



(51) International Patent Classification:

G06F 9/44 (2006.01) G06F 17/30 (2006.01)
G10L 15/22 (2006.01) G06F 3/16 (2006.01)

(21) International Application Number:

PCT/US2015/059355

(22) International Filing Date:

6 November 2015 (06.11.2015)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

62/076,281 6 November 2014 (06.11.2014) US
62/085,852 1 December 2014 (01.12.2014) US
14/700,319 30 April 2015 (30.04.2015) US

(71) Applicant: MICROSOFT TECHNOLOGY LICENSING, LLC [US/US]; Attn: Patent Group Docketing (Bldg. 8/1000), One Microsoft Way, Redmond, Washington 98052-6399 (US).

(72) Inventors: VOGEL, Matthew E.; Microsoft Technology Licensing, LLC, Attn: Patent Group Docketing (Bldg. 8/1000), One Microsoft Way, Redmond, Washington 98052-6399 (US). KRISHNA, Om; Microsoft Technology Licensing, LLC, Attn: Patent Group Docketing (Bldg. 8/1000), One Microsoft Way, Redmond, Washington 98052-6399 (US). YU, Jamie; Microsoft Technology Licensing, LLC, Attn: Patent Group Docketing (Bldg. 8/1000), One Microsoft Way, Redmond, Washington 98052-6399 (US). THOTA, Venkateshwar Rao; Microsoft Technology Licensing, LLC, Attn: Patent Group

Docketing (Bldg. 8/1000), One Microsoft Way, Redmond, Washington 98052-6399 (US). SETO, Julie Christa; Microsoft Technology Licensing, LLC, Attn: Patent Group Docketing (Bldg. 8/1000), One Microsoft Way, Redmond, Washington 98052-6399 (US). YANG, Chenguang; Microsoft Technology Licensing, LLC, Attn: Patent Group Docketing (Bldg. 8/1000), One Microsoft Way, Redmond, Washington 98052-6399 (US). BAI, Annie Y.; Microsoft Technology Licensing, LLC, Attn: Patent Group Docketing (Bldg. 8/1000), One Microsoft Way, Redmond, Washington 98052-6399 (US). STEPANICH, Darron J.; Microsoft Technology Licensing, LLC, Attn: Patent Group Docketing (Bldg. 8/1000), One Microsoft Way, Redmond, Washington 98052-6399 (US). KIKIN GIL, Erez; Microsoft Technology Licensing, LLC, Attn: Patent Group Docketing (Bldg. 8/1000), One Microsoft Way, Redmond, Washington 98052-6399 (US).

(74) Agents: MINHAS, Sandip et al.; Microsoft Corporation, Attn: Patent Group Docketing (Bldg. 8/1000), One Microsoft Way, Redmond, Washington 98052-6399 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, 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, 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.

[Continued on next page]

(54) Title: CONTEXT-BASED ACTIONS IN VOICE USER INTERFACE

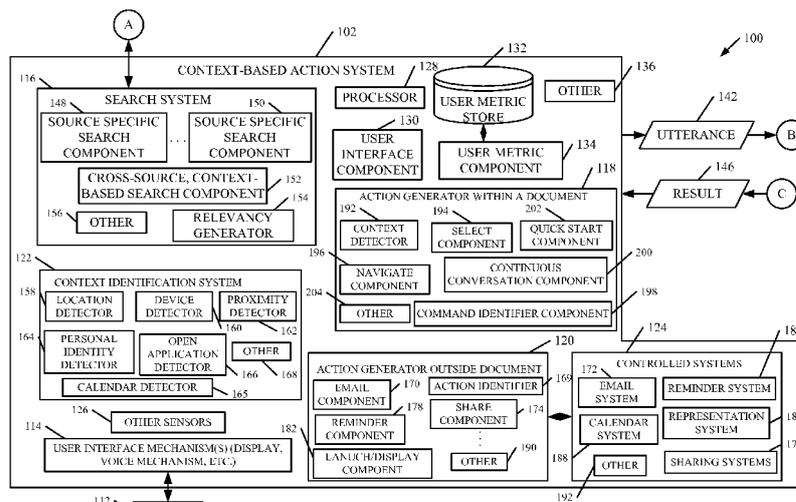


FIG. 1A

(57) Abstract: A computing device receives voice command inputs from a user. The device obtains a language processing result based on the voice command input. The result includes an intent and a set of arguments. The device also obtains a variety of different types of contextual information. An action is identified based on the intent, the arguments, and the contextual information, and the device then suggests the action by displaying a user selectable input mechanism that can be actuated by the user to perform the action. The device can automatically perform the action as well.

WO 2016/073803 A1

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

Declarations under Rule 4.17:

- *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*
- *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

Published:

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

CONTEXT-BASED ACTIONS IN VOICE USER INTERFACE

BACKGROUND

[0001] Computer systems are in wide use. Some such computer systems are deployed
5 on devices that include desktop computers, laptop computers, tablet computers, smart
phones, smart phones with large screens (e.g., phablets) and a variety of other mobile
devices and other computing devices.

[0002] It is not uncommon for a user to operate multiple different computing devices.
For instance, a user may have a desktop computer at work and another computer (such as a
10 laptop computer) at home. The user may also have a mobile device (such as a smartphone)
that the user uses for both work and personal use. The user may have other devices that are
used for both professional and personal work as well.

[0003] In addition, there are many different locations where a user may generate and
store relevant content. For instance, the user may generate and store content local to a device
15 (such as on a smartphone, a laptop computer, a desktop computer, etc.). In a client/server
environment, the user may generate and store information on the server. The user may also
operate in environments where there are multiple different servers, and the user may
generate and store information on any or all of those servers. Similarly, the user may use
one or more cloud-based services, such as cloud-based storage, or other cloud-based
20 services. If so, the user may generate and store content on the cloud as well.

[0004] In order to locate relevant content, a user is often found browsing through many
different files in order to locate the particular document or file that they wish to locate.
Similarly, where a user is in a meeting and wishes to share a document with others in the
meeting, this normally entails the user locating the document himself, or herself, identifying
25 all the different individuals in the meeting, identifying the electronic mail aliases for all of
those individuals, composing an electronic mail message with all of the above information,
then attaching the document and sending it to those individuals.

[0005] These current processes for locating and interacting with content are
cumbersome and time consuming. In addition, especially with relatively small screen
30 devices such as smartphones, they can be error prone.

[0006] The discussion above is merely provided for general background information
and is not intended to be used as an aid in determining the scope of the claimed subject
matter.

SUMMARY

[0007] A computing device receives voice command inputs from a user. The device obtains a language processing result based on the voice command input. The result includes an intent and a set of arguments. The device also obtains a variety of different types of contextual information. An action is identified based on the intent, the arguments, and the contextual information, and the device then suggests the action by displaying a user selectable input mechanism that can be actuated by the user to perform the action. The device can automatically perform the action as well.

[0008] This Summary is provided to introduce a selection of concepts in a simplified form that are further described below in the Detailed Description. This Summary is not intended to identify key features or essential features of the claimed subject matter, nor is it intended to be used as an aid in determining the scope of the claimed subject matter. The claimed subject matter is not limited to implementations that solve any or all disadvantages noted in the background.

BRIEF DESCRIPTION OF THE DRAWINGS

[0009] FIGS. 1A and 1B (collectively referred to herein as FIG. 1) show a block diagram of one example of a context-based command architecture.

[0010] FIG. 2 is a block diagram of one example of a speech processing result.

[0011] FIGS. 3A-3B (collectively referred to herein as FIG. 3) show a flow diagram illustrating one example of the operation of the architecture shown in FIG. 1 in identifying relevant content based on context information.

[0012] FIGS. 3C-3H show examples of user interface displays.

[0013] FIG 4 is a flow diagram illustrating one example of the operation of the architecture shown in FIG. 1 in performing speech processing on a speech input signal.

[0014] FIGS. 4A-4H show examples of user interface displays.

[0015] FIG. 5 is a flow diagram illustrating one example of the operation of the architecture shown in FIG. 1 in processing command inputs within an already-opened document.

[0016] FIGS. 5A-5V show examples of user interface displays.

[0017] FIGS. 6A-6B (collectively referred to herein as FIG. 6) show a flow diagram illustrating one example of the operation of the architecture shown in FIG. 1 in performing a quick start process.

[0018] FIGS. 6C-6H show various examples of user interface displays.

[0019] FIG. 7 is block diagram of one example showing some portions of the architecture shown in FIG. 1 in a remote server environment.

[0020] FIG. 8 is a block diagram showing some portions of the architecture illustrated in FIG. 1 in a client/server environment.

5 [0021] FIG. 9 is a block diagram showing one example of portions of the architecture illustrated in FIG. 1, disposed in a cloud computing architecture.

[0022] FIGS. 10-12 show various examples of mobile devices.

[0023] FIG. 13 is a block diagram of one example of a computing system that can be used in the various architectures in the previous figures.

10

DETAILED DESCRIPTION

[0024] FIGS. 1A and 1B (collectively referred to herein as FIG. 1) show a block diagram of one example of a context-based command architecture 100. Architecture 100 illustratively includes a context-based action system 102 that interacts with a linguistic processing system 104 and a set of sources 106 that illustratively include a plurality of
15 different information sources 108-110. Information sources 108-110 can, themselves, be sources of information such as a store of most recently used documents, electronic mail messages, contacts or people, schedule information (such as calendar information), documents or other information in a collaboration or presentation system, device information, location information, and a wide variety of other information.

20 [0025] It will be noted that all of the items shown in FIG. 1 can be located on a single device or they can be distributed among a variety of different devices or locations. A number of different examples of this are described below with respect to FIGS. 7-9. In the example shown in FIG. 1, context-based action system 102 illustratively receives commands input by user 112 and interprets them in view of a variety of different contextual information. It
25 can then generate a display, such as user interface mechanisms 114, for interaction by user 112. User 112 illustratively interacts with user interface mechanisms 114 in order to control and manipulate system 102. In one example, mechanisms 114 can include user interface display mechanisms, voice mechanisms that are actuated by voice inputs, buttons, keypads, etc., or a wide variety of other mechanisms.

30 [0026] In the example shown in FIG. 1, system 102 illustratively includes search system 116, that searches for information in sources 108-110, action generator 118 that generates actions when the user is already inside of a document, action generator 120 that generates actions when the user is not interacting within a document, context identification system 122, controlled systems 124, a variety of other sensors 126, processor 128, user interface

component 130, user metric store 132, user metric component 134, and it can include other items 136 as well.

[0027] In the example of architecture 100 shown in FIG. 1, linguistic processing system 104 illustratively includes speech recognition system 138 and natural language understanding system 140. Of course, it is also contemplated that either or both of these systems 138, 140 can be part of context-based action system 104 as well, and they are shown separately for the sake of example only.

[0028] Speech recognition system 138 can be any of a wide variety of different types of speech recognition systems that take a speech input (such as utterance 142 that may be generated by a microphone in response to user 112 speaking into the microphone) and generates a contextual representation of the utterance 142. For instance, system 138 can operate using acoustic models and language models in a Hidden Markov process. It can also operate using neural network-based recognition, or other types of recognition.

[0029] Natural language understanding system 140 illustratively receives the textual representation of utterance 142 and generates an interpretation of it. In one example, system 140 includes a plurality of different grammars 143-144 which have rules that are activated based upon the particular words identified in utterance 142. Natural language understanding system 140 generates the interpretation by identifying an intent and various arguments within utterance 142. The intent can be a command desired by user 112 or another type of action or task that user 112 is seeking to perform. The arguments include descriptor information that can identify the particular content on which the action is to be performed. Linguistic processing system 104 illustratively returns the textual representation, the intent and the arguments identified from utterance 142 as linguistic processing result 146. The result is described in greater detail below with respect to FIG. 2.

[0030] Before describing the overall operation of architecture 100 in more detail, a brief description of some of the items in architecture 100, and their corresponding operation, will first be provided. Search system 116 illustratively includes a plurality of different source-specific search components 148-150. In one example, each component 148-150 is configured to search a specific information source 108-110. For instance, where information source 108 is a source of word processing documents, then search component 148 may be configured to search word process documents in source 108. Where information source 110 is an e-mail system that contains e-mail messages and recipients, etc., then search component 150 may be configured to search the source of information 110 that comprises

the e-mail system. Where a source 108-110 is a calendar or contacts directory, a corresponding search component 148-150 can be used.

[0031] In another example, cross-source context-based search component 152 can, either on its own, or by controlling components 148-150, identify (based on the intent and arguments in the result returned by system 104) a plurality of different information sources 5 108-110 that are to be searched for relevant content, and can search those sources 108-110 (or control the corresponding components 148-150 to search them) to generate a set of cross-source, search results. Relevancy generator 154 illustratively obtains a wide variety of different contextual information and generates a relevancy score for each of the items of 10 content returned by search system 116. The relevancy score can be used to generate a ranked list of content (or links to that content) among other things. Search system 116 can include other items 156 as well.

[0032] Context identification system 122 illustratively includes a plurality of different context detectors that each detect a different type of contextual information. For instance, 15 location detector 158 can illustratively be a positioning system (such as a GPS sensor or other location identifying system) that generates an output indicative of a current location of the device that includes context-based action system 102 (and being used by the user). Device detector 160 obtains contextual information about the device upon which it is located (or that user 112 is using), such as whether it is a mobile device, desktop computer, etc. 20 Proximity detector 162 detects a proximity of the device (e.g., the device carried by user 112) to other items, such as other individuals, the user's work location, the user's home location, etc. Personal identity detector 164 illustratively detects the personal identity of user 112 (assuming that the user has opted in or otherwise assented to this type of detection). Calendar detector 165 can detect a current time of day, week, month, fiscal period, year, etc. 25 It can also illustratively identify past, current or future items on the user's calendar. This may include for instance, when the user was in a meeting, the attendees at the meeting, documents presented at the meeting, etc. This information can be provided to relevancy generator 154 or it can be separately generated. Open application detector 166 detects information, such as what applications are currently (or were recently) open on the device 30 being used by user 112, or other devices used by the user. Context identification system 122 can include a wide variety of other detectors 168 that detect a wide variety of other context information as well.

[0033] Action generator 120 illustratively includes an action identifier 169 that identifies actions to take based on the context information and based on the intent and arguments.

Action generator 120 also includes a set of components that can perform the actions, either by themselves or using controlled systems 124, based upon the context information, intent and arguments expressed by user 112 in result 146. For instance, the utterance 142 may be “send the current specification to the attendees of the meeting I am in”. In that case, the intent expressed in utterance 142 will include a number of steps, such as to identify the “current spec”, attach it to an e-mail, (or attaching or otherwise including a link to the document), address the e-mail to the attendees of the “current meeting” that user 112 is in, perhaps adjusting recipient access control rights or permissions, and then send that e-mail to the attendees. Thus, action generator 120 includes items such as e-mail component 170 that can control e-mail system 172 in controlled systems 124. Share component 174 can illustratively perform actions to share content on various sharing systems 176 in controlled systems 124. For instance, it may be that share component 174 posts an item of content to another user’s social network site or work network site, shares it with a work group on a collaboration site, or posts it to a blog, etc. Action generator 120 can also include such items as reminder component 178 that controls a reminder system 180 in controlled systems 124. For instance, it may be that utterance 142 indicates that user 112 wishes to be sent a reminder under certain circumstances. In that case, the intent will be to send a reminder and the arguments will identify the particular circumstances under which the reminder is to be sent. Reminder component 178 thus determines this and performs actions using reminder system 180 to send the desired reminder. Launch/display component 182 can control presentation system 184 or a wide variety of other applications so that it can launch desired applications and generate one or more desired displays, based upon the intent received in result 146. The various components in generator 120 can also control such things as calendar system 188. For instance, the items in action generator 120 can insert an appointment for user 112 or send a meeting request to other meeting invitees on calendar system 188. In addition, where the attendees of a meeting are to be known (such as where user 112 wishes to send an e-mail to all the attendees of a meeting), items in action generator 120 can interact with calendar system 188 to identify the attendees or invitees at the given meeting or to identify other information (such as presentations presented at a given meeting, etc.) and use that information in performing actions. Of course, action generator 120 can include a wide variety of other components 190 that control a wide variety of other controlled systems 192, in order to perform desired actions.

[0034] Action generator 118 illustratively includes a set of components and detectors that can detect a context within an already-opened document and perform actions, or suggest

actions, based upon the context information and based upon various user inputs. For instance, context detector 192 can detect the context of the document that is open. Where the document is a word processing document, it can detect the particular page that is currently being displayed. It may also detect a wide variety of other metadata, such as the different sections in the document, the different authors or commenters on the document, the various revision dates for the document, document formatting, access control permissions, etc. Select component 194 controls interactions by user 112 that indicate that the user wishes to select various items in the open document. Navigation component 196 controls navigation interactions with the document. Command identifier component 198 identifies a subset of commands that have corresponding user input mechanisms that are displayed to the user so the user can quickly execute any of the subset of commands. Continuous conversation component 200 illustratively processes various results 146 or other inputs from user 112, within the context of previous commands that were already identified. Quick start component 202 controls the application that is displaying the document during startup or initial authoring of a document. Of course, action generator 118 can include a wide variety of other items 204 that can be used to perform a wide variety of other actions as well.

[0035] FIG. 2 shows a block diagram of one example of a processing result 146 generated by linguistic processing system 104. It can be seen that result 146 illustratively includes an intent identifier 206 that identifies the intent found in utterance 142. It also includes a set of arguments 208-210 that were recognized in utterance 142. In addition, result 146 can include the textual representation 212 generated from utterance 142 by speech recognition system 138. It can of course include other items 214 as well.

[0036] FIGS. 3A and 3B (collectively referred to herein as FIG. 3) illustrate one example of the operation of system 102 (and relevancy generator 154) in identifying items of content from sources 106 and displaying them for user 112. The relevancy of content can be determined at any desired time, such as when the user first opens or launches system 102, or when the user requests that system 102 refresh the relevancy of the documents being displayed, or automatically based on changes in context information, for example. System 102 thus first determines that it is time to determine which relevant content is to be surfaced for the user. This is indicated by block 216. When it is time to identify relevant content, relevancy generator 154 can access any user preferences or profile information that may be used in the relevancy calculation. This is indicated by block 218. For instance, it may be that the user prefers certain content to be shown first, or otherwise. It also accesses context

identification system 122 and identifies user context. This is indicated by block 220. For instance, personal identity detector 164 can provide the personal identity of user 112 based on login information 222 (again, assuming that the user has assented to this). Location detector 158 and device detector 160 can be used to determine whether user 112 is at home 224 or at work 226, or elsewhere, and what device(s) the user is using. For instance, the user may access a workout log each time he or she is at the gym. Even though the document may not have been accessed in 24 hours, the system senses the user's physical presence at the gym and uses this in the relevance calculation.

[0037] Other user context information 228 (such as usage patterns, etc.) can be identified as well. For instance, it may be close to tax time. The user's previous tax returns may not therefore be surfaced as relevant based on recency or other information alone. However, the system recognizes a pattern in that each year, at this time, the user accesses old tax returns. This pattern can be used in the relevancy calculation. This is only one example.

[0038] Relevancy generator 154 can also identify the items of content that the user has recently interacted with. This is indicated by block 230. For instance, it may be that one of the information sources 106 is a list of the most recently used documents or items of content. Relevancy generator 154 can thus identify the items from that source. It can identify the top n (or a threshold number) of items of content, as indicated by block 232. It can identify the items of content in order of recency as indicated by block 234. It can identify items of content that were most frequently accessed by the user as indicated by block 236 or by other individuals that are related to the user, as indicated by block 238. For instance, the documents can have metadata that defines who accessed documents and when they were accessed. Where the user is accessing information from work, and the user belongs to one or more work groups, or where the user is on one or more teams, or has a supervisor, that information can all be obtained from the various information sources 108-110. Relevancy generator 154 can then identify the various most recently accessed (or most frequently accessed) documents by those other users as well. This is one example only, and the items of content that were most recently interacted with can be identified in other ways as well, and this is indicated by block 240.

[0039] Relevancy generator 154 can also identify the type of interaction. This is indicated by block 242. For instance, it may be that some types of interaction are deemed more important than others. Thus, generator 154 can identify where the interaction was to edit the item of content 244, to simply open it 246, to comment on it 248, to share it 250 or whether the interaction was another type of interaction 252.

[0040] Generator 154 can also identify the frequency of the different types of interaction by user 112 or the other users. This is indicated by block 254.

[0041] In addition, generator 154 can identify what the interactions were, relative to calendar items. This is indicated by block 256. For instance, if the user always accessed a certain item of content in a given recurring meeting, and that meeting is about to occur again, then relevancy generator 154 may determine that particular item of content is particularly relevant at this time. Thus, generator 154 can determine whether the items of content were interacted with before meetings 258, during meetings 260, after meetings 262, relative to various deadlines (such as immediately before certain types of deadlines) 264, during the user's free time 266, or relative to other calendar items 268.

[0042] Generator 154 can also identify the various people involved with user 112, as indicated by block 270. For instance, identification system 122 can identify the various individuals that user 112 frequently messages. This is indicated by block 272. It can identify the various individuals that user 112 is frequently in meetings with as indicated by block 274. It can identify individuals that are closely related to user 112 on an organization relationship chart. This is indicated by block 276. It can identify other users that are in close physical proximity to user 112, as indicated by block 278. It can identify various people that are personally related to user 112 (such as from social network sites in sources 106, etc.) as indicated by block 280. It can identify individuals that user 112 was a co-author with, as indicated by block 282. It can identify the various work groups or social groups or other groups 284 that the user is a member of, and it can identify other people in other ways as well, and this is indicated by block 286.

[0043] Generator 154 can use device detector 160 to identify the particular device that user 112 is currently using. This is indicated by block 288. For instance, it can determine whether the user is using a phone 290, a tablet 292, a laptop computer 294, a desktop computer 296, or another type of device 298. It can also determine or detect patterns on device usage. For example work documents might be more relevant on a desktop computer because that is where the user normally accesses them. There may be other documents (like a gas mileage tracker spreadsheet) that are more relevant on a phone because they are mostly used on the phone.

[0044] In evaluating the relevancy of various items of content, generator 154 can use the various search components 148-152 to identify the content storage location where the items of content are stored. This is indicated by block 300. For instance, they may be stored locally on a local data store as indicated by block 302. They might be stored on a server or

otherwise remotely as indicated by block 304. They may be stored on a particular cloud service 306, or in other locations 308.

[0045] Generator 154 can also use other sensors 126 to identify other sensor inputs 310. For instance, where the other sensors 126 include an accelerometer, generator 154 can determine whether user 112 is currently in motion. This is indicated by block 312. It can also determine the weather where user 112 is located as indicated by block 314, or it can include a wide variety of other sensor inputs 316.

[0046] Relevancy generator 154 then calculates the relevancy of the various items of content that have been identified, based upon all of the information that has been obtained. This is indicated by block 318. For instance, each item of contextual or other information that is obtained relative to an item of content can be assigned a value. The relevancy calculation may be simply aggregating all of those values. This is indicated by block 320.

[0047] It may also be that different items of information are deemed to be more or less important than others. Therefore, the items of information can be weighted based upon their importance, as indicated by block 322. The weighting may change dynamically based on the context information as well. For instance, if the user is at work, then the weights may be set in one way. If the user is at home, the weights may be set differently. This is but one example. There are a wide variety of other types of calculations that can be performed based upon the relevancy and contextual information that was obtained. This is indicated by block 324.

[0048] Relevancy generator 154 can then generate a ranked list of relevant content. This is indicated by block 326. There can also be multiple different types of lists for each different type of content. This is indicated by block 328. For instance, it may be that a certain type of content is a “presentation” in a presentation system (such as a slide presentation system). Thus, the most relevant slide presentations to the user (given the various contextual information) may be ranked in order of relevancy and surfaced for the user in one section of a display. However, it may be that other types of content include electronic mail messages or spreadsheet documents. Those items of content may be ranked based on their relevance and the ranked list for mail messages and the ranked list for spreadsheets can each be surfaced for the user in another, separately scrollable, portion of the user interface display. These are examples only, and the ranked list of relevant content can be generated in other ways as well, and this is indicated by block 330.

[0049] Generator 154 then uses user interface component 130 to surface the relevant content. This is indicated by block 332. For instance, it can surface the top N items of

content as indicated by block 334. It can surface relevant content and divide it out on the display device based on content type. This is indicated by block 336. It can display the content itself as indicated by block 338 or links (e.g., hyperlinks) to the content as indicated by block 340. It can also employ summarization components to generate summaries of the relevant content and display those summaries. This is indicated by block 342. It can surface the relevant content in other ways as well, and this is indicated by block 344.

[0050] System 102 can then receive user interaction with the items of content. This is indicated by block 346. It then performs actions based on the user interactions. This is indicated by block 348, and a number of examples of user interactions are described in greater detail below.

[0051] First, however, a number of examples of user interface displays will be described. FIGS. 3C and 3D show examples of how user interface displays may look on a phone (or other mobile device) and on a tablet computer (or laptop computer). It can be seen that each of the displays illustratively include a content display section 354 and 356, respectively, that either displays content or links to content that has been deemed relevant.

