The keywords obtained from these methods can be ranked based on various information and filtered to identify the most likely keywords. Additionally, the image keywords can be expanded for each image or user based on vertical information (e.g., higher level categories associated with the identified keywords, for example, “France” for the keyword “Paris”). For example, a system that generates keywords for the image can compute a confidence weight associated with an image-keyword pair and provide that to the ad presentation system. These weights can be used to identify dominant keywords for particular images.

**[0049]** In some implementations, the keywords can be also compared with historical logs to identify best performing keywords for the image (or some representation of the concepts the image contains). Different measures of performance can be used to identify best performing keywords including, for example, aggregated click through rate on the ads, user’s refinement of the query in the context of a search etc. From an ad presentation perspective, preference can be given to commercial keywords that the content providers are directing their ads to. Additionally, data from the user’s past queries and/or interactions can be used to rank the keywords. For example, if the system knows (e.g., based on past user information) that the user owns a camera of a particular brand, keywords that represent accessories for the particular camera brand can be preferred.

**[0050]** In some implementations, user information is anonymized to obscure the user’s identity. For example, received information from users (e.g., user interactions, location, device or user identifiers) can be aggregated or removed/obscure (e.g., replace user identifier with random identifier) so that individually identifying information is anonymized while still maintaining the attributes or characteristics associated with particular information (e.g., types of user actions). In some other implementations, the received information is anonymized (so that the originating user device or user device user is unidentifiable) at the user device before transmission to the system that analyzes the received inputs. In this way, the actions of individual users can be obscured or unobservable while still permitting analysis of user information. Additionally, users can opt-in or opt-out of making data for images associated with the user available to the ad system (e.g., for use in identifying appropriate ads). In some implementations, a notification procedure is used to prompt the user and a user interface can be provided to transparently view and modify privacy settings.

**[0051]** In some implementations, users can specify types of images that can be analyzed. For example, the user can specify that only images uploaded in particular contexts are included. Thus, the user can specify that images uploaded to a search system can be used but not those uploaded to a private album or to a social site. In some implementations, the system has default settings that can be modified by a user. The default settings can be designed to be overly limiting on the criteria for using images associated with the user.

**[0052]** In some other implementations, the system or user can specify particular types of images that are not analyzed or used to identify keywords. For example, images can be processed, on a network location or on the client device, to identify images containing people. As a result of the pro-

cessing, for example, images recognized as having people can be excluded while images recognized as containing objects can be used.

**[0053]** The user devices **314** and **316** each include a respective display **318** and **320** for presenting content to the user of the device. For example, the display can be used to present application content. The display can also be used to present publisher content or search result content using a web browser. The display of content can include the presentation of one or more ads e.g., ad **322** or ad **324**.

**[0054]** FIG. 4 is an example process **400** for presenting ads to users. The process **400** can be performed by a system including one or more computing devices. For example, the process **400** can be performed by an ad system, e.g., the ad system **300** of FIG. 3.

**[0055]** An ad request is received (**402**). The ad request can be received by an ad system from a publisher. For example, a publisher can request an ad to be provided for content to be served to the user device. The ad can be provided to the publisher for inclusion in the content prior to sending to the user device or, alternatively, the ad can be served directly to the user device.

**[0056]** In some other implementations, the client device requests the ad from the ad system. For example, the user device can request an ad in response to rendering content from the publisher. A script in received web page content from a publisher can prompt the user device browser to send an ad request. Alternatively, presenting particular application content (e.g., content in a mobile application) can prompt an ad request to be sent to the ad system.

**[0057]** In response to the received ad request, candidate ads are identified using image keywords (**404**). Image keywords associated with the user can be identified based on the ad request. For example, the ad request can include a user identifier for the user. Image keywords associated with that user can be identified and compared with ad keywords. Ads can be scored based on the degree of match between the ad keywords and the identified image keywords associated with the user. Candidate ads can include ads having a threshold score or a specified number of top scoring ads. In some implementations, the score of the image keyword matching alone is used to identify candidate ads.

**[0058]** In some other implementations, the image keyword matching score is combined with one or more other scores to generate overall ad scores (e.g., scores based on the content to be presented or user demographic information).

**[0059]** In some implementations, the score for the image keyword matching can be discounted according to the age of the image keyword (e.g., how long ago the particular image keywords were added). Thus, keywords associated with more recent user image activity is weighted more heavily than older user image activity. A specified decay function can be used to adjust the image keyword weight.