[0052] The displays also include a filter section 358 and 360 that have a set of user selectable filter mechanisms that can be actuated to filter the displayed content. For instance, the currently displayed content is filtered as being “relevant”. However, the user can also actuate the “recent” user input mechanism 362 to filter the displayed content based on how recently it was accessed. The user can actuate the “frequent” user input mechanism 364 to filter the displayed content based on how frequently it is interacted with. The user can actuate the “favorites” user input mechanism 366 to filter the displayed content based upon what the user has deemed a favorite, and the user can actuate the “updated” user input mechanism 368 to update the relevant content.

[0053] FIGS. 3E and 3F are similar to FIGS. 3C and 3D, except that the user is now actuating the “update” user input mechanism 368. FIGS. 3G and 3H show that relevancy generator 154 has received the input indicating that the user wishes to have it recalculate or update the relevancy of the various items of content. This may be, for instance, that the user has switched locations or is in a different meeting now (which is different from when the user last had the relevancy calculated) or otherwise. FIGS. 3G and 3H show that the relevant content that is being displayed has changed, because the user has changed physical locations.

[0054] For instance, in one example, relevancy generator 154 has used location detector 158 to detect that user 112 is now at work, instead of at home. Therefore, the relevancy of

the various items of content has changed so that the work-related content is now more relevant, and is displayed higher in the list, than the personal or home-related content.

[0055] FIG. 4 is a flow diagram illustrating one example of the operation of architecture 100 in processing a speech input (or utterance 142) received by user 112 at system 102. It is assumed for the sake of the present discussion that the user's device is currently displaying relevant content (such as that shown above with respect to FIGS. 3A-3H).

[0056] The user then illustratively provides an actuation input indicating that the user is about to provide a voice command. Alternatively, the device might always be listening and activated by a particular voice command. As one example, the user can touch a touch sensitive display screen or otherwise launch the speech recognition functionality of the system. System 102 then illustratively receives a natural language user voice input (or utterance 142). This is indicated by block 380 in FIG. 4. For instance, this can be received at a microphone on the user's device, or otherwise. System 102 then provides the utterance 142 to linguistic processing system 104. It should be noted that system 104, or the portions 138 and 140 of system 104, can be located locally on the user's device, at a server, in the cloud, in multiple different places, etc. For instance, it may be that the user's device has a relatively small and simple speech recognition system and natural language understanding system, so simple utterances are processed locally. The cloud or server may have a relatively large and comprehensive speech recognition system and natural language processing system so complex utterances are sent to the cloud. In another example, utterances can be processed both places and the result with the highest confidence score can be used. This is indicated by block 384.

[0057] Speech recognition system 138 generates a textual representation of the utterance, as indicated by block 382. Once a textual representation is generated, natural language understanding system 140 identifies an intent 386 in the utterance 142, based upon the textual representation. It also identifies arguments 388 from the textual representation as well. It can identify other information as well, and this is indicated by block 390.

[0058] The intent illustratively corresponds to an action that the user wishes to perform. For instance, the user may utter a phrase such as "share this document with Joe." In that case, natural language understanding system 140 will identify the word "share" as an action or command that the user wishes the system to perform. As an example, the word "share" may activate one or more rules in the various grammars 143-144. Those rules may activate other rules, each of which have a set of arguments that are to be matched before the rule fires. When the textual representation matches a given rule, the intent for that rule and the

various arguments for the rule are output as part of result 146, along with the textual representation.

[0059] User interface component 130 then displays the textual representation to the user, as indicated by block 392. Action identifier 169 in action generator 120 identifies actions to
5 take based upon the intent and it can also do this based on the context information and the arguments. This is indicated by block 394. For instance, it can access a set of rules 143-144 as well. The rules illustratively map the intent to a given action or set of actions. This is indicated by block 396.

[0060] The rules can be nested or arranged in a hierarchal or dependency structure in order
10 to accomplish a task that requires multiple different commands or actions. Or, they can reference one another or interact in other ways. This is indicated by block 398. For instance, sending a document to a set of attendees at a meeting requires multiple different actions. First, it requires the document to be located. It also requires a set of attendees to be located. It then requires an e-mail to be generated with those attendees as recipients. It then requires
15 that the located document be attached to the e-mail and that the e-mail be sent. All of these actions or commands can be identified through a sequence of rules that are active based upon the intent expressed in the utterance. The actions can be identified in other ways as well, and this is indicated by block 400.

[0061] Search system 116 then identifies the item of content on which the actions are to be
20 performed based upon the arguments and context information. This is indicated by block 402. For instance, where the intent is to send an e-mail with attachments, the arguments will illustratively identify the attachment to the e-mail and the recipients of the e-mail. The arguments may explicitly identify the attachment and attendees or they may do so in an implicit way (such as where the attachment is identified as “the spec I revised this morning”
25 or where the recipients are identified as “the attendees of this meeting”). In that case, search system 116 illustratively searches the various sources of information 106 to identify the attachment to be sent, as well as the recipients to the e-mail. As it searches the various sources of information 106, it will illustratively identify the most relevant search results (such as the most likely recipients, the most likely attachments, etc.). In doing so, it can use
30 the relevancy generator to generate a relevancy score for each identified item, or it can generate relevancy or estimate the most likely recipients and attachment in other ways as well. Considering the relevancy is indicated by block 404. The item of content on which the action is to be performed can be identified in other ways 406 as well.

[0062] Once action generator 120 has identified the action to be taken, and has used search system 116 to identify the items of content needed to perform the action, it illustratively uses one of the components in action generator 120 to perform the action as indicated by block 408. It can do so by performing one or more actions in one of the controlled systems 124, such as by attaching a document to an e-mail in an email system, sending the e-mail and simultaneously sharing the document to a group or social network site, etc. Performing multiple actions is indicated by block 410. The actions can be performed in other ways as well, and this is indicated by block 412.

[0063] FIGS. 4A-4H show various examples of user interface displays that can be generated in performing commands based on a voice input. For instance, FIGS. 4A and 4B show examples of how user interface displays may be on the phone 350 and tablet computer 352, respectively. In the example illustrated, the user has provided a voice input of “Find me the Contoso sales deck I sent last week for Rob Young I think I updated it last week from home.” In doing so, the user first illustratively actuates a user input mechanism (such as mechanism 414) to indicate to system 102 that the user is about to speak an utterance. The utterance was captured and sent to linguistic processing system 104 where speech recognition system 138 generates a textual representation of it. It is displayed on the corresponding user interface display, such as at block 416 or 418, respectively.

[0064] FIGS. 4C and 4D are similar to FIGS. 4A and 4B, except that it is seen that search system 116 has identified a number of different relevant items of content, that are relevant to the user’s utterance. In order to do this, action generator 120 first received an indication that the intent in the utterance was to “find” a document. It then used search system 116 to locate documents relevant to the user’s query. In doing so, it used arguments such as “Contoso sales deck”, “last week”, “Rob Young”, and “home” to identify relevant documents. The displays shown in FIGS. 4C and 4D indicate that the displayed items of content are what was found in response to the user’s query. They are illustratively displayed in order of relevance to the user’s query, as determined either by relevancy generator 154, or in other ways. Display component 182 then generates the user interface displays on phone 350 and/or tablet 352 (whichever the user is using) and the relevant content is displayed.

[0065] FIGS. 4E and 4F are similar to FIGS. 4C and 4D, except that they show that the user is now actuating a “share” user input mechanism 420 to share the document. This can be done, for instance, by touching the user input mechanism on a touch sensitive screen. Share component 174 then uses search system 116 to identify recommendations as to which people

or groups the document should be shared with. This can be done based upon the contextual information of the user, the document, the other users or other context information as discussed above. Share component 174 then uses user interface component 130 to display the share recommendations to the user. This is illustrated in FIGS. 4G-4H.

5 [0066] In the example shown in FIGS. 4G and 4H, the share recommendations 422 are displayed to the user for selection. In the example shown, the share recommendations were identified based upon a proximity of user 112 to the various people that were analyzed, and based upon other contextual information. This is but one example. If the share recommendations 422 are accurate, the user simply actuates the send user input mechanism
10 424 and the selected item of content is shared with the people identified in the share recommendation.

[0067] Some of the “intents” in the utterance described thus far represent commands or actions that are taken outside the context of any given document. For instance, the commands can include such things as open, download, print, share, create, send, schedule
15 an appointment, delete, post, find, etc. These actions can be performed on documents or other objects across multiple different data sources (e.g., word processing documents, presentations, emails, etc.) and across different devices (e.g., the command “open, on my desktop, the documents I received from Nicole yesterday.” may be input by the user on a mobile device, but the document may be opened on the user’s desktop).

20 [0068] It should also be noted that system 102 can perform actions on a document, when the user has a document open. For instance, the system can perform actions within the document (such as navigation actions, formatting actions, selecting actions, etc.). FIG. 5 is a flow diagram illustrating the operation of system 102 in doing this. It is first assumed that the user has a document open. This is indicated by block 440. The document may, for
25 instance, be a word processing document 442, a spreadsheet 444, a presentation document (such as a slide presentation) 446, an electronic mail message 448, an image 450, or another document 452.

[0069] Context detector 192 illustratively detects the document context. This is indicated by block 454. For instance, it may detect the page that the document is scrolled to, the
30 formatting of the document, or a wide variety of other contextual information.

[0070] Action generator 118 then accesses user metric component 134 to identify a user learning metric corresponding to user 112. This is indicated by block 456. By way of example, user metric component 134 can illustratively generate a user metric indicative of how advanced or experienced user 112 is in using the particular application that the user is

using, in using the particular device the user is using, in using context-based action system 102, or in using other items. As one example, user metric component may keep track of how many times the user has used formatting features in a given application. It may keep track of which types of features the user has used, how often the user has used system 104,
5 or a wide variety of other things. This can serve as a metric (or a metric can be calculated based on this information) indicative of the user's experience. As the user gains experience in using a given application or system 102, it maybe that the system need not generate as many help prompts for the user. However, if the user is a new user, then it may be that the system wishes to generate more help prompts or to display them more frequently. In any
10 case, generator 118 accesses the user learning metric for the current user 112 as an indication of how experienced or knowledgeable the user is in interacting with the document.

[0071] Command identifier component 198 then identifies a context-specific subset of commands and a display frequency based upon the context and the user learning metric. This is indicated by block 458. For instance, it may be that a user is simply viewing a
15 document. In that case, command identifier component 198 may predict that the user may wish to perform one of a first subset of commands or actions. In another example, the user may have selected either a word, phrase, paragraph, object or image in the document. In that case, command identifier component 198 illustratively predicts that the user may desire to perform any of a different subset of commands.

[0072] In identifying the subset of commands, component 198 may access a set of command prediction rules, or a command prediction component that predicts the various commands that the user may wish to use. In any case, component 198 identifies a subset of commands that the user may wish to take, based upon the context of the application, the user's context,
20 and the user's experience level (represented by the learning metric). Component 198 may also consider user-specific usage or preference data. This is indicated by block 460. The component 198 will also modify the subset of commands that are identified based upon changes in the metric, context data, etc. This is indicated by block 462. The subset of
25 commands can be identified in other ways as indicated by block 464.

[0073] Command identifier component 198 then uses user interface component 130 to
30 display mechanisms for the identified subset of commands. They are displayed at a frequency based upon the user learning metric. The user can then interact with the display mechanisms to perform the subset of commands. Displaying the identified subset of commands for user interaction is indicated by block 466. An example may be helpful.

[0074] FIGS. 5A-5D show various examples of user interface displays and command display mechanisms. In FIG. 5A, for instance, it can be seen that the user is simply viewing a document in display 468. Therefore, command identifier component 198 identifies a first set of commands that the user may wish to execute, given the fact that the user is simply
5 viewing the document. One example of a subset of commands includes a “go to” command, a “find” command, a formatting “themes” command, among others. It then generates a set of display elements 470, with one display element corresponding to each of the identified subset of commands. Display elements 470 are illustratively user actuatable display elements that can be actuated by the user (such as by tapping on them) to perform the
10 corresponding action or command. Therefore, for instance, if the user actuates the “go to” display element, navigate component 196 will navigate the user to an identified section of the document. If the user actuates the “themes” display element, then action generator 118 will navigate the user through a formatting user experience that allows the user to format the document according to various themes. These are examples only.

[0075] FIG. 5B shows that the user has selected the words “Redmond, Washington”. Thus, command identifier component 198 has identified a second subset of commands and displayed a set of display elements 472 corresponding to those commands. Because the user has selected a set of words, command identifier component 198 has predicted that the user is likely to want to perform different commands than when the user had not selected any
20 words (as shown in FIG. 5A). Thus, the display elements 472 correspond to different commands than the display elements 470.

[0076] It can be seen in FIG. 5C that the user has now selected an entire paragraph. Therefore, command identifier component 198 has predicted that the user may wish to perform any of yet another subset of commands. Component 198 thus displays another set
25 of display elements 474 corresponding to those commands.

[0077] FIG. 5D shows that the user has now selected an image in the document. Based on this fact, command identifier component 198 predicts that the user is likely to perform one of yet a different subset of commands, and it has displayed a different set of display elements 476 corresponding to those commands. The user can interact with one of the display
30 elements, and action generator 118 illustratively uses one of its components to perform the corresponding action or command. This is indicated by blocks 478 and 480 in FIG. 5.

[0078] When the user does this, the user may then provide further interaction inputs. In that case, in one example, continuous conversation component 200 in action generator 118 interprets the subsequent inputs in the context of the previous command. This is indicated

by block 482. Of course, the commands can be voice commands 484, or other commands 486. FIGS. 5E-5K illustrate one example of how the system interprets commands in view of previous commands.

[0079] It can be seen in FIG. 5E that the user has selected all of the text in the document being displayed. The user then actuates user input mechanism 484 indicating that the user is about to provide a voice input. The user then speaks the input “make the line spacing wider”. FIG. 5F shows that the textual representation of that utterance is displayed at 488. In response, linguistic processing system 104 identifies the intent as changing the line spacing and provides that along with result 146, to continuous conversation component 200. Command identifier component 198 uses this as context information and identifies a subset of commands that are related to changing line spacing and displays a set of display elements 486 corresponding to those commands.

[0080] The user taps a display element 486 indicating that the user wishes to make the line spacing wider. Action generator 118 then makes the line spacing wider, as shown in FIG. 5G. The user then provides another voice input “more”. This is provided to system 104 where the word is returned as result 146. Component 200 identifies that this is a continued instruction to further increase the line spacing, based upon the previous instruction to increase the line spacing. Thus, it controls action generator 118 to increase the line spacing even more. This is shown in FIG. 5H. FIG. 5I shows that command identifier component 198 maintains two display elements 490 that have to do with line spacing. This is because it has been interpreted that the user is still attempting to modify the line spacing. Thus, as shown in FIG. 5J, the user can simply tap one of those user input mechanisms, and action generator 118 performs the corresponding action by navigating the user through another user experience that allows the user to manually set the line spacing to a given level, as indicated by FIG. 5K.

[0081] Returning to FIG. 5, the user may provide more interactions, as indicated by block 492. If this is the case, processing reverts to block 454 where the document context is again detected, the user learning metric is accessed, etc.

[0082] At block 478, it may be that the user does not ever actuate one of the display elements that is displayed by command identifier component 198 (such as 470 in FIG. 5A). In that case, the system detects whether the user provides any other interaction inputs within the document. This is indicated by block 494. If not, then the display elements corresponding to the identified commands are eventually removed from the display (e.g., they disappear after a certain amount of time). This is indicated by block 496. For instance, if the user is

simply viewing a document as shown in FIG. 5A, and the user does not actuate any of the display elements 470, then after a predetermined period of time, they are removed. The predetermined period of time may vary based upon the user learning metric for this particular user. For instance, if the user is a fairly experienced user of the system, then the system will determine that it need not display the display elements for very long, because the user already knows how to use them. However, if the user is a relatively inexperienced or new user, then the display elements may be displayed for a longer period of time, so that the user can review them, try them, consider various options, etc. In addition, it should be noted that, in one example, the user can access an entire set of commands at any time. For instance, the user can do this by navigating into commanding menus, so the full set of commands may be hidden or buried, but not inaccessible.

[0083] If the user does perform another type of interaction (other than actuating one of the display elements) then action generator 118 performs any corresponding action. This is indicated by block 498. Processing then again reverts to block 454. In doing so, the user interaction will be part of the document context that is considered by command identifier component 198 in predicting commands that the user will likely use, and displaying the corresponding display elements. FIGS. 5L-5O illustrate one example of this.

[0084] FIG. 5L is similar to FIG. 5A, except that it can be seen that the user is now providing a touch gesture to scroll the document being displayed. Given that user interaction as context information, command identifier component 198 predicts that the user is likely trying to find a particular place in the document. Therefore, command identifier component 198 displays a set of display elements 501 corresponding to commands that would allow the user to do this. For instance, one of the display elements is a “go to page” display element. When the user actuates this, as indicated in FIG. 5N, action generator 118 displays a menu 503 that allows the user to select a particular page of the document that the user wishes to scroll to. Navigation component 196 then navigates the user to that portion of the document. Again, this is but one example, and there are a wide variety of other user actions or interactions that can be used as context information by component 198 to predict likely commands that the user wishes to use.

[0085] FIGS. 5P-5V show one example of user interactions where the user is selecting different items within a document. Selecting items on a relatively small screen device can be very cumbersome. The user often inadvertently selects the wrong items, or inadvertently selects multiple items when only one is intended. In addition, the user often has to scroll through long documents selecting various items of the same type. FIG. 5P shows that the

user is simply viewing a document that has pictures in it. The user actuates a user input mechanism 535 that indicates to the system that the user is about to provide a voice command. FIG. 5Q shows that the user has provided a voice command “I want to format all my pictures”. FIG. 5R shows that the system is now displaying the textual representation of that input. FIG. 5S shows that select component 194 has received the result 146 indicating that an intent is that the user wishes to select all of the pictures in the opened document, and format them. Thus, select component 194 automatically locates and selects all of the pictures in the displayed document. Command identifier component 198 also identifies a subset of commands that have to do with formatting pictures and displays display elements corresponding to them. FIG. 5S shows that the user is tapping one of those display elements. The same is shown in FIG. 5T. FIG. 5U shows that the user has selected a different picture to be substituted for the selected pictures, and FIG. 5V shows that the newly selected picture has been inserted.

[0086] FIGS. 6A and 6B (collectively referred to herein as FIG. 6) show a flow diagram illustrating one example of the operation of quick start component 202 in allowing a user to quickly begin authoring a document. It is believed that some users have a difficult time in beginning to author a document, in that they are unsure of how to set up the formatting, etc. Thus, component 202 assists the user in doing this. It first receives a user input opening a new document. This is indicated by block 505. Again, the document can be a word processing document 507, a spreadsheet document 509, a presentation document 511, or a wide variety of other documents 513.

[0087] The action generator 120 uses launch/display component 182 to launch the relevant application and to display a blank canvas for the document. This is indicated by block 515. Action generator 118 then receives a user input requesting formatting that is similar to a previous document. This is indicated by block 517. Again, this can be a voice command 519, or another natural language command 521. The system then performs natural language understanding on the input to identify a document, identify its formatting, and apply that formatting to the present document.

[0088] Performing natural language understanding on the input is indicated by block 519. In one example, quick start component 202 displays a plurality of different, relevant documents and allow the user to select the particular document from which the formatting is to be identified and applied to the new document. Determining whether the user is to have more than one option is indicated by block 521. This determination can be made based on user preferences, based on administrative set up, or in other ways.

[0089] Where the user is not to be presented with a plurality of different options, then the document identified by the user in the user's natural language input is retrieved, and the format settings are identified. This is indicated by block 523. Processing then skips to block 525 where the identified formatting is applied to the displayed, new document.

5 [0090] However, if, at block 521, a plurality of different options are to be provided (or if the natural language processing result is ambiguous and multiple documents are identified) then the system identifies the previous document and other similar documents as indicated by block 527. The documents can be similar because they are recent documents 529, they can be of a similar document type 531, or they can be similar in other ways 533.

10 [0091] For each of the identified documents, quick start component 202 generates and displays a corresponding display element. The display element illustratively identifies the document and displays some content of the document so that the user can have some sense of the formatting that was applied to the underlying document. Generating and displaying a selectable display element corresponding to the identified documents is indicated by block
15 535.

[0092] Component 202 then receives user selection of one of those display elements as indicated by block 537. It identifies the formatting for the corresponding document as indicated by block 539. This can be done by parsing the corresponding document as indicated by block 541, by accessing formatting metadata corresponding to the document as
20 indicated by block 543, or in other ways 545. Then, the identified formatting is applied to the displayed, new document at block 525.

[0093] Quick start component 202 can also display a user input mechanism that can be actuated by the user to allow the user to see the particular formatting that has been applied. When this happens, component 202 exposes the formatting information on the new,
25 displayed document, so the user can see how it is formatted. This is indicated by block 527. The system then goes on to process other user inputs, such as further formatting inputs, authoring inputs, or other inputs. This is indicated by block 529.

[0094] FIGS. 6C-6H show various examples of this. It can be seen in FIG. 6C that the user has opened a document and the corresponding application has displayed a blank canvas.
30 The user then actuates a user input mechanism so that the user can provide a voice command. In FIG. 6D, the user has provided the voice command "set this up like my last term paper". Linguistic processing system 104 generates the textual representation and the natural language understanding for that utterance. The intent is to find the user's last term paper, extract the formatting from that paper and apply it to the newly opened document.

[0095] FIG. 6E shows that action generator 118 has used search system 116 to identify a plurality of relevant documents that correspond to the user's query. It displays a display element (such as a thumbnail or other display element) corresponding to each of the relevant documents. FIG. 6E shows that the user is selecting one of them. FIG. 6F shows that the user is providing a user input to apply that formatting to the newly opened document. FIG. 6G shows a user input mechanism 531 that can be actuated by the user to see the formatting of the different sections of the document. When the user actuates user input mechanism 531, the formatting information is displayed. FIG. 6H shows that a plurality of different sections can be displayed separately, with their own format information.

5 [0096] FIGS. 7-9 show various examples indicating that the components or items in the architecture 100 shown in FIG. 1 can be located in different places. FIG. 7 shows that a variety of the different items are located on a mobile device 601 that is connected to a search service that implements search system 116 over a network 603. Network 603 can be a local area network, a wide area network, a cellular network, a near field communication network, or a wide variety of other networks.

10 [0097] FIG. 8 shows that the items are disposed in a client/server architecture with context-based action system 102 deployed on a server 604, along with linguistic processing system 104. User 112 accesses system 102 on server 604 through a client device 606 that generates the user interface mechanisms 114 for interaction by user 112.

15 [0098] The present discussion has mentioned processors and servers. In one embodiment, the processors and servers include computer processors with associated memory and timing circuitry, not separately shown. They are functional parts of the systems or devices to which they belong and are activated by, and facilitate the functionality of the other components or items in those systems.

20 [0099] Also, a number of user interface displays have been discussed. They can take a wide variety of different forms and can have a wide variety of different user actuatable input mechanisms disposed thereon. For instance, the user actuatable input mechanisms can be text boxes, check boxes, icons, links, drop-down menus, search boxes, etc. They can also be actuated in a wide variety of different ways. For instance, they can be actuated using a point and click device (such as a track ball or mouse). They can be actuated using hardware buttons, switches, a joystick or keyboard, thumb switches or thumb pads, etc. They can also be actuated using a virtual keyboard or other virtual actuators. In addition, where the screen on which they are displayed is a touch sensitive screen, they can be actuated using touch

25

30

gestures. Also, where the device that displays them has speech recognition components, they can be actuated using speech commands.

[00100] A number of data stores have also been discussed. It will be noted they can each be broken into multiple data stores. All can be local to the systems accessing them, all
5 can be remote, or some can be local while others are remote. All of these configurations are contemplated herein.

[00101] Also, the figures show a number of blocks with functionality ascribed to each block. It will be noted that fewer blocks can be used so the functionality is performed by fewer components. Also, more blocks can be used with the functionality distributed among
10 more components.

[00102] FIG. 9 is a block diagram of architecture 100, shown in FIG. 1, except that its elements are disposed in a cloud computing architecture 500. Cloud computing provides computation, software, data access, and storage services that do not require end-user knowledge of the physical location or configuration of the system that delivers the services.
15 In various embodiments, cloud computing delivers the services over a wide area network, such as the internet, using appropriate protocols. For instance, cloud computing providers deliver applications over a wide area network and they can be accessed through a web browser or any other computing component. Software or components of architecture 100 as well as the corresponding data, can be stored on servers at a remote location. The
20 computing resources in a cloud computing environment can be consolidated at a remote data center location or they can be dispersed. Cloud computing infrastructures can deliver services through shared data centers, even though they appear as a single point of access for the user. Thus, the components and functions described herein can be provided from a service provider at a remote location using a cloud computing architecture. Alternatively,
25 they can be provided from a conventional server, or they can be installed on client devices directly, or in other ways.

[00103] The description is intended to include both public cloud computing and private cloud computing. Cloud computing (both public and private) provides substantially seamless pooling of resources, as well as a reduced need to manage and configure
30 underlying hardware infrastructure.

[00104] A public cloud is managed by a vendor and typically supports multiple consumers using the same infrastructure. Also, a public cloud, as opposed to a private cloud, can free up the end users from managing the hardware. A private cloud may be managed by the organization itself and the infrastructure is typically not shared with other

organizations. The organization still maintains the hardware to some extent, such as installations and repairs, etc.

[00105] In the example shown in FIG. 9, some items are similar to those shown in FIG. 1 and they are similarly numbered. FIG. 9 specifically shows that various items in
5 FIG. 1 can be located in cloud 502 (which can be public, private, or a combination where portions are public while others are private). Therefore, user 112 uses a user device 504 to access those items through cloud 502.

[00106] FIG. 9 also depicts another embodiment of a cloud architecture. FIG. 9 shows that it is also contemplated that some elements of architecture 100 can be disposed in cloud
10 502 while others are not. By way of example, data store 132 and sources 106 can be disposed outside of cloud 502, and accessed through cloud 502. In another example, system 102 or search system 116 or other portions can also be outside of cloud 502. Regardless of where they are located, they can be accessed directly by device 504, through a network (either a wide area network or a local area network), they can be hosted at a remote site by
15 a service, or they can be provided as a service through a cloud or accessed by a connection service that resides in the cloud. All of these architectures are contemplated herein.

[00107] It will also be noted that architecture 100, or portions of it, can be disposed on a wide variety of different devices. Some of those devices include servers, desktop computers, laptop computers, tablet computers, or other mobile devices, such as palm top
20 computers, cell phones, smart phones, multimedia players, personal digital assistants, etc.

[00108] FIG. 10 is a simplified block diagram of one illustrative embodiment of a handheld or mobile computing device that can be used as a user's or client's hand held device 16, in which the present system (or parts of it) can be deployed. FIGS. 11-12 are examples of handheld or mobile devices.

[00109] FIG. 10 provides a general block diagram of the components of a client
25 device 16 that can run components of architecture 100 or that interacts with architecture 100, or both. Not all elements of FIG. 1 are shown in device 16 although it is contemplated that they could be disposed on device 16 in the illustrated components, or added to them. In the device 16, a communications link 13 is provided that allows the handheld device to
30 communicate with other computing devices and under some examples provides a channel for receiving information automatically, such as by scanning. Examples of communications link 13 include an infrared port, a serial/USB port, a cable network port such as an Ethernet port, and a wireless network port allowing communication though one or more communication protocols including General Packet Radio Service (GPRS), LTE, HSPA,

HSPA+ and other 3G and 4G radio protocols, 1Xrtt, and Short Message Service, which are wireless services used to provide cellular access to a network, as well as Wi-Fi protocols, and Bluetooth protocol, which provide local wireless connections to networks.

5 [00110] Under other examples, applications or systems are received on a removable Secure Digital (SD) card that is connected to a SD card interface 15. SD card interface 15 and communication links 13 communicate with a processor 17 (which can also embody processor 128 or other processors from FIG. 1) along a bus 19 that is also connected to memory 21 and input/output (I/O) components 23, as well as clock 25 and location system 27.

10 [00111] I/O components 23, in one embodiment, are provided to facilitate input and output operations. I/O components 23 for various embodiments of the device 16 can include input components such as buttons, touch sensors, multi-touch sensors, optical or video sensors, voice sensors, touch screens, proximity sensors, microphones, tilt sensors, and gravity switches and output components such as a display device, a speaker, and or a printer 15 port. Other I/O components 23 can be used as well.

[00112] Clock 25 illustratively comprises a real time clock component that outputs a time and date. It can also, illustratively, provide timing functions for processor 17.

20 [00113] Location system 27 illustratively includes a component that outputs a current geographical location of device 16. This can include, for instance, a global positioning system (GPS) receiver, a LORAN system, a dead reckoning system, a cellular triangulation system, or other positioning system. It can also include, for example, mapping software or navigation software that generates desired maps, navigation routes and other geographic functions.

25 [00114] Memory 21 stores operating system 29, network settings 31, applications 33, application configuration settings 35, data store 37, communication drivers 39, and communication configuration settings 41. Memory 21 can include all types of tangible volatile and non-volatile computer-readable memory devices. It can also include computer storage media (described below). Memory 21 stores computer readable instructions that, when executed by processor 17, cause the processor to perform computer-implemented 30 steps or functions according to the instructions. Application 154 or the items in data store 156, for example, can reside in memory 21. Similarly, device 16 can have a client business system 24 which can run various business applications or embody parts or all of tenant 104. Processor 17 can be activated by other components to facilitate their functionality as well.

[00115] Examples of the network settings 31 include things such as proxy information, Internet connection information, and mappings. Application configuration settings 35 include settings that tailor the application for a specific enterprise or user. Communication configuration settings 41 provide parameters for communicating with other
5 computers and include items such as GPRS parameters, SMS parameters, connection user names and passwords.

[00116] Applications 33 can be applications that have previously been stored on the device 16 or applications that are installed during use, although these can be part of operating system 29, or hosted external to device 16, as well.

10 [00117] FIG. 11 shows one example in which device 16 is a tablet computer 600. In FIG. 6, computer 600 is shown with user interface display screen 602. Screen 602 can be a touch screen (so touch gestures from a user's finger can be used to interact with the application) or a pen-enabled interface that receives inputs from a pen or stylus. It can also use an on-screen virtual keyboard. Of course, it might also be attached to a keyboard or
15 other user input device through a suitable attachment mechanism, such as a wireless link or USB port, for instance. Computer 600 can also illustratively receive voice inputs as well.

[00118] Additional examples of devices 16 can be used as well. Device 16 can be a feature phone, smart phone or mobile phone. The phone includes a set of keypads for dialing phone numbers, a display capable of displaying images including application images, icons,
20 web pages, photographs, and video, and control buttons for selecting items shown on the display. The phone can include an antenna for receiving cellular phone signals such as General Packet Radio Service (GPRS) and 1Xrtt, and Short Message Service (SMS) signals. In some examples, the phone also includes a Secure Digital (SD) card slot that accepts a SD card.

25 [00119] The mobile device can also be a personal digital assistant or a multimedia player or a tablet computing device, etc. (hereinafter referred to as a PDA). The PDA can include an inductive screen that senses the position of a stylus (or other pointers, such as a user's finger) when the stylus is positioned over the screen. This allows the user to select, highlight, and move items on the screen as well as draw and write. The PDA can also
30 include a number of user input keys or buttons which allow the user to scroll through menu options or other display options which are displayed on the display, and allow the user to change applications or select user input functions, without contacting the display. The PDA can include an internal antenna and an infrared transmitter/receiver that allow for wireless communication with other computers as well as connection ports that allow for hardware

connections to other computing devices. Such hardware connections are typically made through a cradle that connects to the other computer through a serial or USB port. As such, these connections are non-network connections.

[00120] FIG. 12 also shows that the phone is a smart phone 71. Smart phone 71 has
5 a touch sensitive display 73 that displays icons or tiles or other user input mechanisms 75. Mechanisms 75 can be used by a user to run applications, make calls, perform data transfer operations, etc. In general, smart phone 71 is built on a mobile operating system and offers more advanced computing capability and connectivity than a feature phone.

[00121] Note that other forms of the devices 16 are possible.

10 [00122] FIG. 13 is one embodiment of a computing environment in which architecture 100, or parts of it, (for example) can be deployed. With reference to FIG. 13, an exemplary system for implementing some embodiments includes a general-purpose computing device in the form of a computer 810. Components of computer 810 may include, but are not limited to, a processing unit 820 (which can comprise processor 128 or
15 those in device 504 or other devices), a system memory 830, and a system bus 821 that couples various system components including the system memory to the processing unit 820. The system bus 821 may be any of several types of bus structures including a memory bus or memory controller, a peripheral bus, and a local bus using any of a variety of bus architectures. By way of example, and not limitation, such architectures include Industry
20 Standard Architecture (ISA) bus, Micro Channel Architecture (MCA) bus, Enhanced ISA (EISA) bus, Video Electronics Standards Association (VESA) local bus, and Peripheral Component Interconnect (PCI) bus also known as Mezzanine bus. Memory and programs described with respect to FIG. 1 can be deployed in corresponding portions of FIG. 13.

[00123] Computer 810 typically includes a variety of computer readable media.
25 Computer readable media can be any available media that can be accessed by computer 810 and includes both volatile and nonvolatile media, removable and non-removable media. By way of example, and not limitation, computer readable media may comprise computer storage media and communication media. Computer storage media is different from, and does not include, a modulated data signal or carrier wave. It includes hardware storage
30 media including both volatile and nonvolatile, removable and non-removable media implemented in any method or technology for storage of information such as computer readable instructions, data structures, program modules or other data. Computer storage media includes, but is not limited to, RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM, digital versatile disks (DVD) or other optical disk storage,

magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other medium which can be used to store the desired information and which can be accessed by computer 810. Communication media typically embodies computer readable instructions, data structures, program modules or other data in a transport mechanism and includes any information delivery media. The term “modulated data signal” means a signal that has one or more of its characteristics set or changed in such a manner as to encode information in the signal. By way of example, and not limitation, communication media includes wired media such as a wired network or direct-wired connection, and wireless media such as acoustic, RF, infrared and other wireless media. Combinations of any of the above should also be included within the scope of computer readable media.

[00124] The system memory 830 includes computer storage media in the form of volatile and/or nonvolatile memory such as read only memory (ROM) 831 and random access memory (RAM) 832. A basic input/output system 833 (BIOS), containing the basic routines that help to transfer information between elements within computer 810, such as during start-up, is typically stored in ROM 831. RAM 832 typically contains data and/or program modules that are immediately accessible to and/or presently being operated on by processing unit 820. By way of example, and not limitation, FIG. 13 illustrates operating system 834, application programs 835, other program modules 836, and program data 837.

[00125] The computer 810 may also include other removable/non-removable volatile/nonvolatile computer storage media. By way of example only, FIG. 13 illustrates a hard disk drive 841 that reads from or writes to non-removable, nonvolatile magnetic media, and an optical disk drive 855 that reads from or writes to a removable, nonvolatile optical disk 856 such as a CD ROM or other optical media. Other removable/non-removable, volatile/nonvolatile computer storage media that can be used in the exemplary operating environment include, but are not limited to, magnetic tape cassettes, flash memory cards, digital versatile disks, digital video tape, solid state RAM, solid state ROM, and the like. The hard disk drive 841 is typically connected to the system bus 821 through a non-removable memory interface such as interface 840, and optical disk drive 855 are typically connected to the system bus 821 by a removable memory interface, such as interface 850.

[00126] Alternatively, or in addition, the functionality described herein can be performed, at least in part, by one or more hardware logic components. For example, and without limitation, illustrative types of hardware logic components that can be used include Field-programmable Gate Arrays (FPGAs), Program-specific Integrated Circuits (ASICs),

Program-specific Standard Products (ASSPs), System-on-a-chip systems (SOCs), Complex Programmable Logic Devices (CPLDs), etc.

[00127] The drives and their associated computer storage media discussed above and illustrated in FIG. 13, provide storage of computer readable instructions, data structures, program modules and other data for the computer 810. In FIG. 13, for example, hard disk drive 841 is illustrated as storing operating system 844, application programs 845, other program modules 846, and program data 847. Note that these components can either be the same as or different from operating system 834, application programs 835, other program modules 836, and program data 837. Operating system 844, application programs 845, other program modules 846, and program data 847 are given different numbers here to illustrate that, at a minimum, they are different copies.

[00128] A user may enter commands and information into the computer 810 through input devices such as a keyboard 862, a microphone 863, and a pointing device 861, such as a mouse, trackball or touch pad. Other input devices (not shown) may include a joystick, game pad, satellite dish, scanner, or the like. These and other input devices are often connected to the processing unit 820 through a user input interface 860 that is coupled to the system bus, but may be connected by other interface and bus structures, such as a parallel port, game port or a universal serial bus (USB). A visual display 891 or other type of display device is also connected to the system bus 821 via an interface, such as a video interface 890. In addition to the monitor, computers may also include other peripheral output devices such as speakers 897 and printer 896, which may be connected through an output peripheral interface 895.

[00129] The computer 810 is operated in a networked environment using logical connections to one or more remote computers, such as a remote computer 880. The remote computer 880 may be a personal computer, a hand-held device, a server, a router, a network PC, a peer device or other common network node, and typically includes many or all of the elements described above relative to the computer 810. The logical connections depicted in FIG. 13 include a local area network (LAN) 871 and a wide area network (WAN) 873, but may also include other networks. Such networking environments are commonplace in offices, enterprise-wide computer networks, intranets and the Internet.

[00130] When used in a LAN networking environment, the computer 810 is connected to the LAN 871 through a network interface or adapter 870. When used in a WAN networking environment, the computer 810 typically includes a modem 872 or other means for establishing communications over the WAN 873, such as the Internet. The

modem 872, which may be internal or external, may be connected to the system bus 821 via the user input interface 860, or other appropriate mechanism. In a networked environment, program modules depicted relative to the computer 810, or portions thereof, may be stored in the remote memory storage device. By way of example, and not limitation, FIG. 13
5 illustrates remote application programs 885 as residing on remote computer 880. It will be appreciated that the network connections shown are exemplary and other means of establishing a communications link between the computers may be used.

[00131] It should also be noted that the different embodiments described herein can be combined in different ways. That is, parts of one or more embodiments can be combined
10 with parts of one or more other embodiments. All of this is contemplated herein.

[00132] Example 1 is a computing system, comprising:

[00133] a user interface mechanism that receives an utterance, sends the utterance to a linguistic processing system and receives a linguistic processing result indicative of an intent and a set of arguments recognized in the utterance;

15 [00134] a context identification system that identifies a current context of the computing system;

[00135] an action identifier that identifies an action based on the current context and based on the intent in the linguistic processing result; and

[00136] a cross-source search component that searches multiple different information
20 sources based on the current context and based on the set of arguments in the linguistic processing result to identify content on which the action is to be performed.

[00137] Example 2 is the computing system of any or all previous examples and further comprising:

[00138] an action component that controls a controlled system to take the action on
25 the identified content.

[00139] Example 3 is the computing system of any or all previous examples wherein the utterance identifies the content implicitly and wherein the cross-source search component uses the arguments as search terms to search the multiple different information sources to identify a plurality of different search results, each identifying a different item of
30 content.

[00140] Example 4 is the computing system of any or all previous examples and further comprising:

[00141] a relevancy generator that generates a relevancy measure for each of the different search results to obtain a given search result.

[00142] Example 5 is the computing system of any or all previous examples wherein the relevancy generator uses the arguments and current context to disambiguates otherwise ambiguous search results to obtain the given search result, based on the current context and the arguments.

5 [00143] Example 6 is the computing system of any or all previous examples wherein the linguistic processing system comprises:

[00144] a first linguistic processing system disposed on the computing system and including a speech recognition system and a natural language understanding system, the speech recognition system generating a textual representation of the utterance and the
10 natural language understanding system identifying the intent and arguments in a first linguistic processing result, and further disambiguating otherwise ambiguous arguments based on the current context.

[00145] Example 7 is the computing system of any or all previous examples wherein the user interface mechanism further sends the utterance to a second linguistic processing
15 system remote from the computing system, the second linguistic processing system generating a second linguistic processing result, the computing system combining the first and second linguistic processing results to obtain a final linguistic processing result.

[00146] Example 8 is the computing system of any or all previous examples wherein the context identification system comprises:

20 [00147] a personal identity detector that identifies a personal identity of the user.

[00148] Example 9 is the computing system of any or all previous examples wherein the context identification system comprises:

[00149] a proximity detector that identifies other people or things within a given proximity to the computing system.

25 [00150] Example 10 is the computing system of any or all previous examples wherein the context identification system comprises:

[00151] an open application detector that identifies open applications on the computing system.

[00152] Example 11 is the computing system of any or all previous examples wherein
30 the context identification system comprises:

[00153] a device detector that identifies a particular device that the computing system is disposed on.

[00154] Example 12 is the computing system of any or all previous examples wherein the context identification system comprises:

[00155] a calendar detector that identifies a particular time and items on a calendar of the user.

[00156] Example 13 is the computing system of any or all previous examples wherein the particular device comprises a mobile device.

5 [00157] Example 14 is the computing system of any or all previous examples wherein the context identification system comprises:

[00158] a location detector that identifies a location of the particular device on which the computing system is disposed.

[00159] Example 15 is the computing system of any or all previous examples wherein
10 the action identifier identifies the action as one of send, open, schedule a meeting, create, attach, email, set a reminder, share, present, launch and display.

[00160] Example 16 is a computer-implemented method, comprising:

[00161] receiving an utterance;

[00162] transmitting the utterance to a linguistic processing system;

15 [00163] receiving a linguistic processing result indicative of an intent and a set of arguments recognized in the utterance;

[00164] detecting a context of the computing system;

[00165] analyzing the context and the intent in the linguistic processing result to identify an action to take;

20 [00166] searching multiple different information sources based on the context and based on the set of arguments in the linguistic processing result to identify content on which the action is to be performed; and

[00167] automatically controlling a controlled system to perform the action on the identified content.

25 [00168] Example 17 is the computer implemented method of any or all previous examples and further comprising:

[00169] generating a relevancy measure corresponding to each of the different search results; and

[00170] ranking the different search results based on the corresponding relevancy
30 measures.

[00171] Example 18 is the computer implemented method of any or all previous examples and further comprising:

[00172] using the arguments and context to disambiguate otherwise ambiguous search results to obtain the given search result, based on the current context and the arguments.

[00173] Example 19 is a mobile device, comprising:

5 [00174] a user interface mechanism that receives an utterance, sends the utterance to a linguistic processing system and receives a linguistic processing result indicative of an intent and a set of arguments recognized in the utterance;

[00175] a context identification system that identifies a context of the computing system;

10 [00176] an action identifier that identifies an action based on the current context and based on the intent in the linguistic processing result;

[00177] a cross-source search component that searches multiple different information sources based on the context and based on the set of arguments in the linguistic processing result to identify content on which the action is to be performed; and

15 [00178] an action component that controls a controlled system to take the action on the identified content.

[00179] Example 20 is the mobile device of any or all previous examples wherein the utterance identifies the content implicitly and wherein the cross-source search component uses the arguments as search terms to search the multiple different information sources to
20 identify a plurality of different search results, each identifying a different item of content, and further comprising:

[00180] a relevancy generator that generates a relevancy measure for each of the different search results to obtain a given search result and that uses the arguments and current context to disambiguate otherwise ambiguous search results to obtain the given
25 search result, based on the context and the arguments.

[00181] Although the subject matter has been described in language specific to structural features and/or methodological acts, it is to be understood that the subject matter defined in the appended claims is not necessarily limited to the specific features or acts described above. Rather, the specific features and acts described above are disclosed as
30 example forms of implementing the claims.

CLAIMS

1. A computing system, comprising:
 - a user interface mechanism that receives an utterance, sends the utterance to a linguistic processing system and receives a linguistic processing result indicative of an intent and a set of arguments recognized in the utterance;
 - a context identification system that identifies a current context of the computing system;
 - an action identifier that identifies an action based on the current context and based on the intent in the linguistic processing result;
 - a cross-source search component that searches multiple different information sources based on the current context and based on the set of arguments in the linguistic processing result to identify content on which the action is to be performed; and
 - an action component that controls a controlled system to take the action on the identified content.
2. The computing system of claim 1 wherein the utterance identifies the content implicitly and wherein the cross-source search component uses the arguments as search terms to search the multiple different information sources to identify a plurality of different search results, each identifying a different item of content.
3. The computing system of claim 1 and further comprising:
 - a relevancy generator that generates a relevancy measure for each of the different search results to obtain a given search result.
4. The computing system of claim 3 wherein the relevancy generator uses the arguments and current context to disambiguates otherwise ambiguous search results to obtain the given search result, based on the current context and the arguments.
5. The computing system of claim 3 wherein the linguistic processing system comprises:
 - a first linguistic processing system disposed on the computing system and including a speech recognition system and a natural language understanding system, the speech recognition system generating a textual representation of the utterance and the natural language understanding system identifying the intent and arguments in a first linguistic processing result, and further disambiguating otherwise ambiguous arguments based on the current context.

6. The computing system of claim 5 wherein the user interface mechanism further sends the utterance to a second linguistic processing system remote from the computing system, the second linguistic processing system generating a second linguistic processing result, the computing system combining the first and second linguistic processing results to obtain a final linguistic processing result.
7. The computing system of claim 3 wherein the context identification system comprises:
 - a personal identity detector that identifies a personal identity of the user;
 - a proximity detector that identifies other people or things within a given proximity to the computing system; and
 - an open application detector that identifies open applications on the computing system.
8. The computing system of claim 3 wherein the context identification system comprises:
 - a device detector that identifies a particular device that the computing system is disposed on.
9. The computing system of claim 3 wherein the context identification system comprises:
 - a calendar detector that identifies a particular time and items on a calendar of the user.
10. The computing system of claim 8 wherein the particular device comprises a mobile device.
11. The computing system of claim 8 wherein the context identification system comprises:
 - a location detector that identifies a location of the particular device on which the computing system is disposed.
12. The computing system of claim 3 wherein the action identifier identifies the action as one of send, open, schedule a meeting, create, attach, email, set a reminder, share, present, launch and display.
13. A computer-implemented method, comprising:
 - receiving an utterance;
 - transmitting the utterance to a linguistic processing system;
 - receiving a linguistic processing result indicative of an intent and a set of arguments recognized in the utterance;

detecting a context of the computing system;
analyzing the context and the intent in the linguistic processing result to identify an action to take;
searching multiple different information sources based on the context and based on the set of arguments in the linguistic processing result to identify content on which the action is to be performed; and
automatically controlling a controlled system to perform the action on the identified content.

14. A mobile device, comprising:
a user interface mechanism that receives an utterance, sends the utterance to a linguistic processing system and receives a linguistic processing result indicative of an intent and a set of arguments recognized in the utterance;
a context identification system that identifies a context of the computing system;
an action identifier that identifies an action based on the current context and based on the intent in the linguistic processing result;
a cross-source search component that searches multiple different information sources based on the context and based on the set of arguments in the linguistic processing result to identify content on which the action is to be performed; and
an action component that controls a controlled system to take the action on the identified content.
15. The mobile device of claim 14 wherein the utterance identifies the content implicitly and wherein the cross-source search component uses the arguments as search terms to search the multiple different information sources to identify a plurality of different search results, each identifying a different item of content, and further comprising:
a relevancy generator that generates a relevancy measure for each of the different search results to obtain a given search result and that uses the arguments and current context to disambiguate otherwise ambiguous search results to obtain the given search result, based on the context and the arguments.