**[0060]** In some implementations, the ad request is sent contemporaneously with analysis of a particular image such that the image keywords are associated with recent user activity. For example, an ad can be presented directly in response to a received image, where the ad is selected based on image keywords derived from the received image. This can occur, for example, when the image is uploaded through a content provider that includes one or more ads. For example, a user can attach an image to an e-mail message in an e-mail interface that includes one or more ads. The attached image can be analyzed and used to determine a next

ad to present to the user (e.g., when the ads are refreshed in the e-mail interface). In some other implementations, the received image keywords are used to identify subsequent ad or ads to be presented to the user.

**[0061]** A particular ad of the candidate ads is selected for presentation to the user (**406**). In some implementations, a candidate ad having a highest overall score is selected for presentation to the user. The score can be calculated using a combination of scores for a variety of types of information including a score for the matching of image keywords to ads. In some other implementations, an auction process is used to identify a particular ad of the candidate ads for presentation.

**[0062]** The selected ad is sent for presentation (**408**). In some implementations, the ad is sent directly to the client device associated with the ad request. The ad can then be incorporated into particular content (e.g., application content). In some other implementations, the ad is sent to a content provider that incorporates the ad into the content before providing the content to the user device.

**[0063]** Embodiments of the subject matter and the operations described in this specification can be implemented in digital electronic circuitry, or in computer software, firmware, or hardware, including the structures disclosed in this specification and their structural equivalents, or in combinations of one or more of them. Embodiments of the subject matter described in this specification can be implemented as 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 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-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).

**[0064]** 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.

**[0065]** 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 environ-

ment can realize various different computing model infrastructures, such as web services, distributed computing and grid computing infrastructures.

**[0066]** 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.

**[0067]** 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).

**[0068]** 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 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. The essential elements of a computer are 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 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 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 processor and the memory can be supplemented by, or incorporated in, special purpose logic circuitry.

**[0069]** To provide for interaction with a user, embodiments 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 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 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.

**[0070]** Embodiments 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 internet network (e.g., the Internet), and peer-to-peer networks (e.g., ad hoc peer-to-peer networks).

**[0071]** 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 embodiments, 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.

**[0072]** While this specification contains many specific implementation details, these should not be construed as limitations on the scope of any inventions or of what may be claimed, but rather as descriptions of features specific to particular embodiments of particular inventions. Certain features that are described in this specification in the context of separate embodiments can also be implemented in combination in a single embodiment. Conversely, various features that are described in the context of a single embodiment can also be implemented in multiple embodiments separately or in any suitable subcombination. 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 subcombination or variation of a subcombination.

**[0073]** 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 embodiments described above should not be understood as requiring such separation in all embodiments, and it should be understood that the described program components and systems can generally be integrated together in a single software product or packaged into multiple software products.

**[0074]** Thus, particular embodiments of the subject matter have been described. Other embodiments 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 necessarily require the particular order shown, or sequential order, to achieve desirable results. In certain implementations, multitasking and parallel processing may be advantageous.

What is claimed is:

- 1.-27. (canceled)
- 28.** A computing system, the system comprising:
  - one or more processors; and
  - one or more non-transitory computer-readable media that collectively store instructions that, when executed by the one or more processors, cause the computing system to perform operations, the operations comprising:
    - receiving, via an application executed by a computing device connected to a network, an image captured by a camera of the computing device;
    - processing the image captured by the camera to determine one or more keywords associated with the image;
    - determining a plurality of content items responsive to the image capture by the camera based at least in part on the one or more keywords, wherein the plurality of content items comprise one or more sponsored content items; and
    - providing the plurality of content items to the computing device for display in a mobile application.
- 29.** The system of claim **28**, wherein processing the image captured by the camera to determine the one or more keywords comprises:
  - identifying, via optical character recognition, text in the image captured by the camera.
- 30.** The system of claim **28**, wherein processing the image captured by the camera to determine the one or more keywords comprises:
  - determining one or more existing objects in an image database are associated with the image captured by the camera; and
  - processing metadata associated with the one or more existing objects to obtain the one or more keywords.
- 31.** The system of claim **28**, wherein processing the image captured by the camera to determine the one or more keywords comprises:
  - determining a matching image that matches the image captured by the camera;
  - obtaining a search history log associated with the matching image;
  - processing the search history log to determine one or more search queries associated with the matching image; and
  - determining the one or more keywords based on one or more search terms in the one or more search queries.
- 32.** The system of claim **31**, wherein processing the search history log to determine the one or more search queries associated with the matching image comprises:

- processing the search history log to determine the matching image was previously provided as a search result responsive to the one or more search queries.