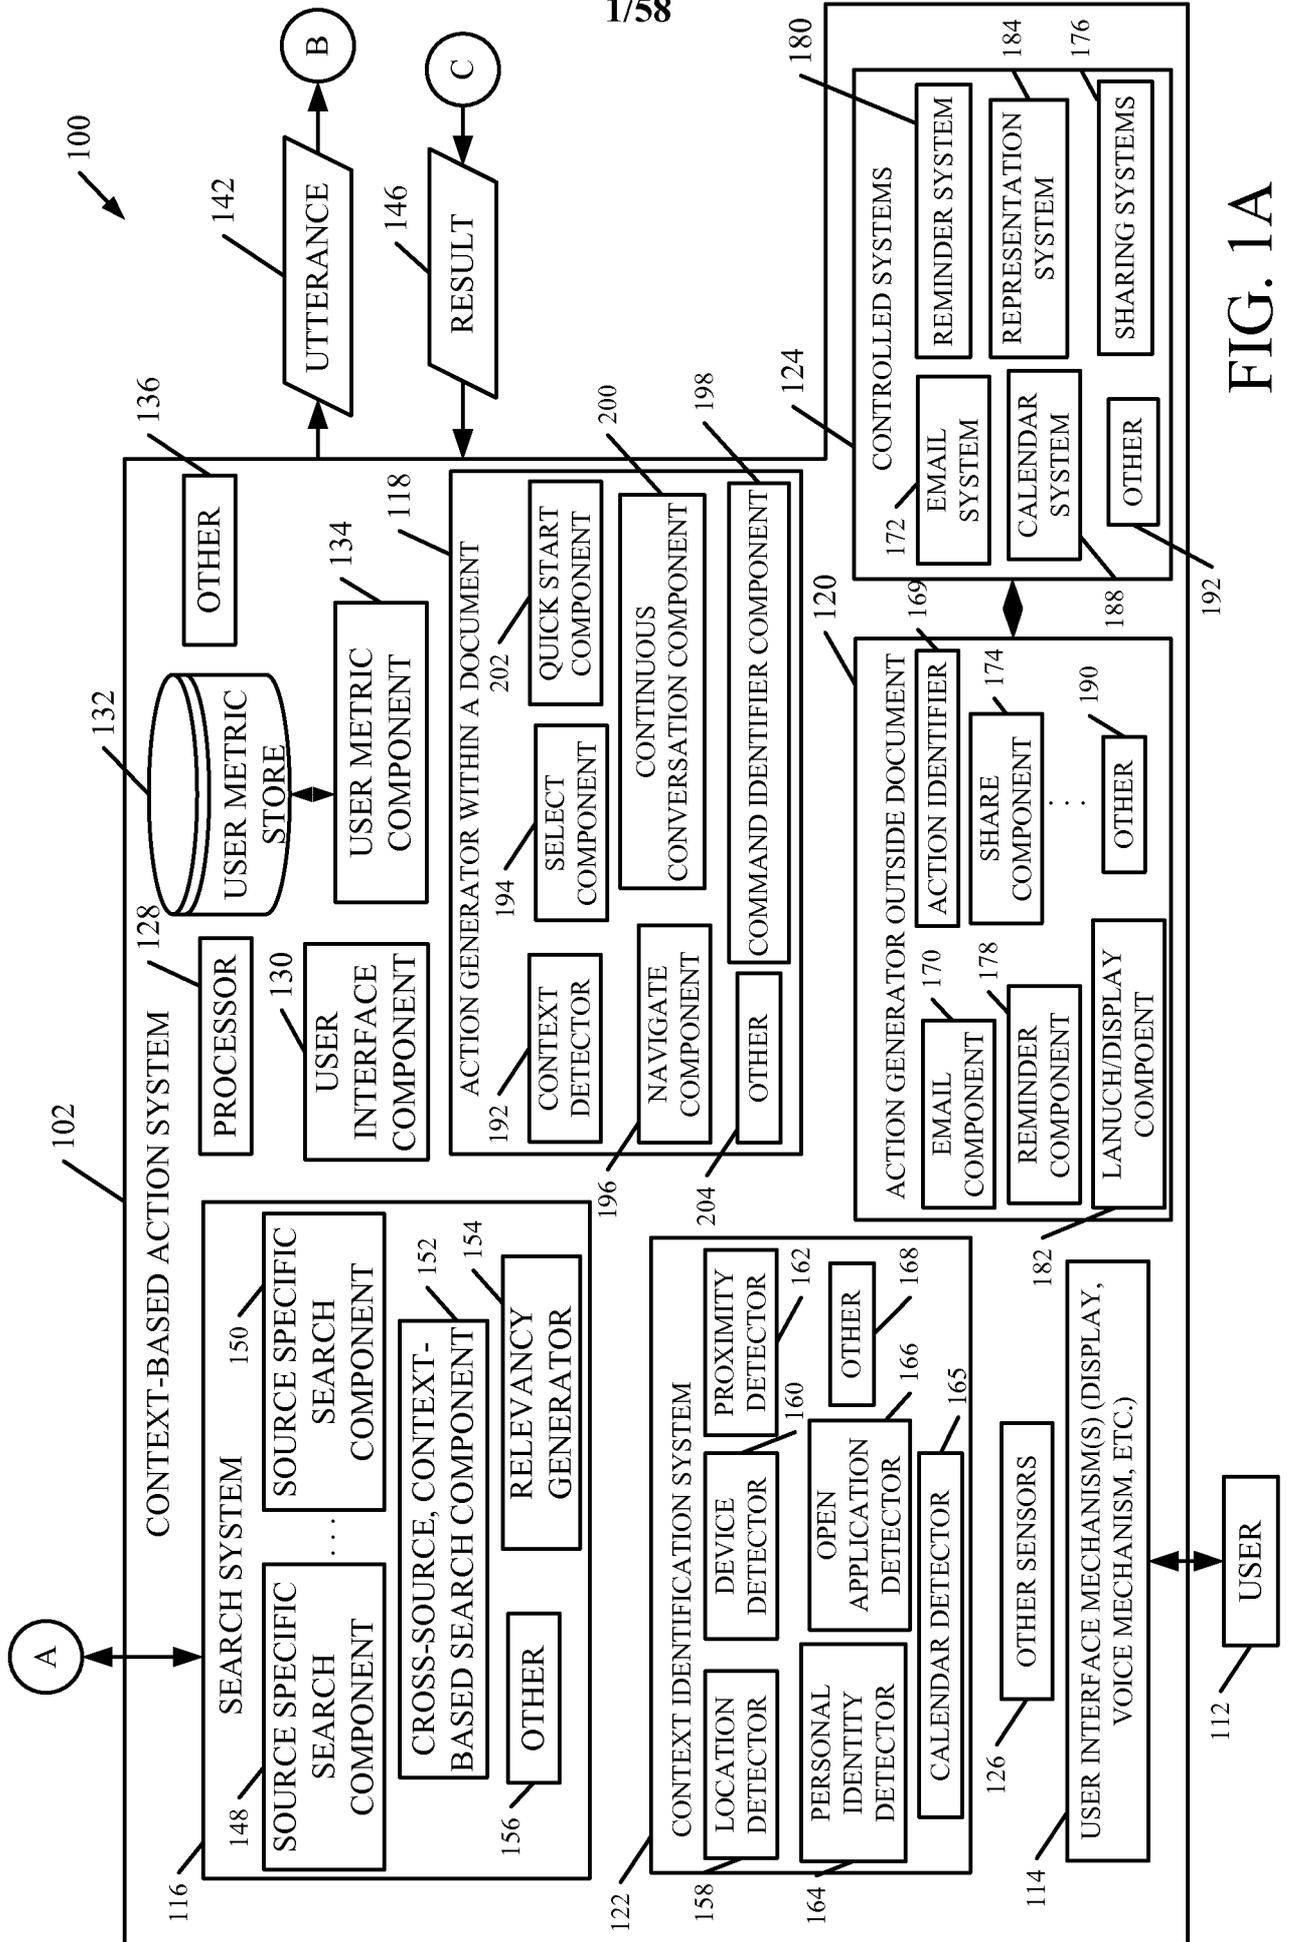


FIG. 1A

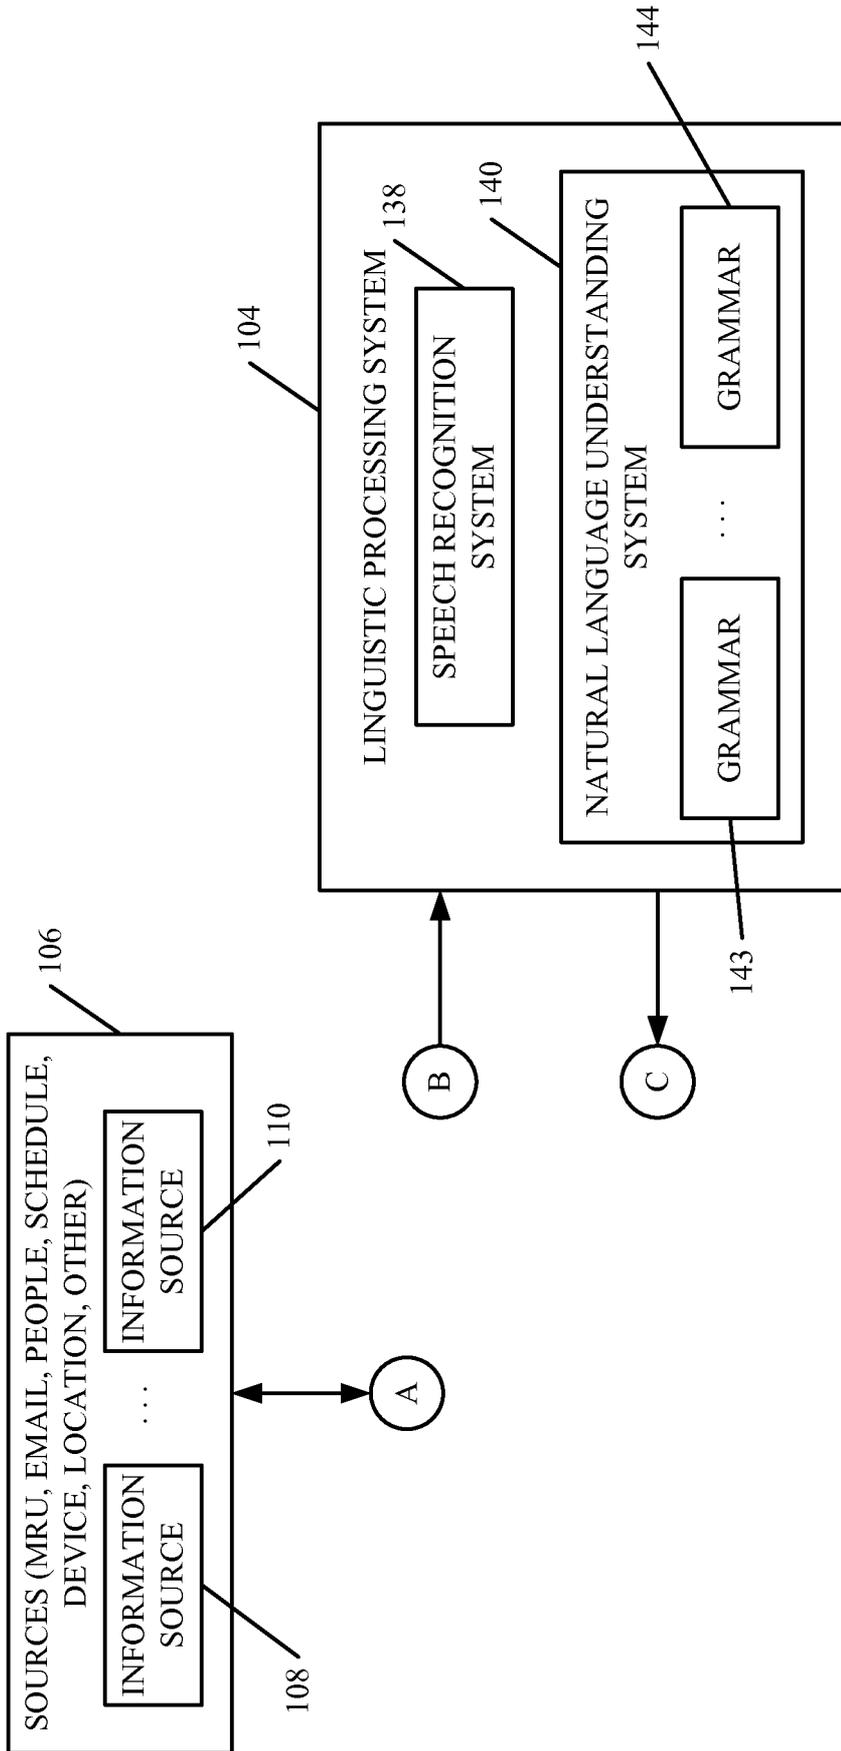


FIG. 1B

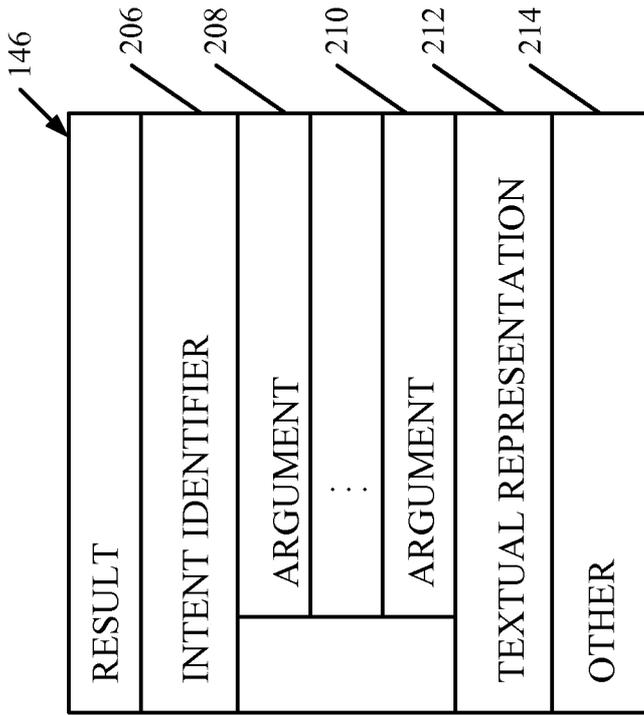


FIG. 2

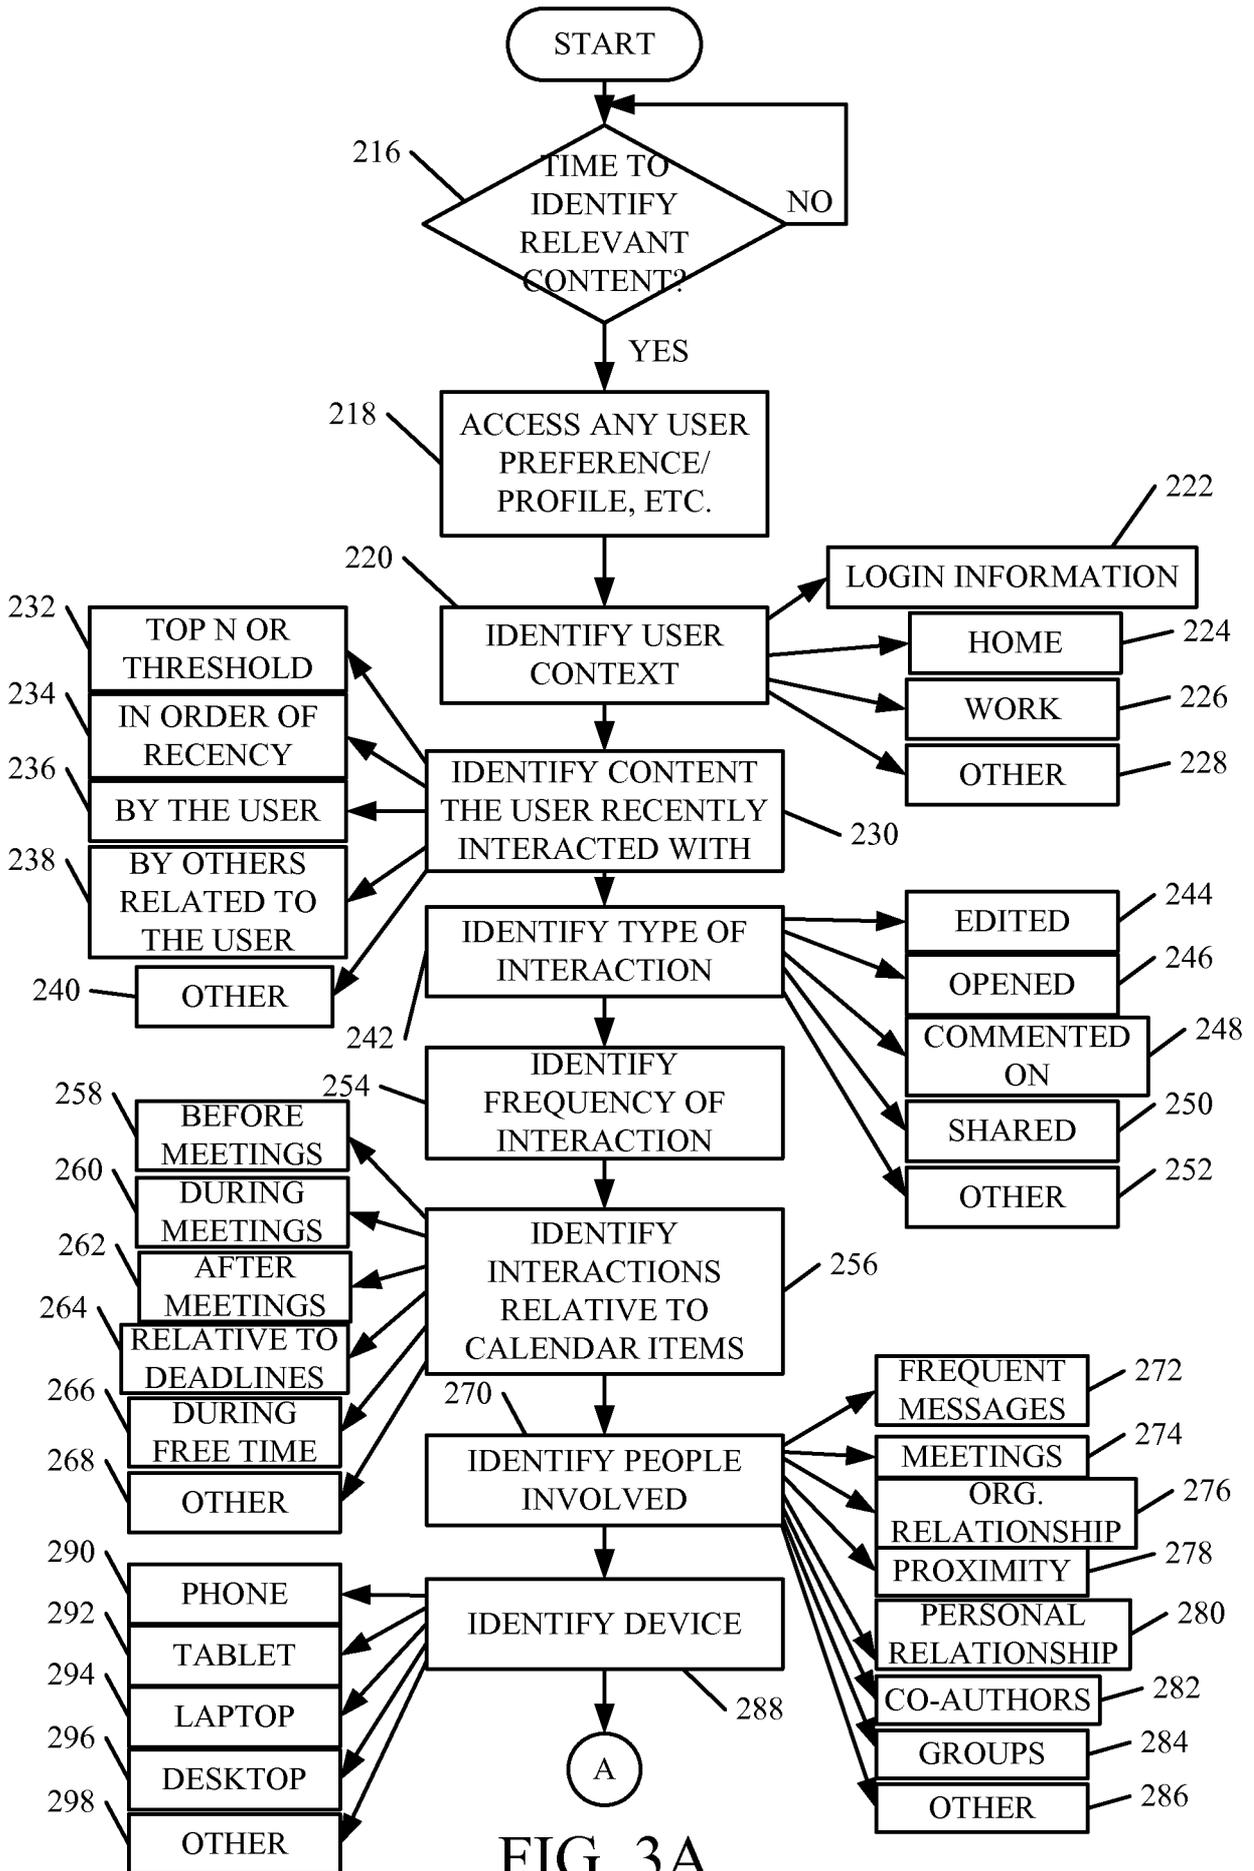


FIG. 3A

5/58

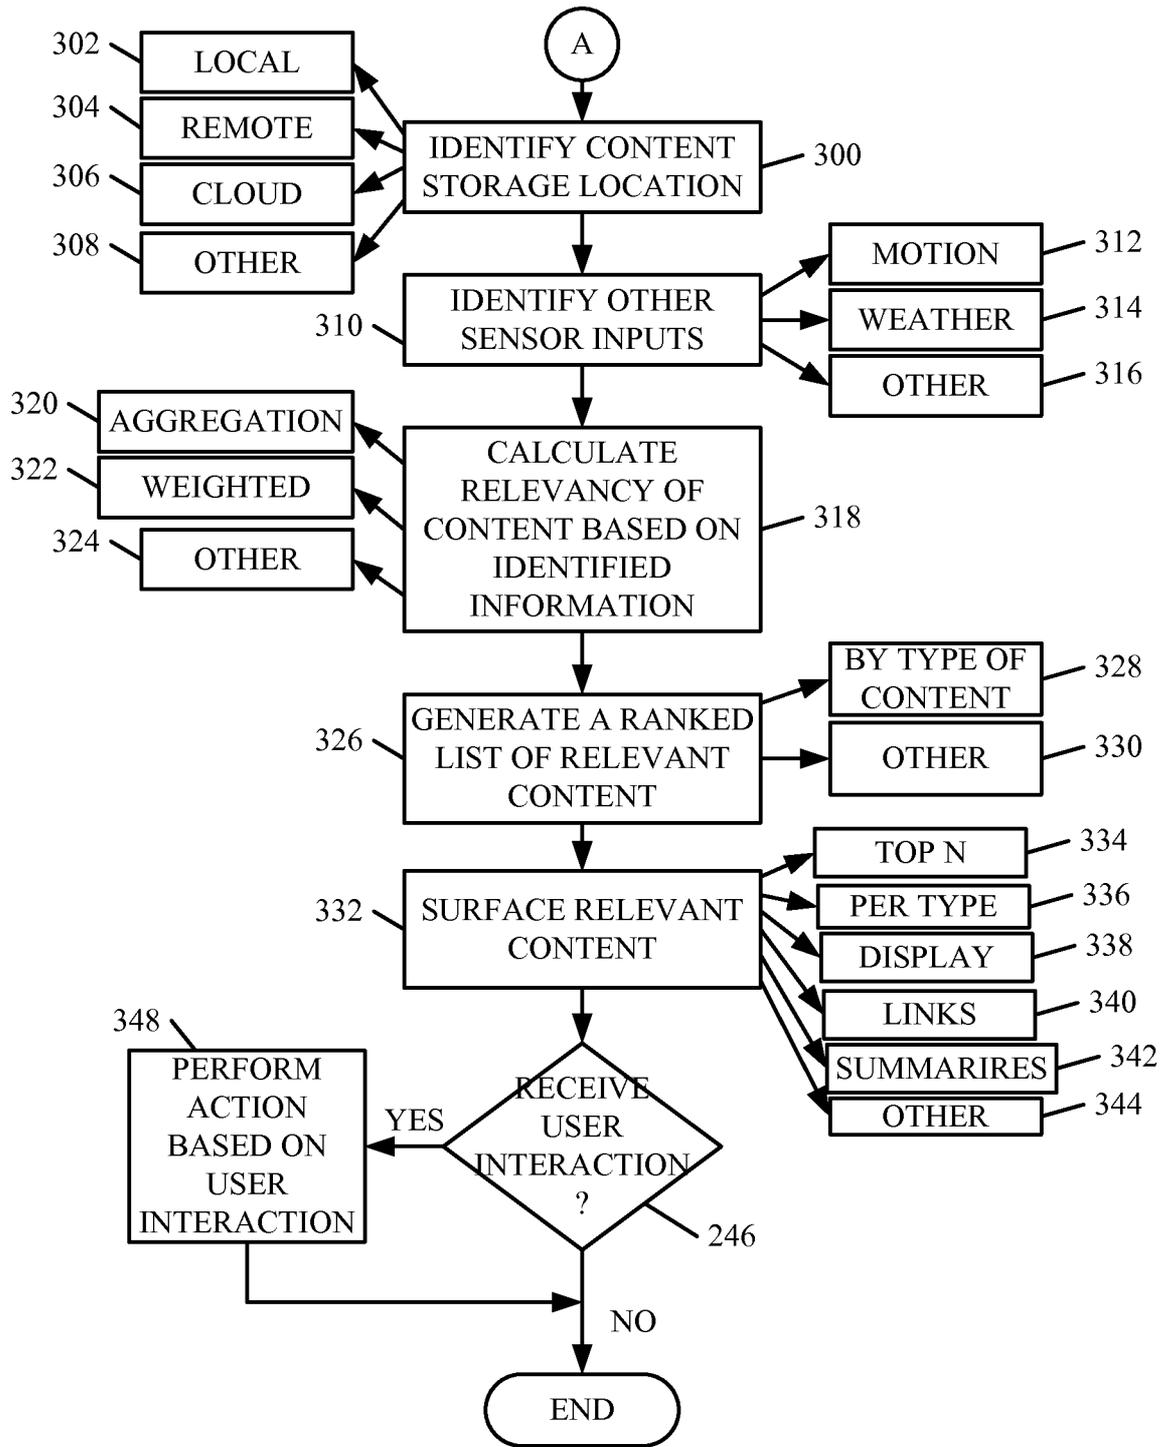


FIG. 3B

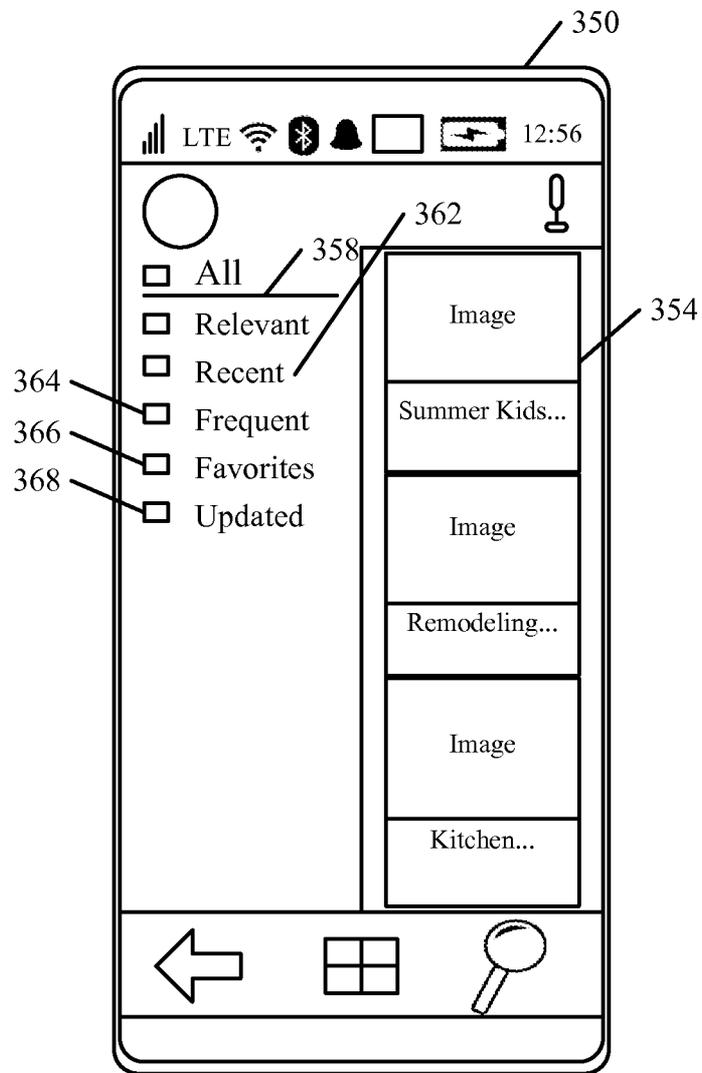


FIG. 3C

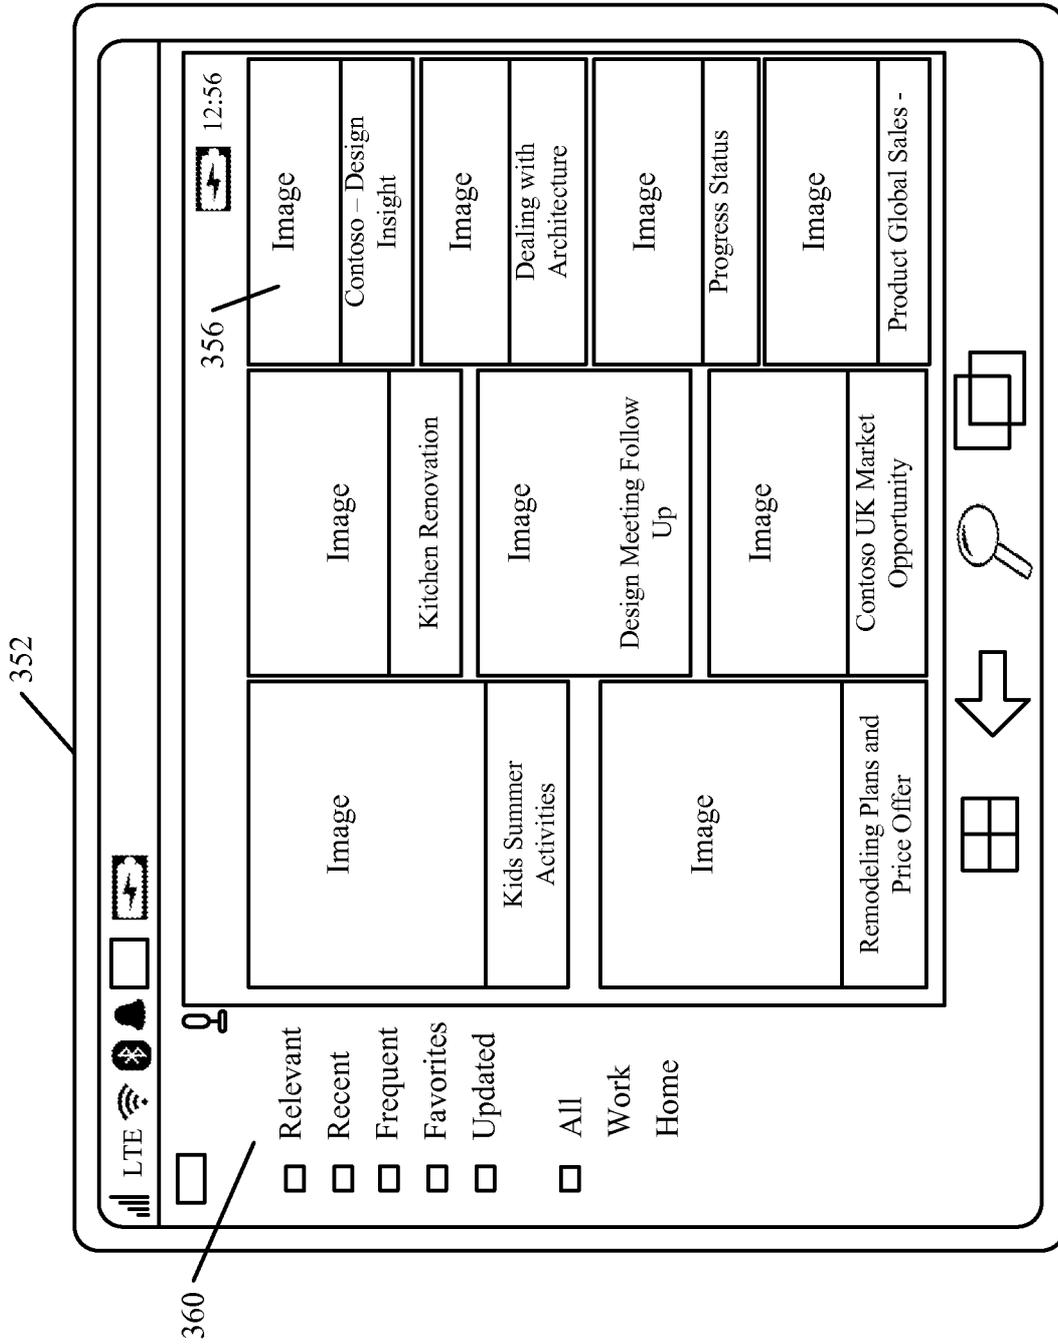


FIG. 3D

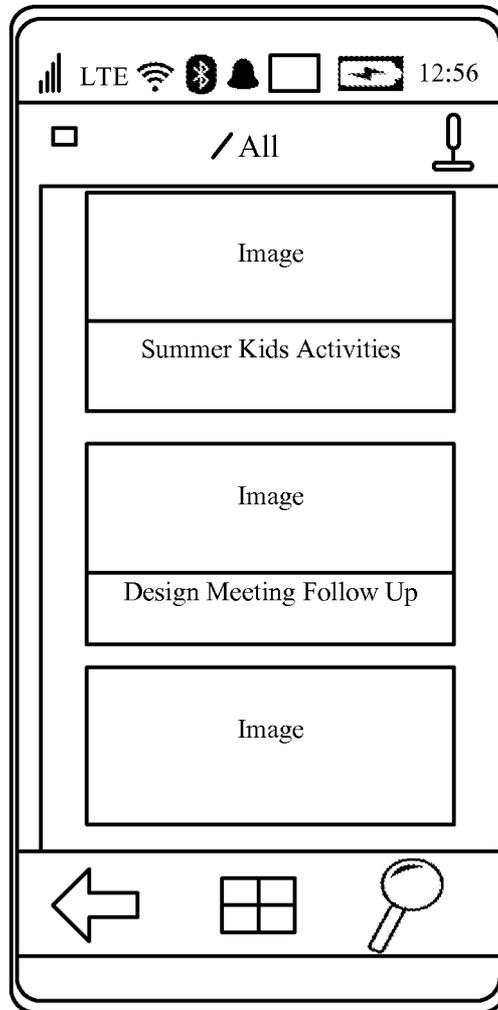
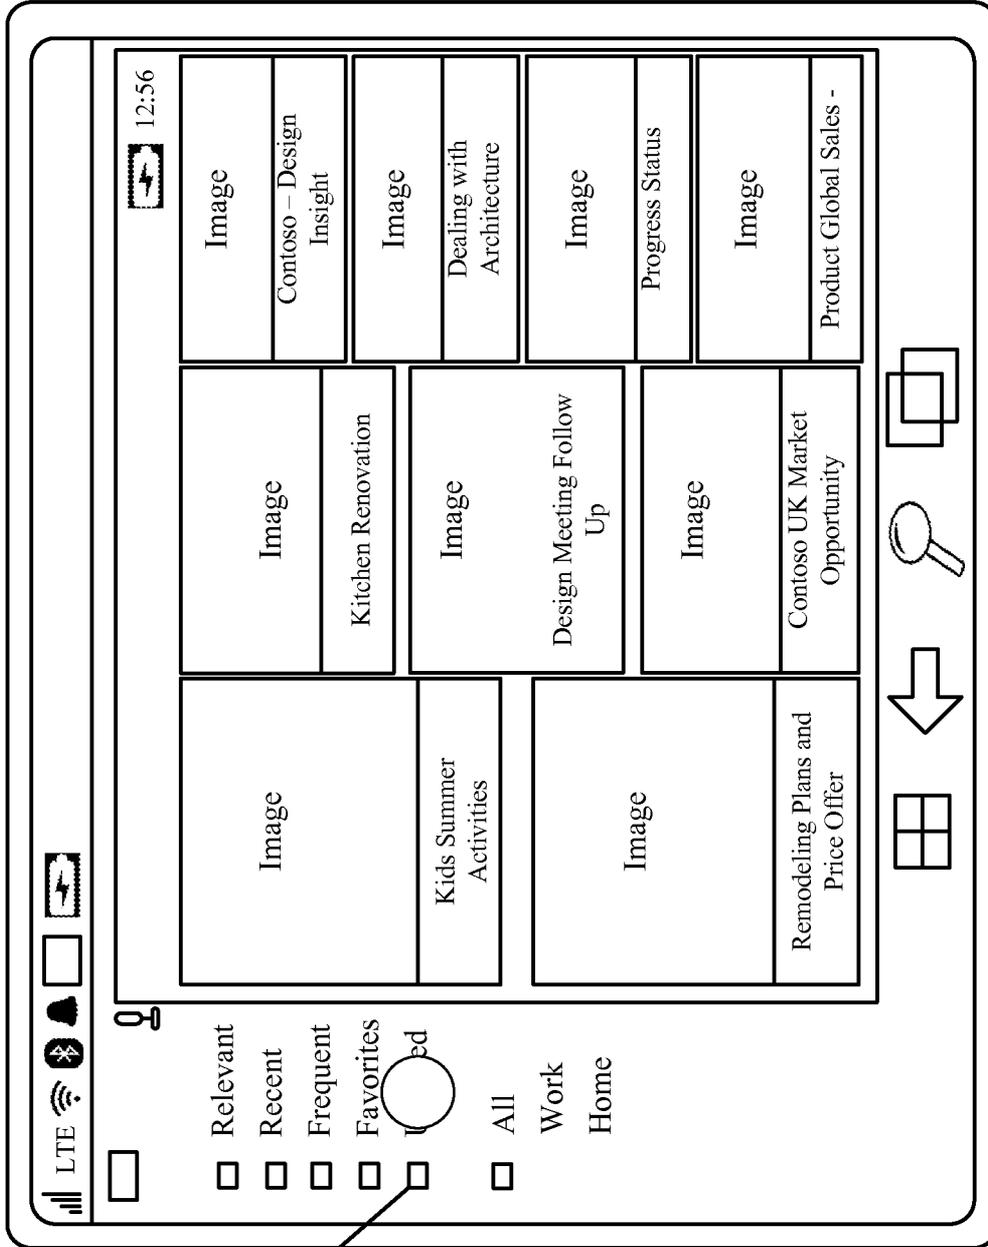


FIG. 3E



368

FIG. 3F

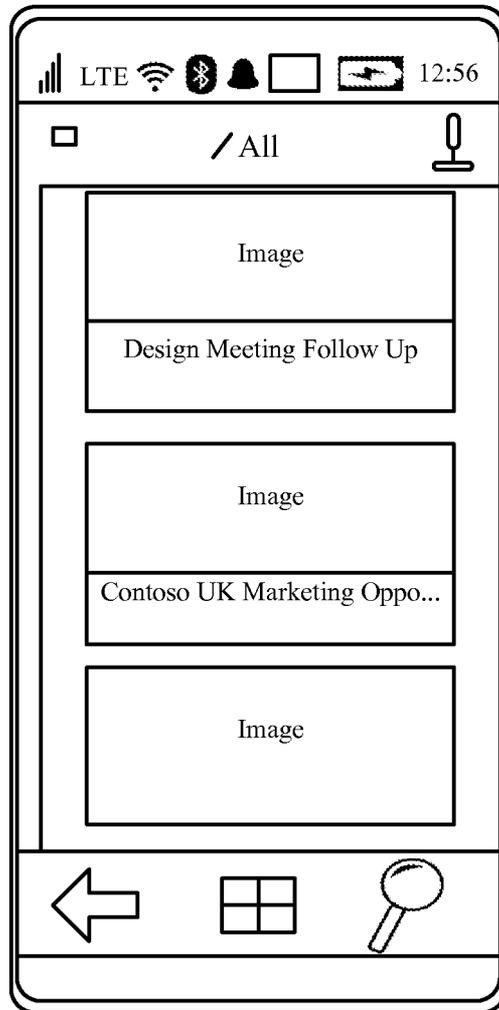


FIG. 3G

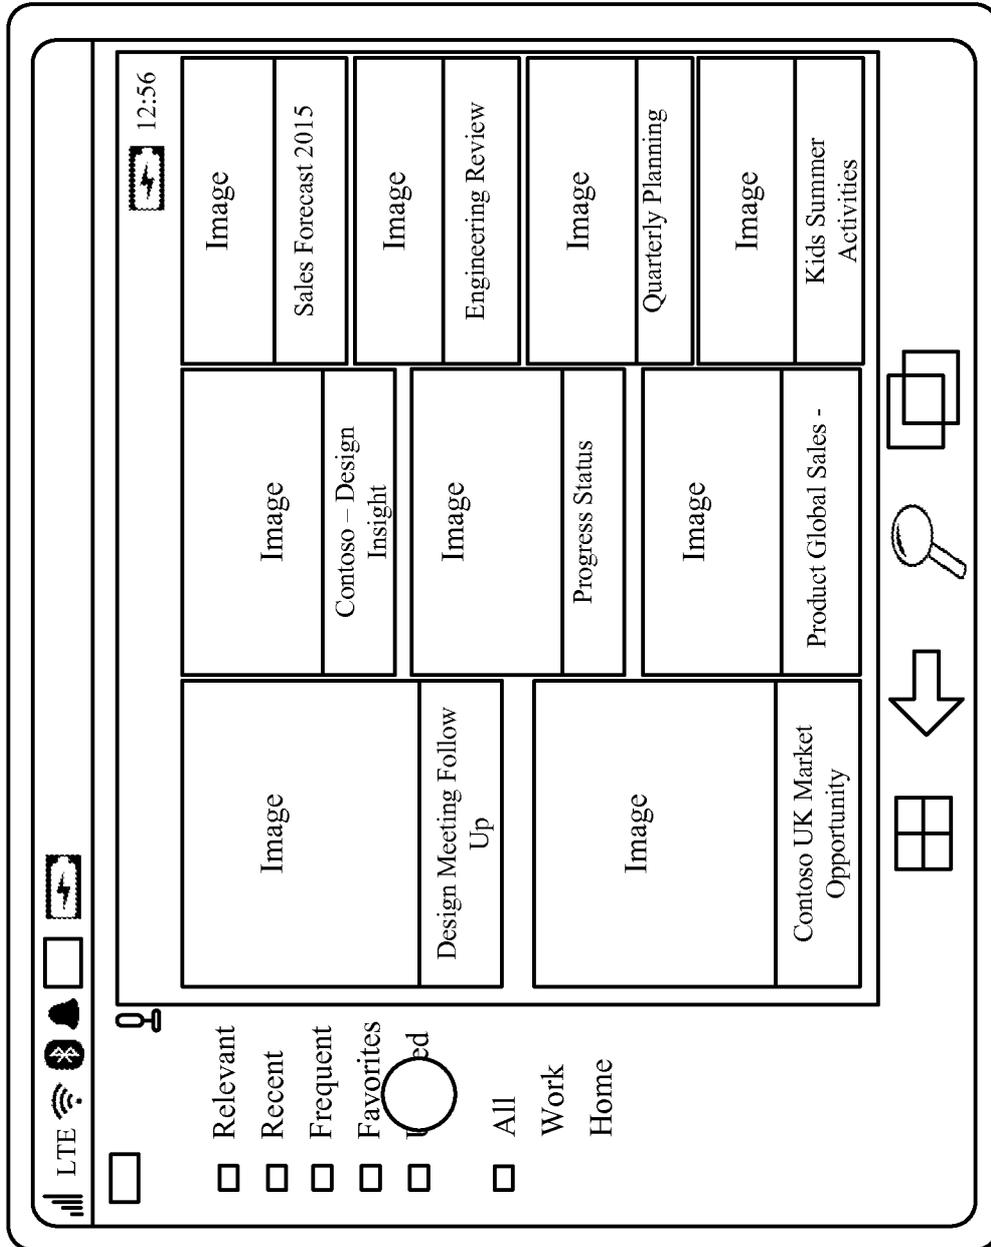


FIG. 3H

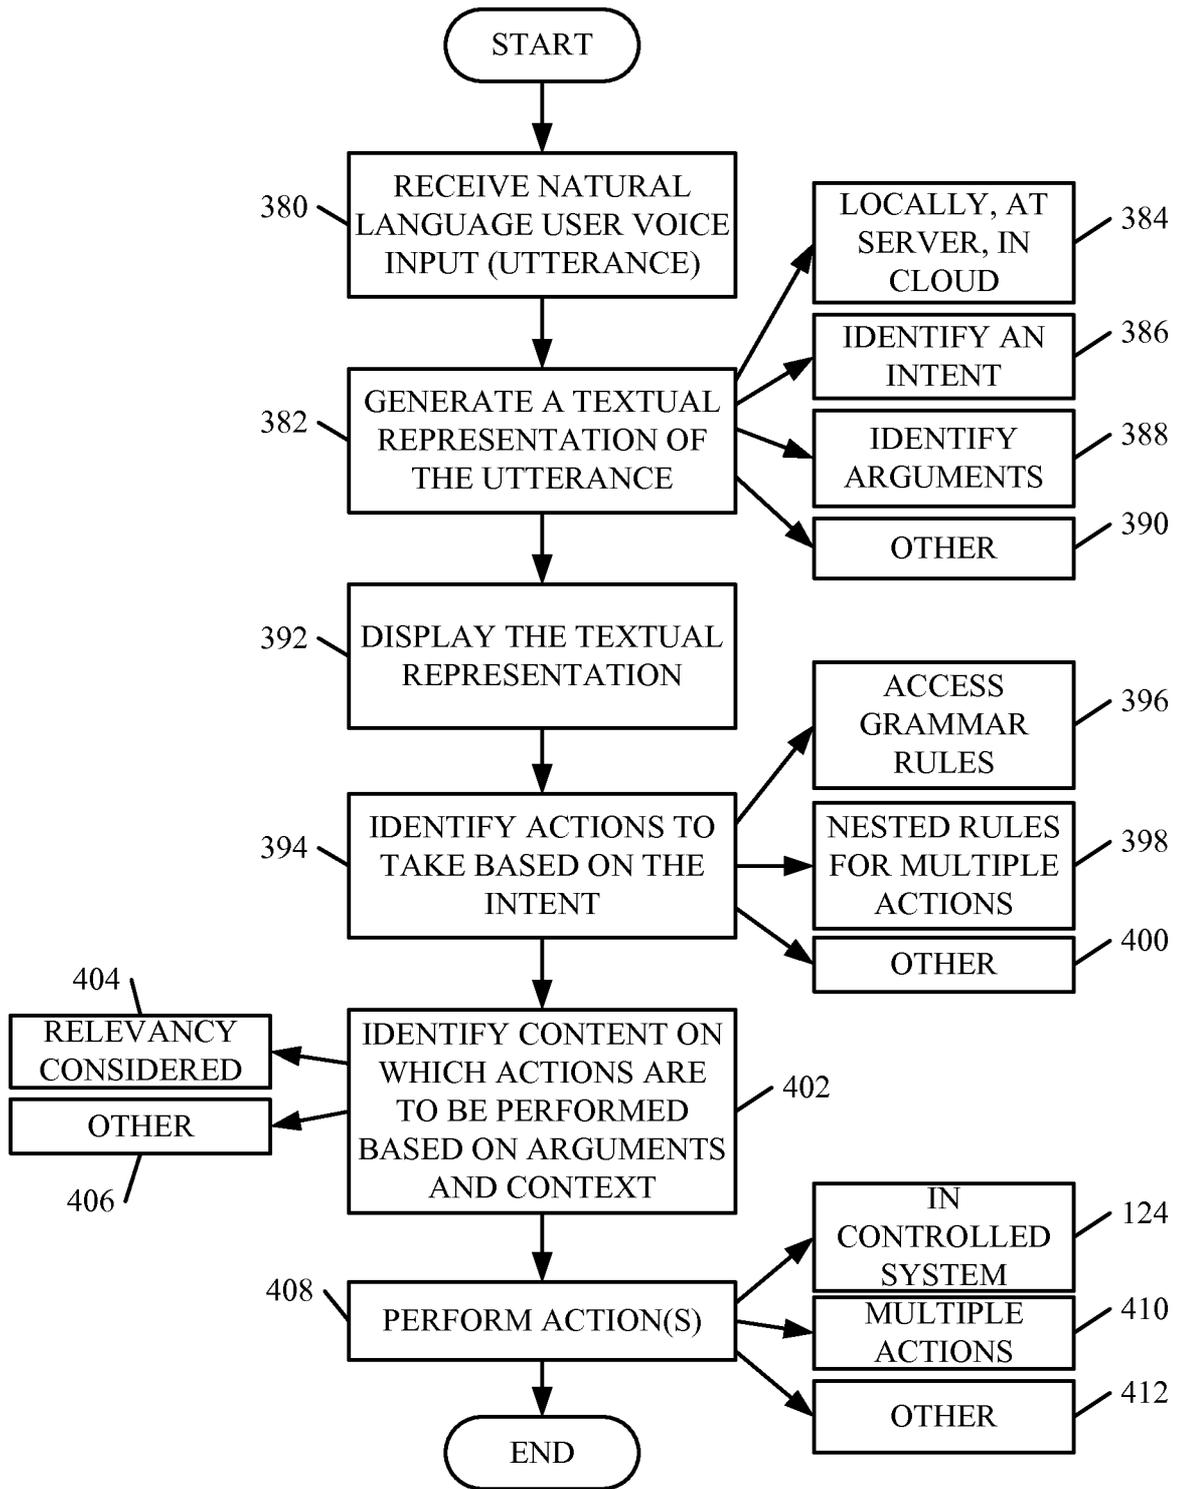


FIG. 4

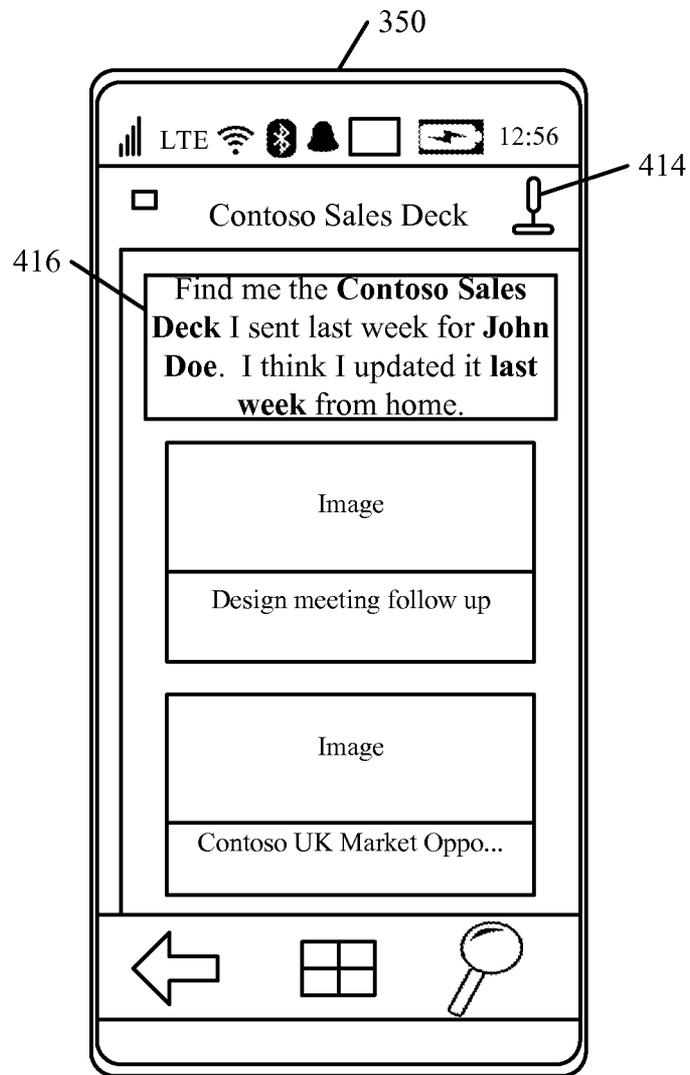


FIG. 4A

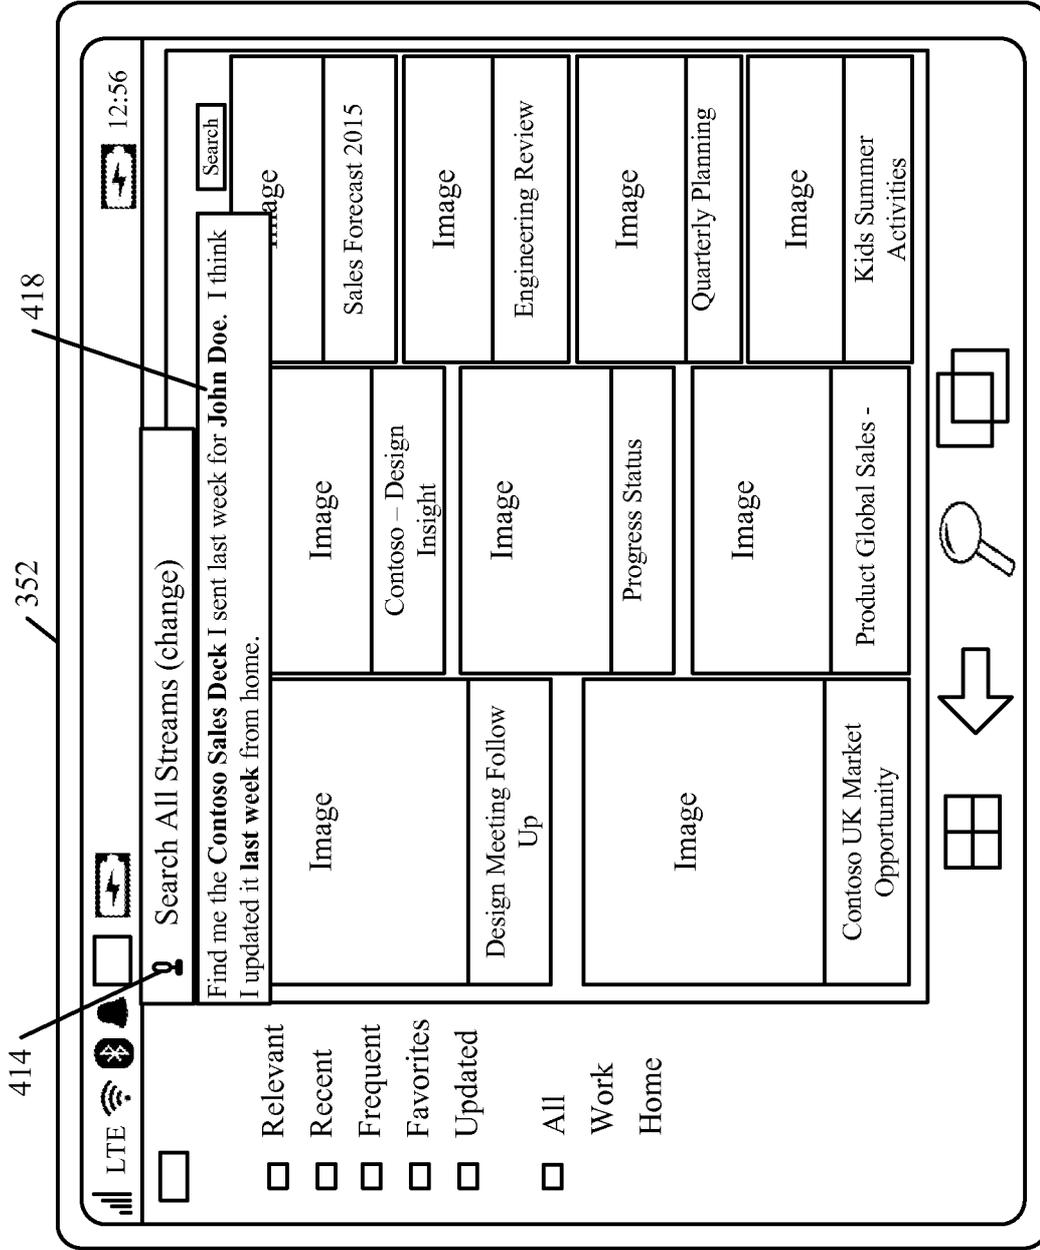


FIG. 4B

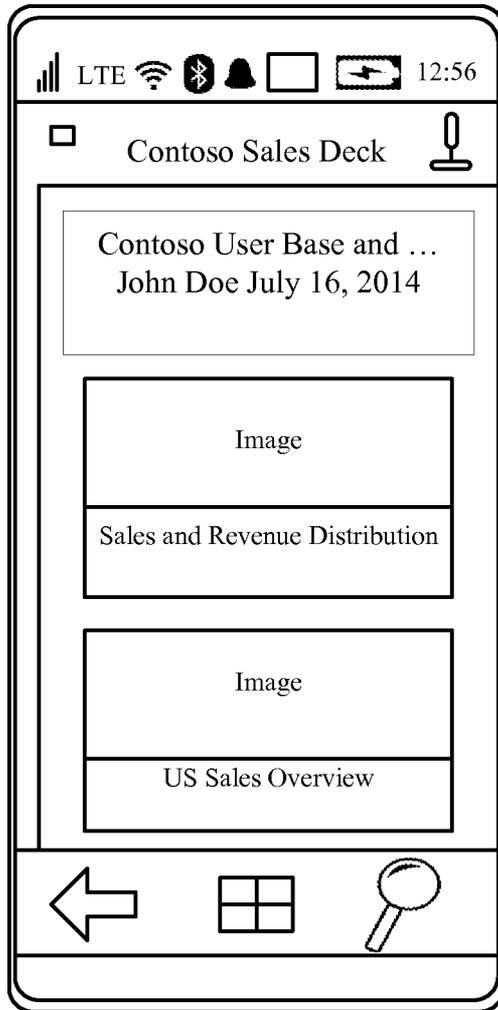
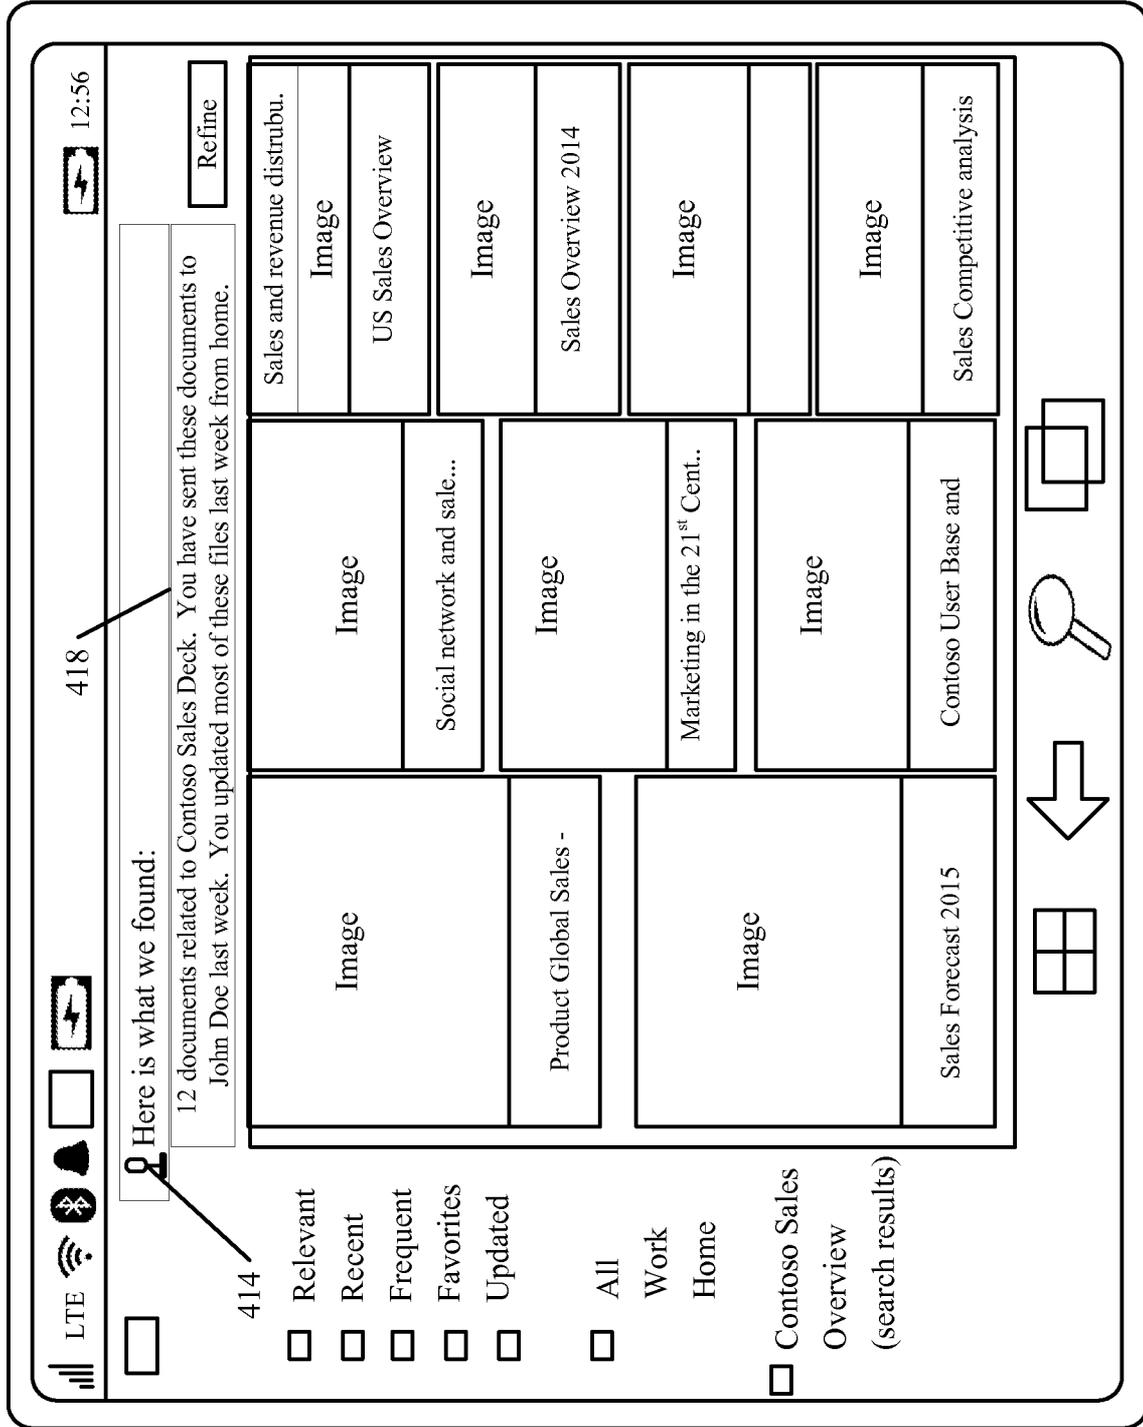