- 33.** The system of claim **31**, wherein processing the search history log to determine the one or more search queries associated with the matching image comprises:

- processing the search history log to determine the matching image was previously selected after being provided as a search result responsive to the one or more search queries.

- 34.** The system of claim **28**, wherein the operations further comprise:

- expanding the one or more keywords to include categories;
- determining a score for each of the plurality of content items based on matching the one or more keywords for each of the plurality of content items with the one or more keywords including the categories; and
- ranking the plurality of content items based on the scores.

- 35.** The system of claim **28**, wherein the computing device is a mobile computing device with the camera integrated into the mobile computing device.

- 36.** The system of claim **28**, wherein processing the image captured by the camera to determine the one or more keywords comprises:

- processing the image with an image analysis module of a search system.

- 37.** The system of claim **28**, wherein the one or more sponsored content items are determined by:

- obtaining an ad request for an ad to be presented to a user associated with the computing device;
- identifying one or more candidate ads of a collection of ads for presentation to the user, wherein identifying the one or more candidate ads includes scoring each candidate ad based at least in part on a degree of match between the respective ad and the one or more keywords; and
- selecting an ad responsive to the ad request from the candidate ads based on the scores.

- 38.** The system of claim **37**, where selecting an ad responsive to the ad request includes determining a similarity between one or more keywords associated with each ad of the collection of ads and the image keywords, and wherein determining the similarity include calculating a similarity score for each ad, and wherein selecting the ad responsive to the ad request includes using a combination of the similarity score and one or more other scores for each.

- 39.** A computer-implemented method, the method comprising:

- receiving, via an application executed by a computing device connected to a network and by a data processing system including one or more processors and memory, an image captured by a camera of the computing device;

- processing, by the data processing system, the image captured by the camera to determine one or more keywords;

- determining, by the data processing system, a plurality of search results responsive to the image capture by the camera based at least in part on the one or more keywords;

- determining, by the data processing system, one or more ads based on at least a subset of the plurality of search results; and

providing, by the data processing system, the plurality of search results and the one or more ads to the computing device.

**40.** The computer-implemented method of claim **39**, further comprising:

receiving a request for an ad to be presented to a user associated with the computing device;

for a collection of ads, identifying one or more candidate ads for presentation to the user, wherein identifying the one or more candidate ads includes scoring each candidate ad based at least in part on a degree of match between the respective ad and the one or more keywords; and

wherein the one or more ads are determined based at least in part on the scores of the candidate ads.

**41.** The computer-implemented method of claim **39**, further comprising:

processing the image captured by the camera to determine an intent of a user associated with the computing device.

**42.** The computer-implemented method of claim **41**, wherein the plurality of search results are determined based at least in part on the intent of the user.

**43.** The computer-implemented method of claim **39**, wherein the plurality of search results are determined by a search system.

**44.** The computer-implemented method of claim **39**, wherein the one or more ads are determined by a search system.

**45.** One or more non-transitory computer-readable media comprising instructions that when executed by one or more

computing devices cause the one or more computing devices to perform operations comprising:

receiving, via an application executed by a computing device connected to a network, an image captured by a camera of the computing device;

processing the image captured by the camera to determine one or more keywords associated with the image captured by the camera;

determining a plurality of search results responsive to the image capture by the camera based at least in part on the one or more keywords;

determining one or more ads based on at least one of the image captured by the camera or a subset of the plurality of search results; and

providing the plurality of search results and the one or more ads to the computing device.

**46.** The one or more non-transitory computer-readable media of claim **45**, wherein the one or more ads comprise an image ad with embedded information associated with a link and machine executable instructions.

**47.** The one or more non-transitory computer-readable media of claim **45**, wherein providing the plurality of search results and the one or more ads to the computing device comprises:

providing the plurality of search results and the one or more ads for display in a mobile application associated with a search system.

\* \* \* \* \*