FIG. 4C



352

418

414

FIG. 4D

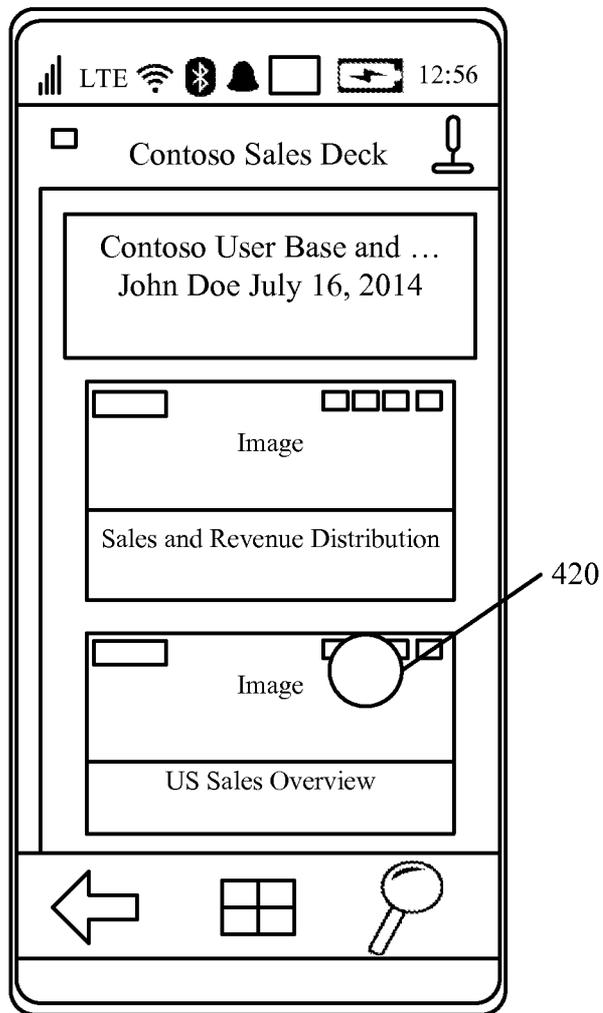


FIG. 4E

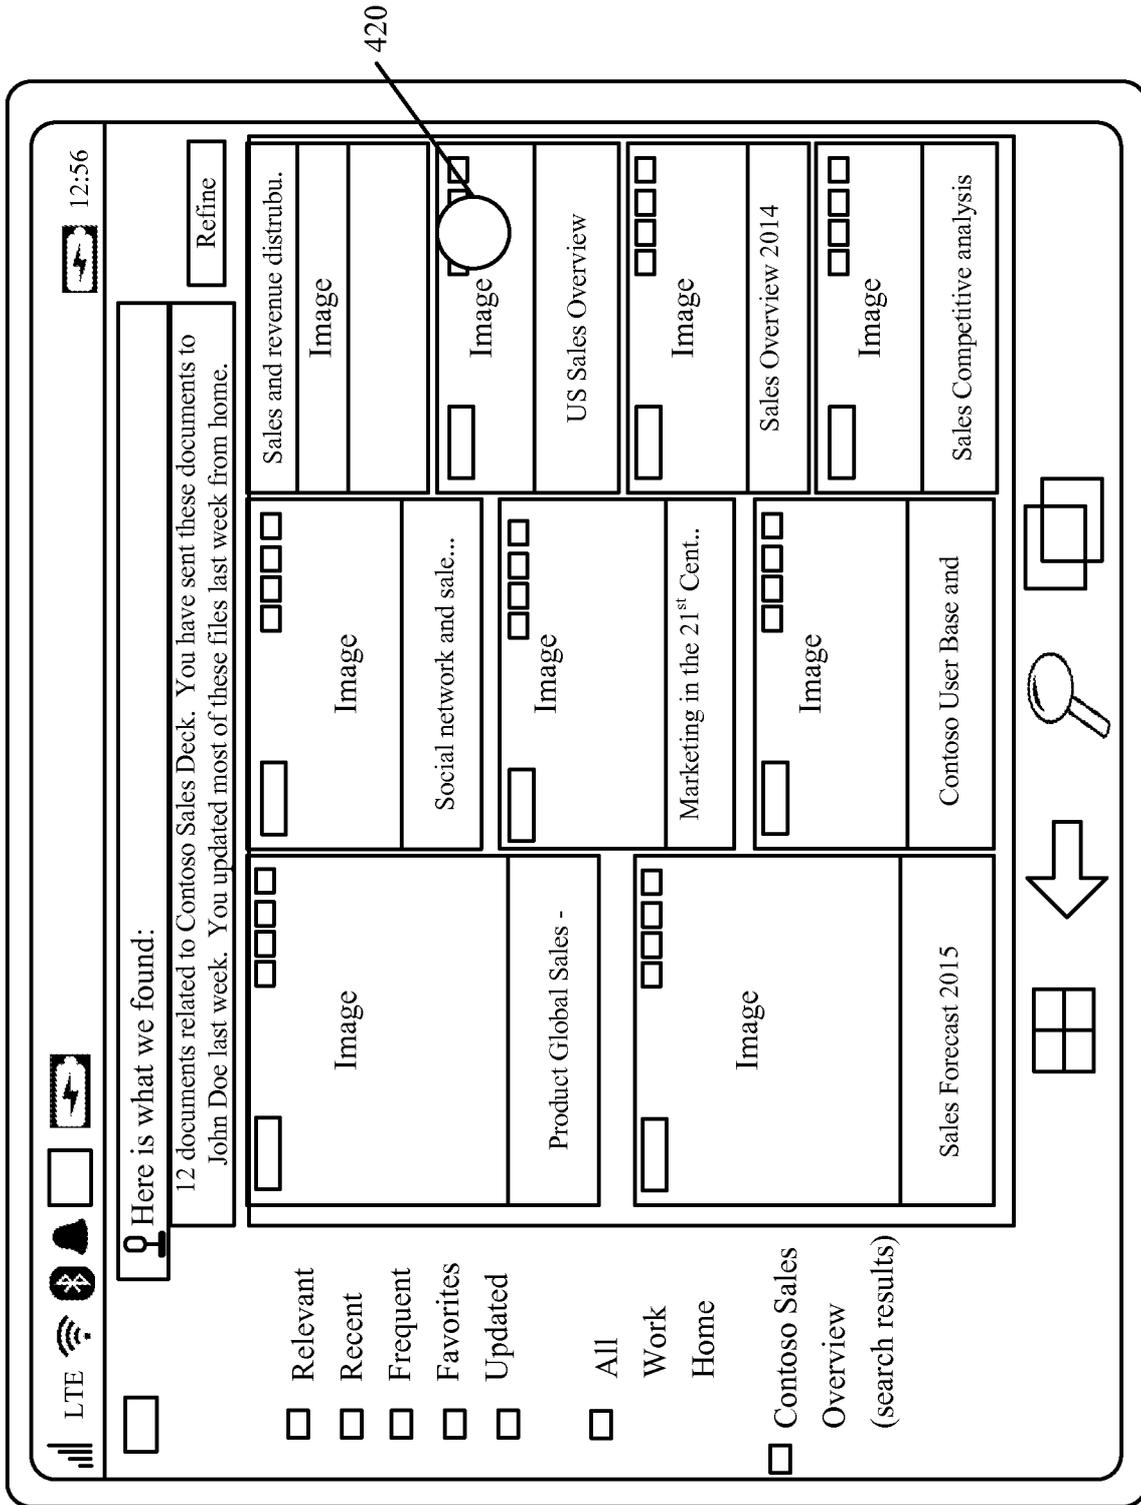


FIG. 4F

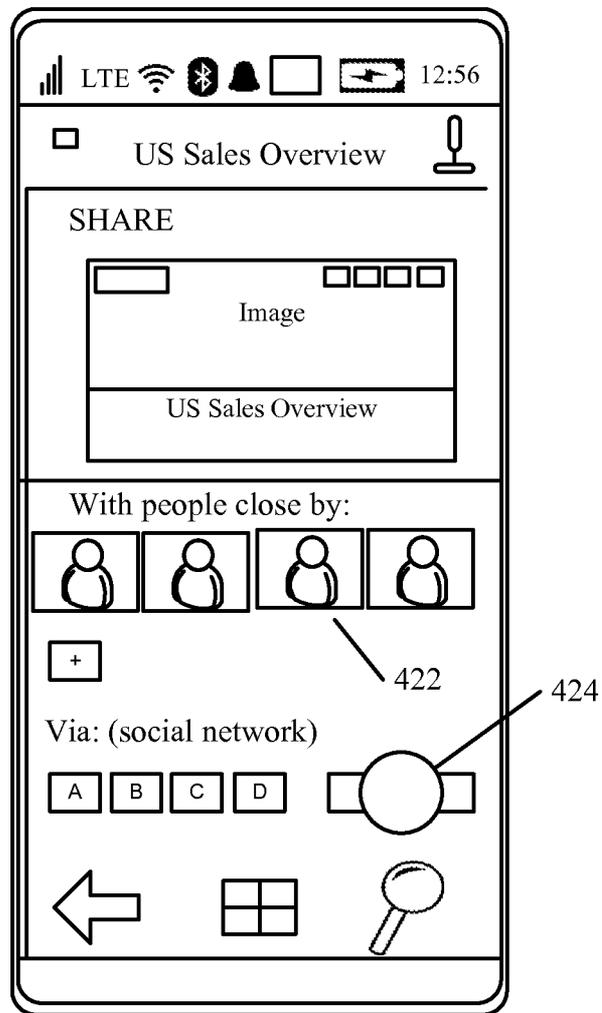


FIG. 4G

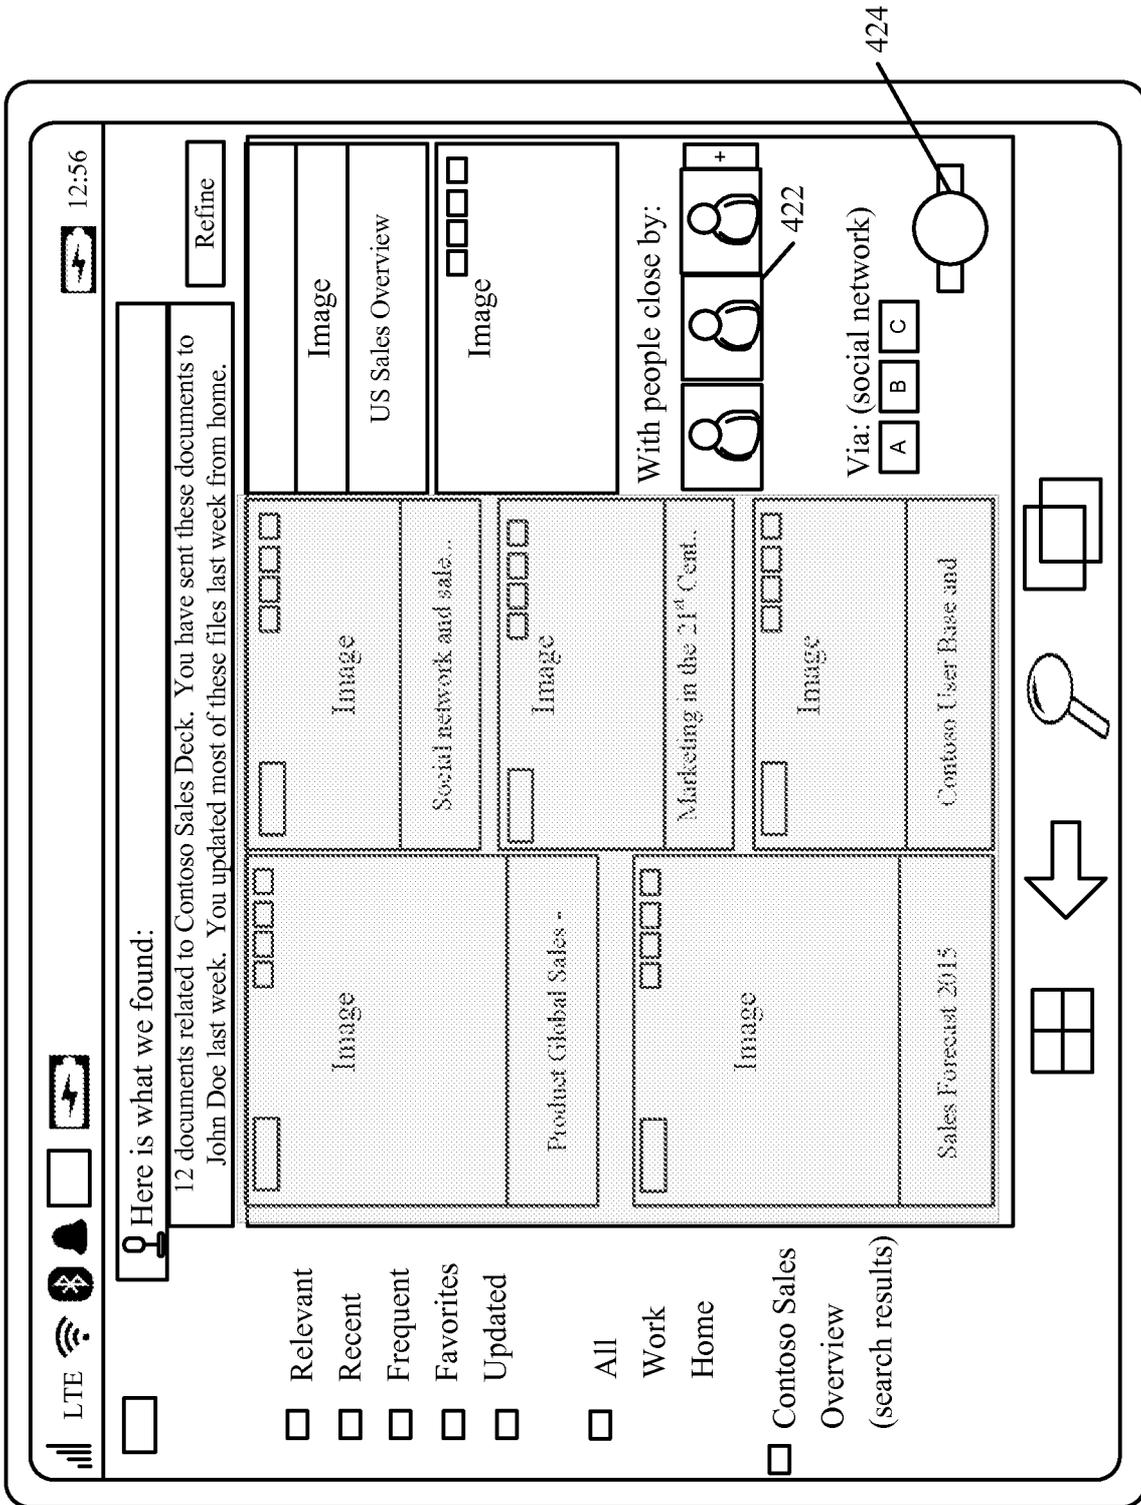


FIG. 4H

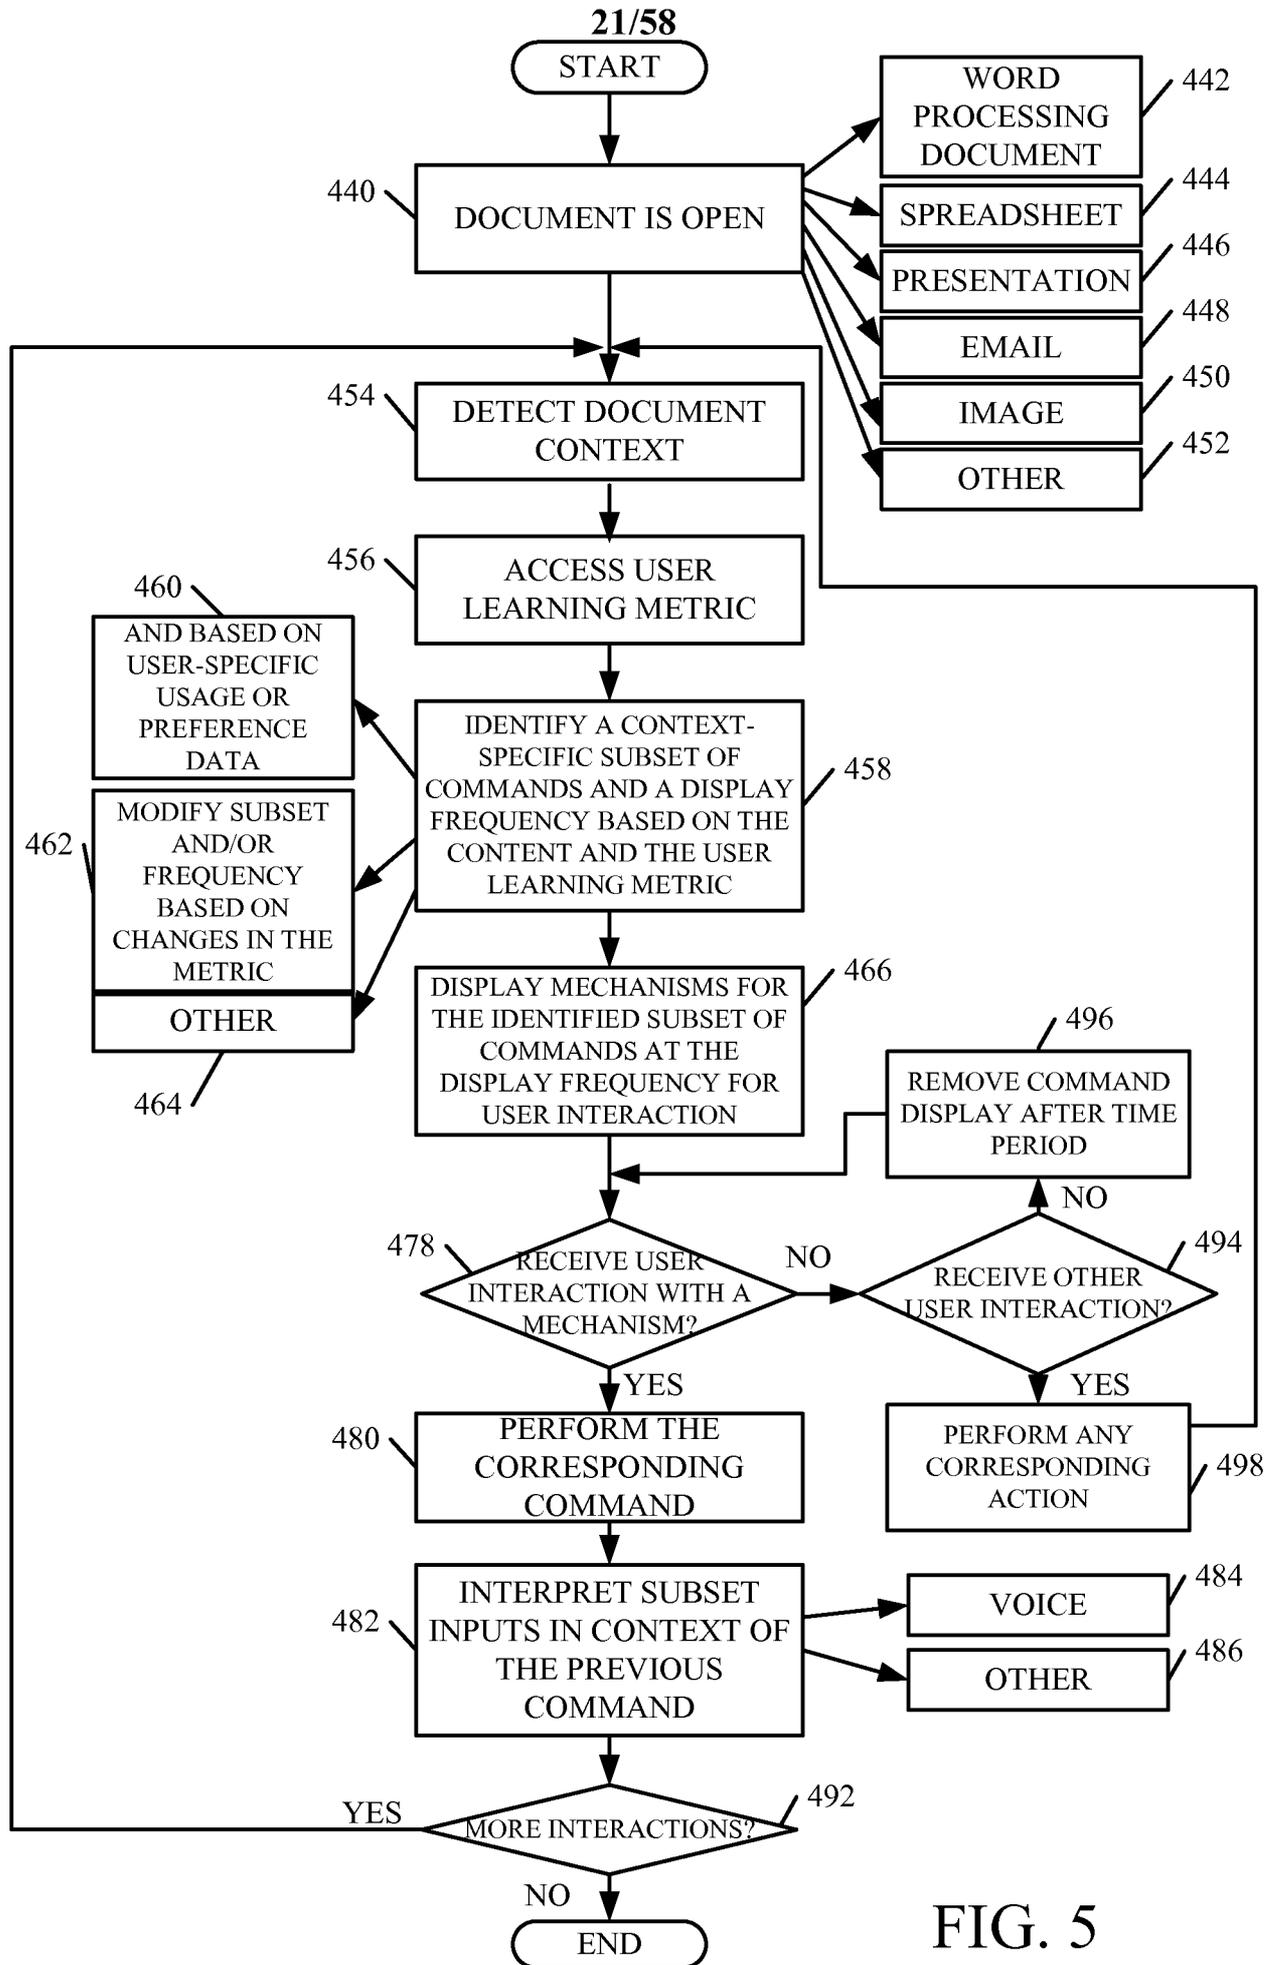


FIG. 5

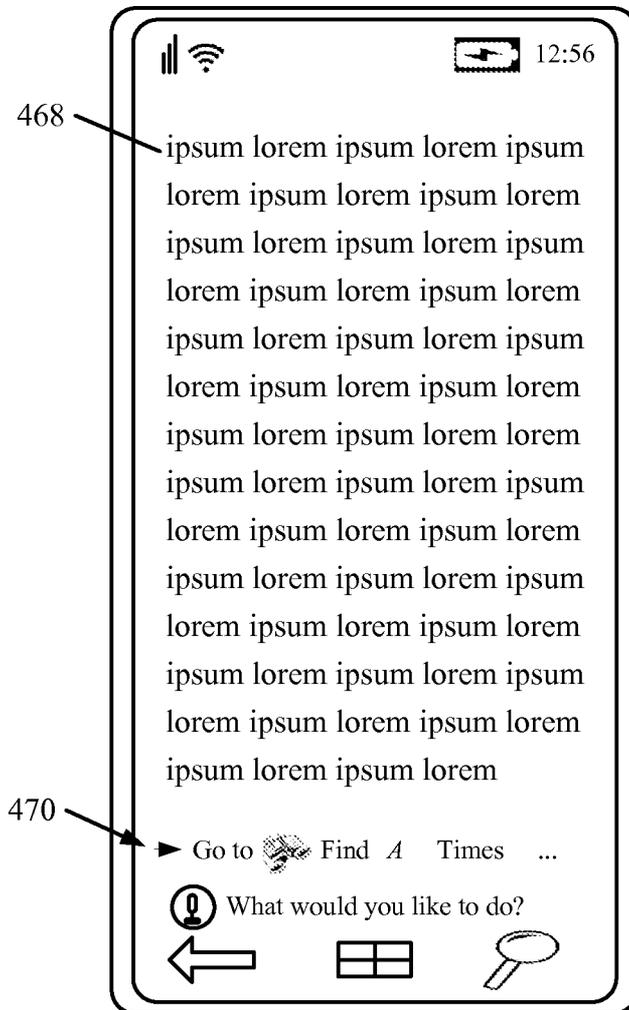


FIG. 5A

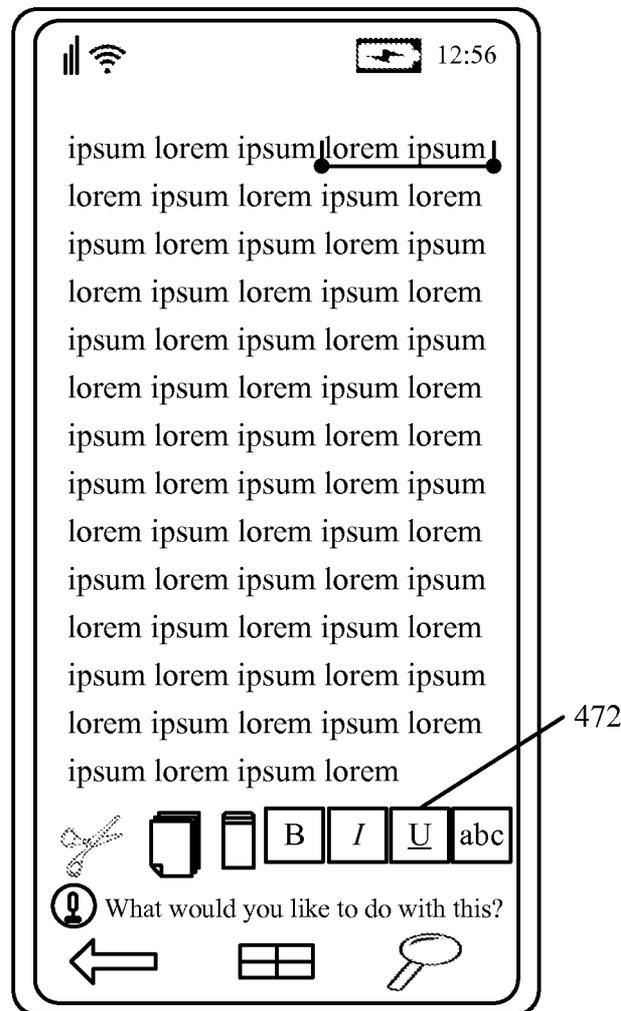


FIG. 5B

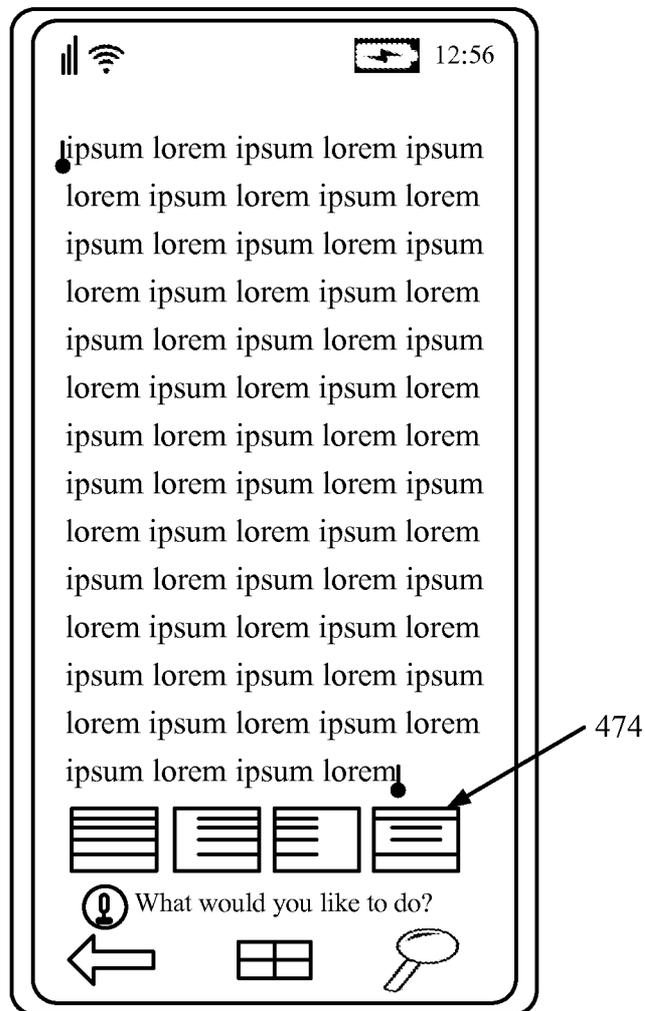


FIG. 5C

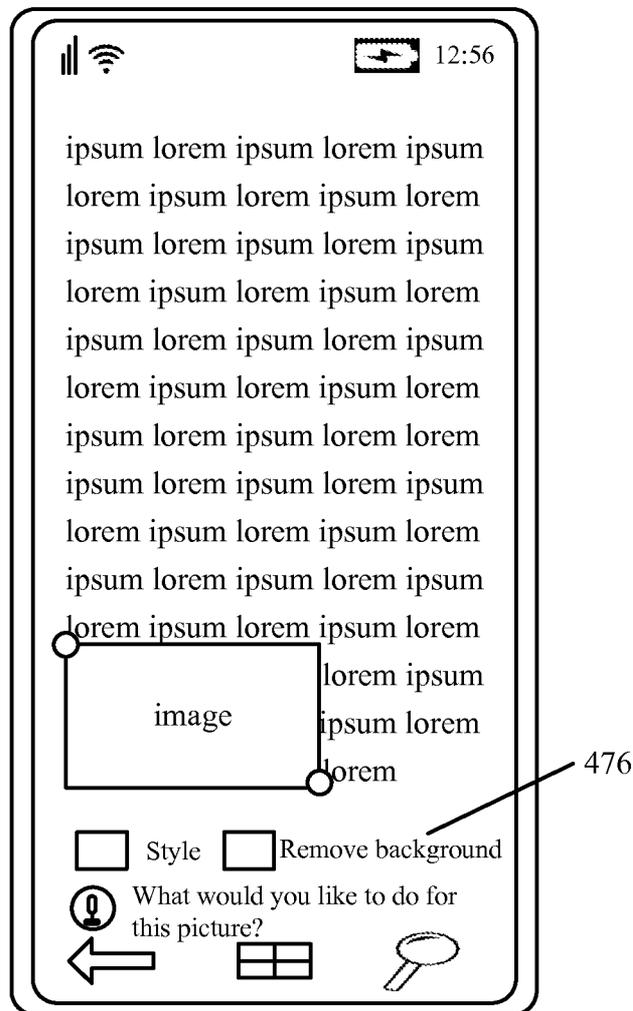


FIG. 5D

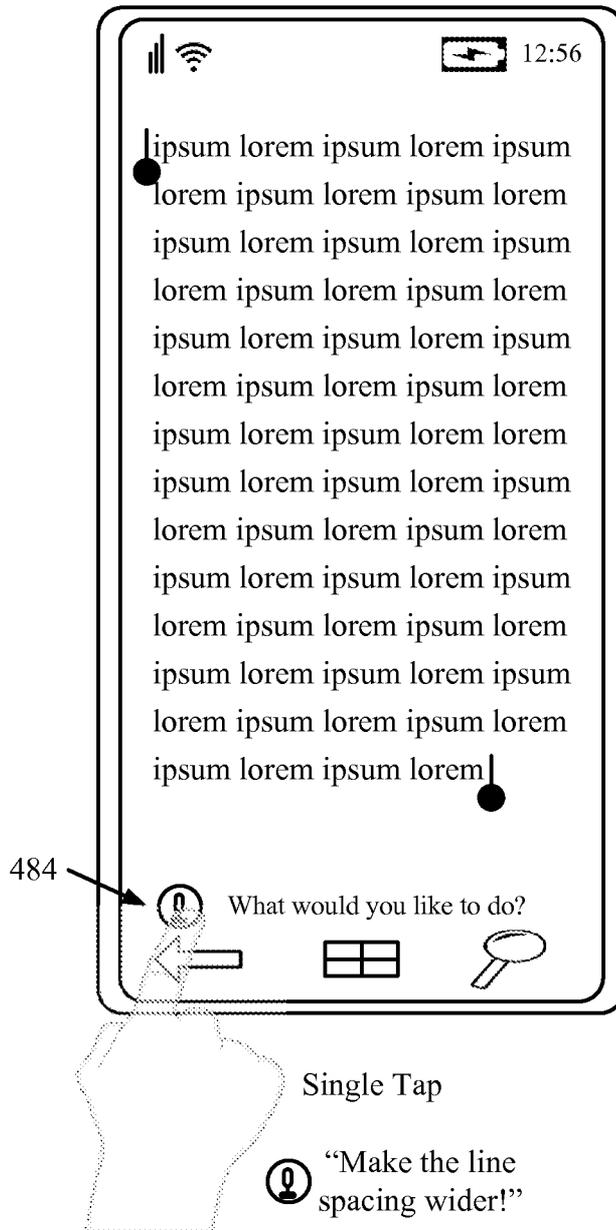


FIG. 5E

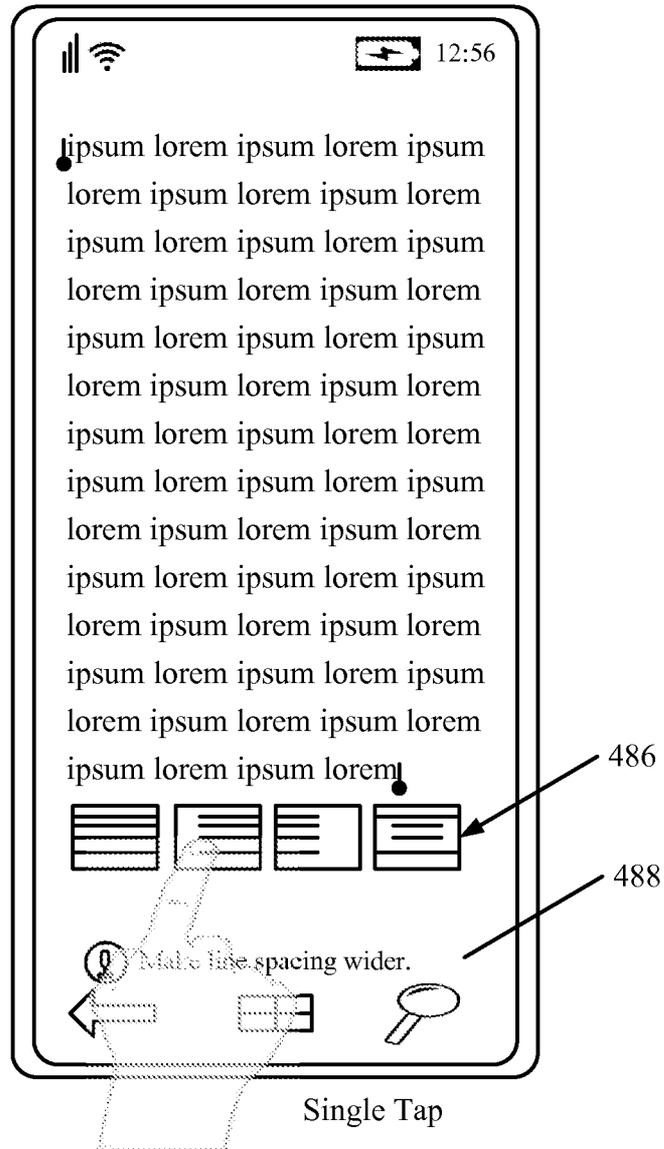


FIG. 5F

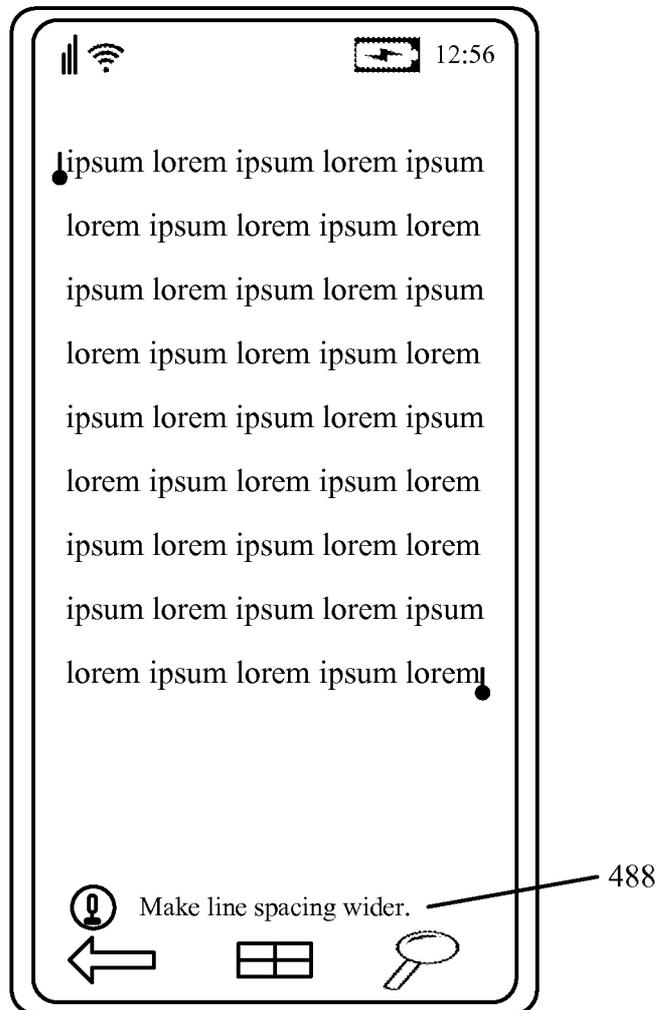
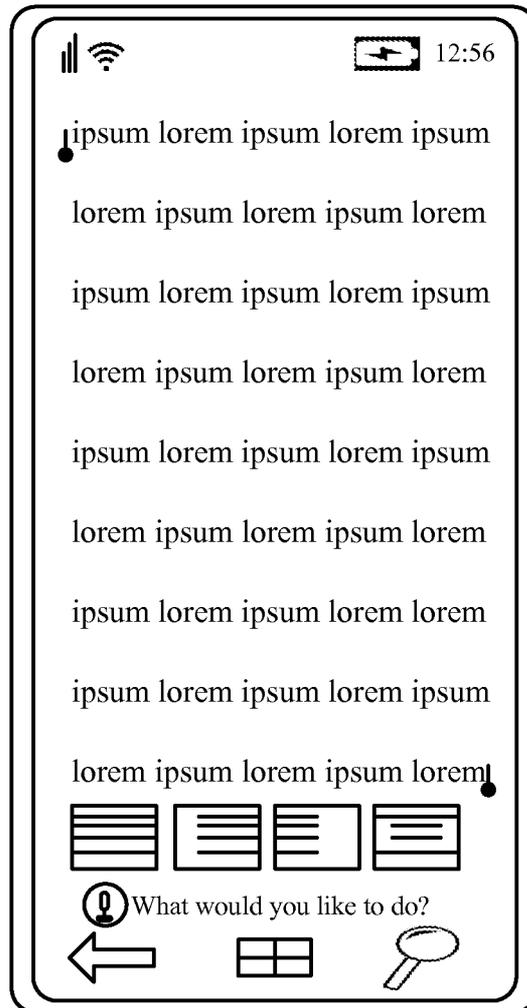


FIG. 5G



ⓘ "More"

FIG. 5H

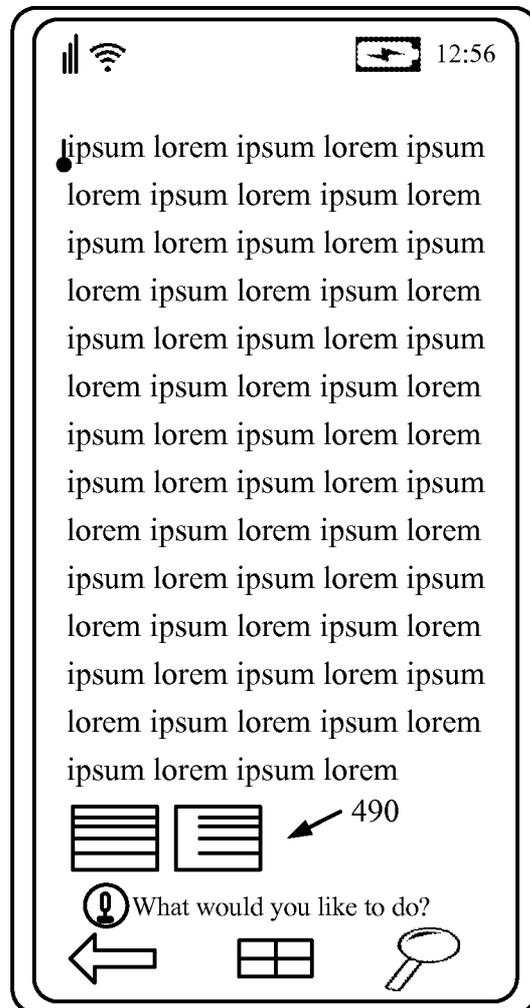


FIG. 5I

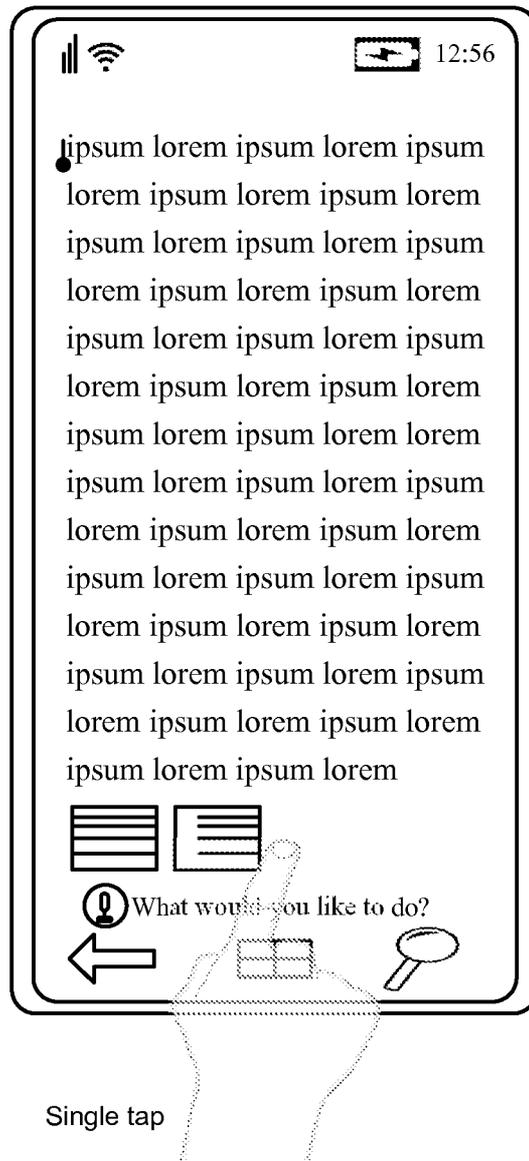


FIG. 5J

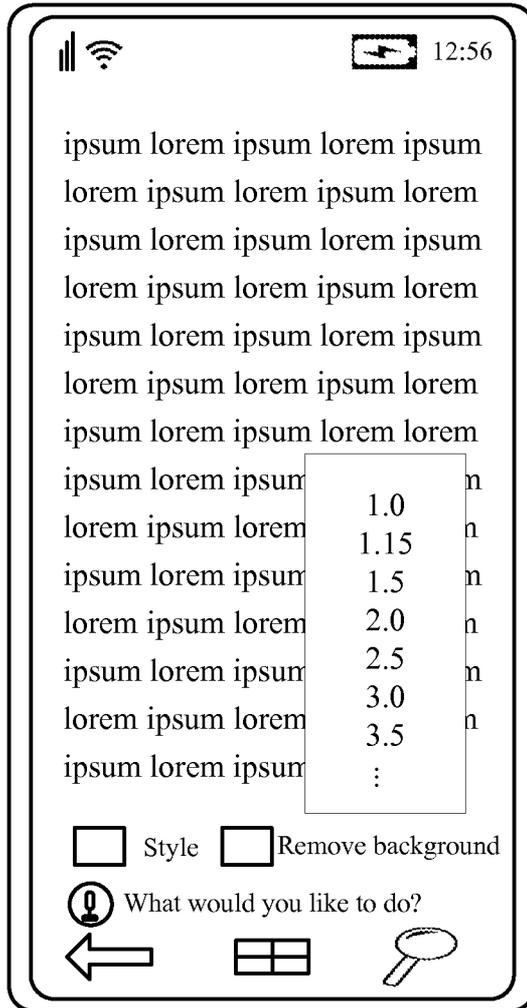
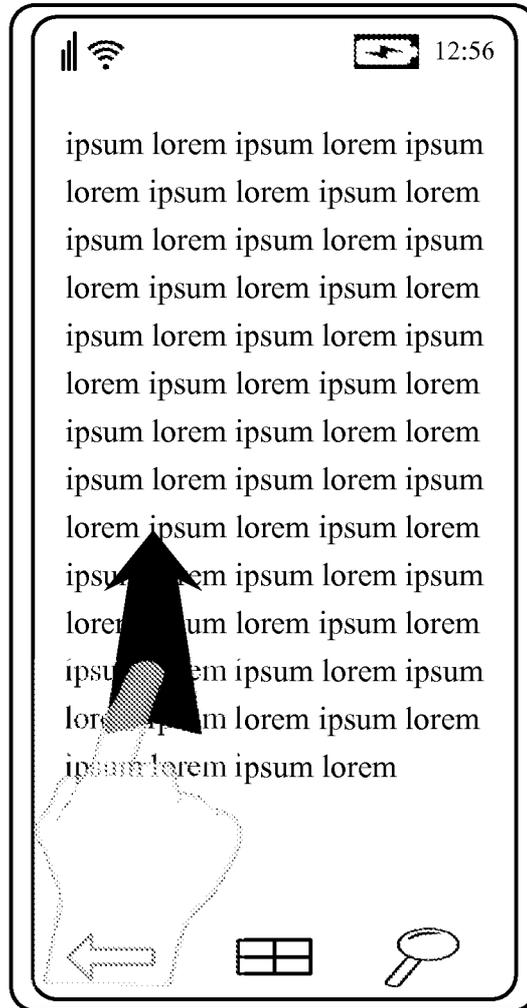
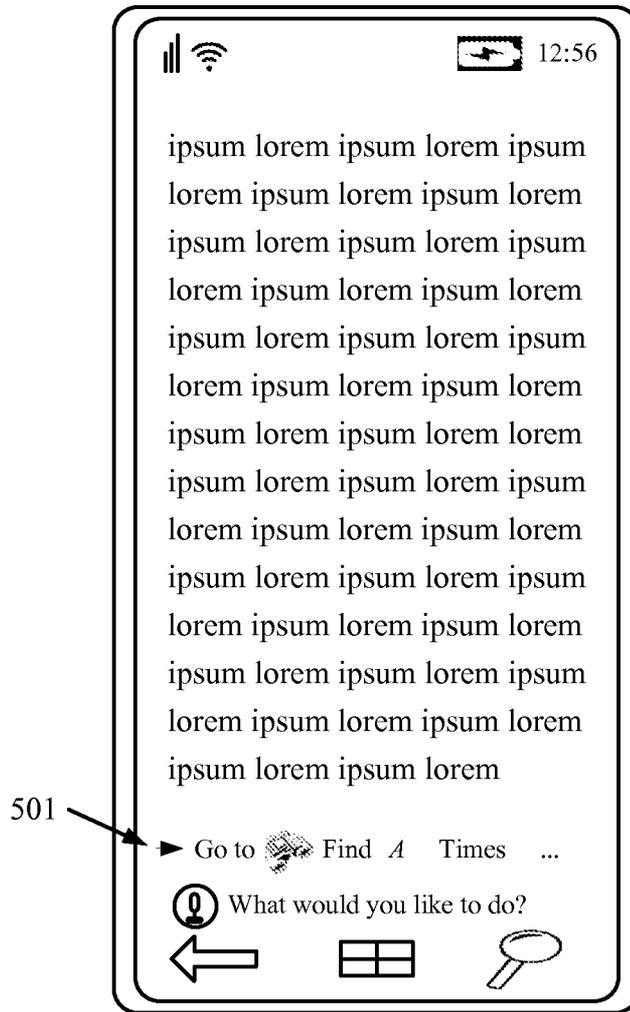


FIG. 5K



Scroll, scroll, scroll

FIG. 5L



Predictive UI appears

FIG. 5M

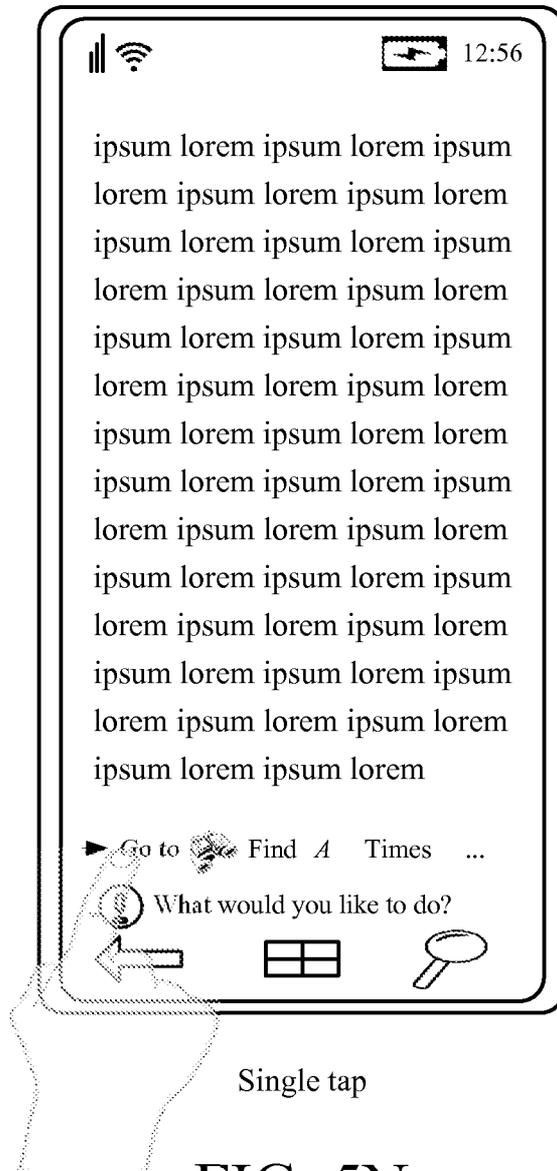


FIG. 5N

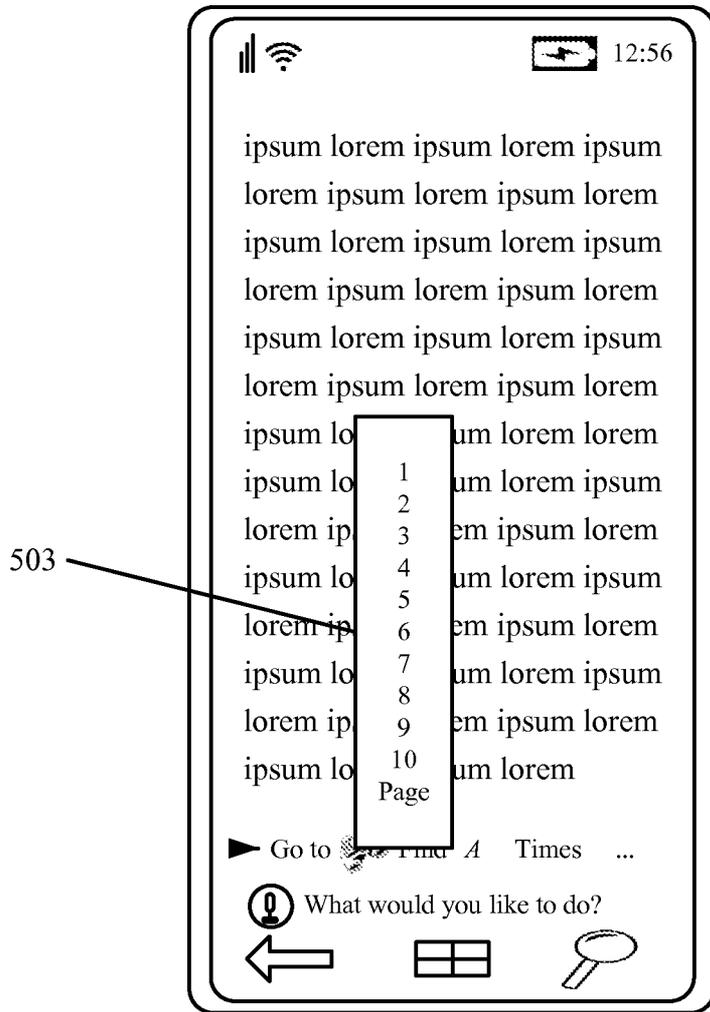


FIG. 50

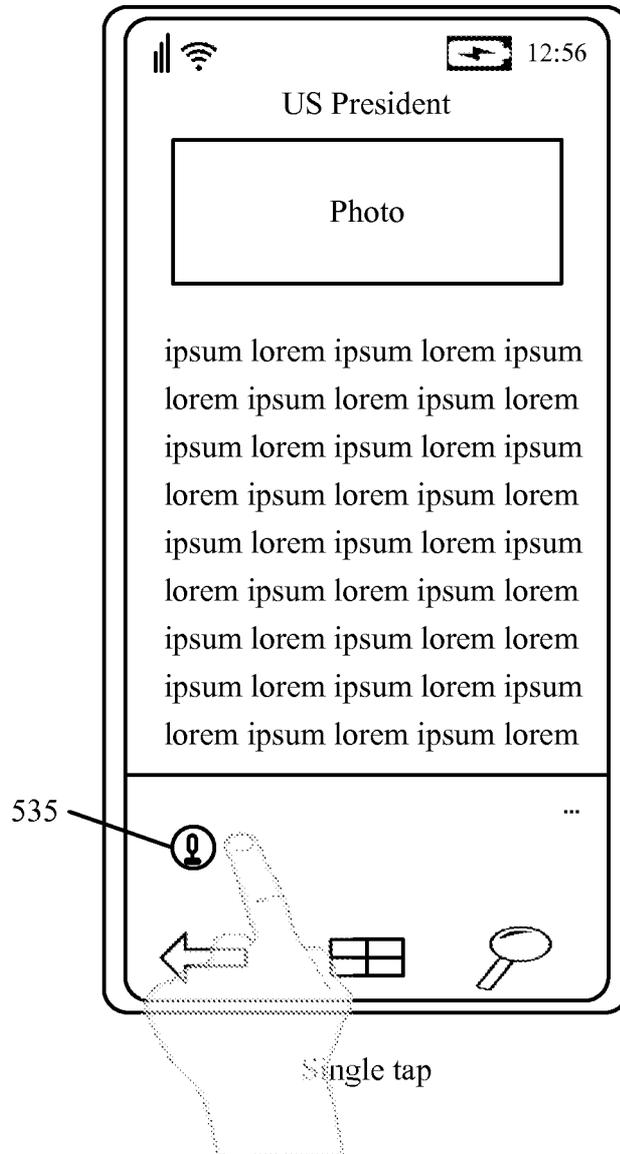
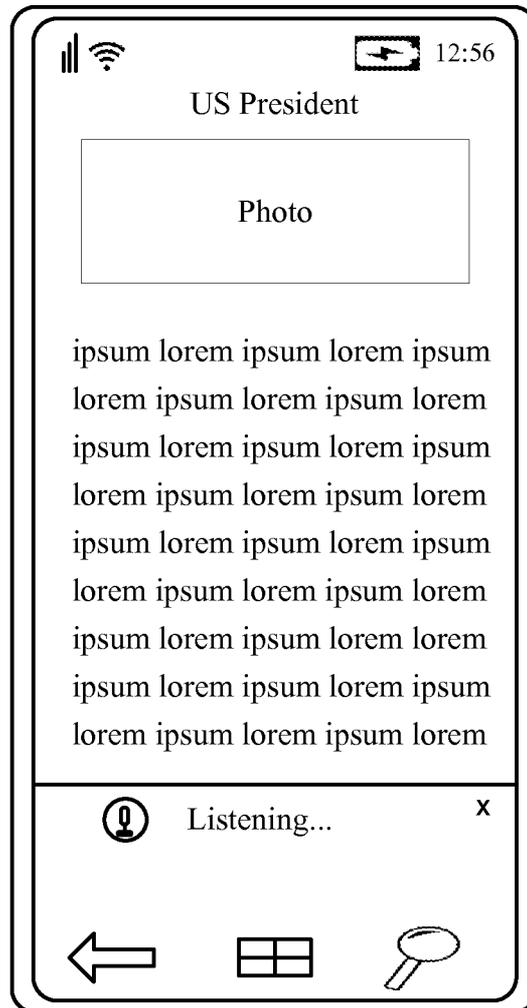


FIG. 5P



“I want to format
all my pictures.”

FIG. 5Q

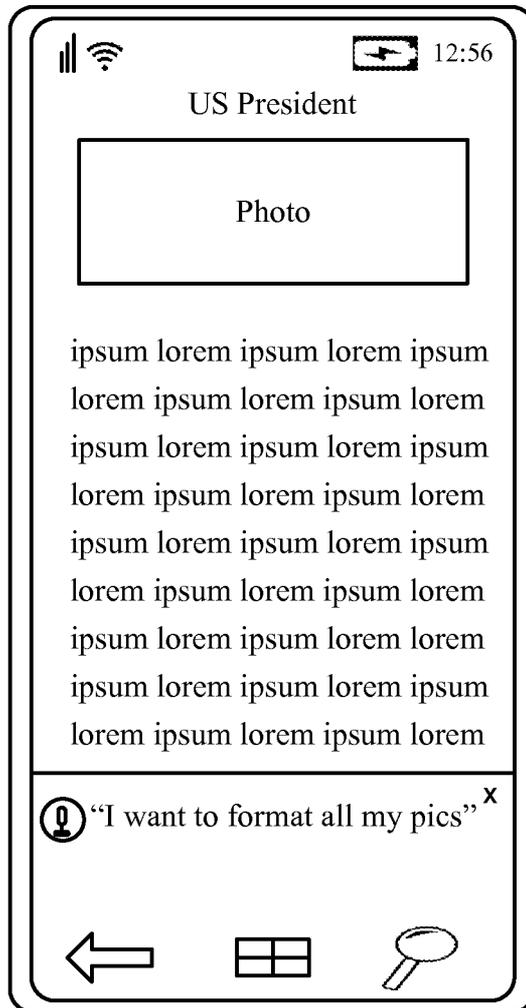


FIG. 5R

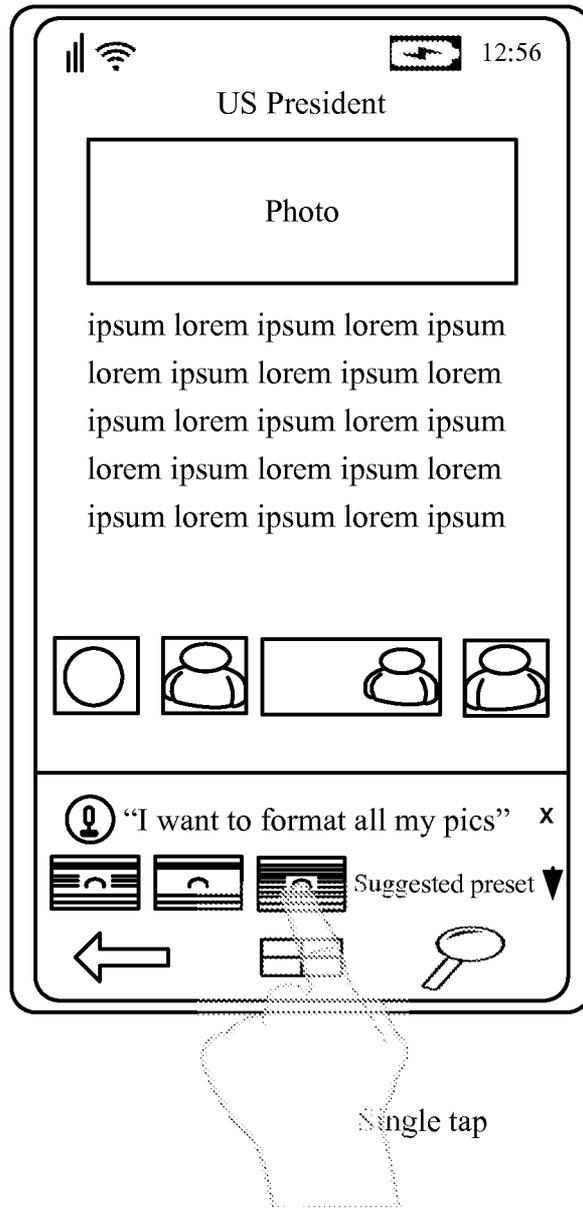


FIG. 5S

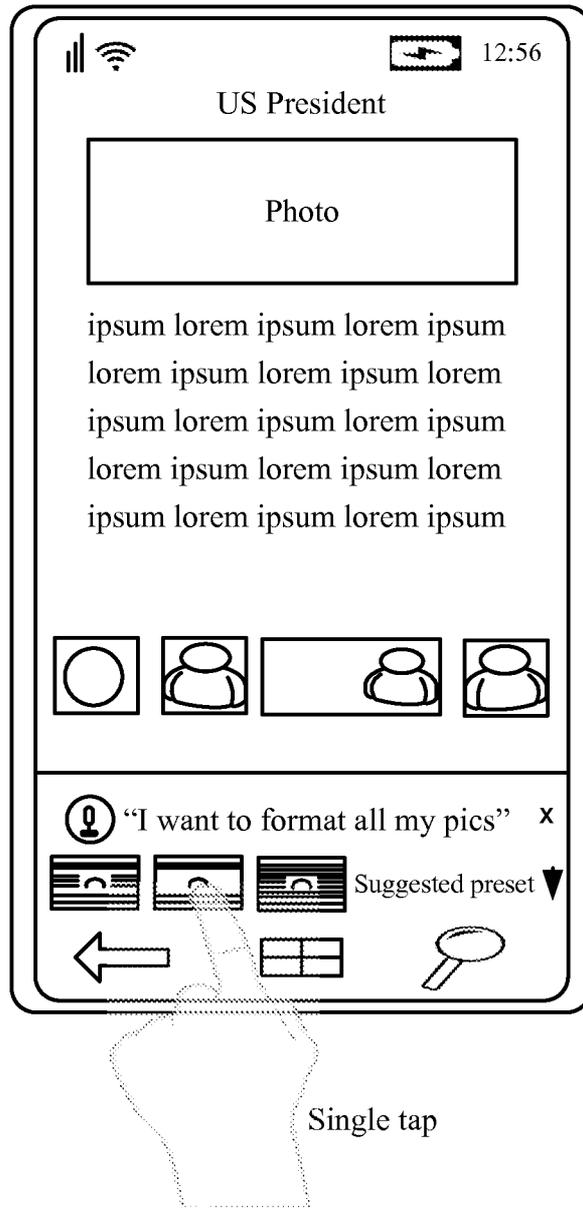
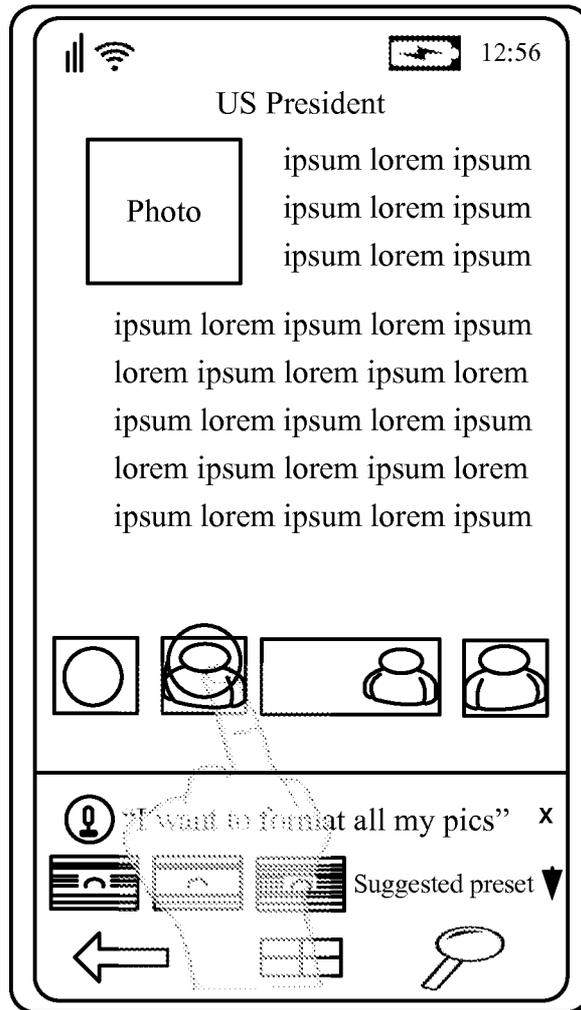


FIG. 5T



Tap and Hold

FIG. 5U

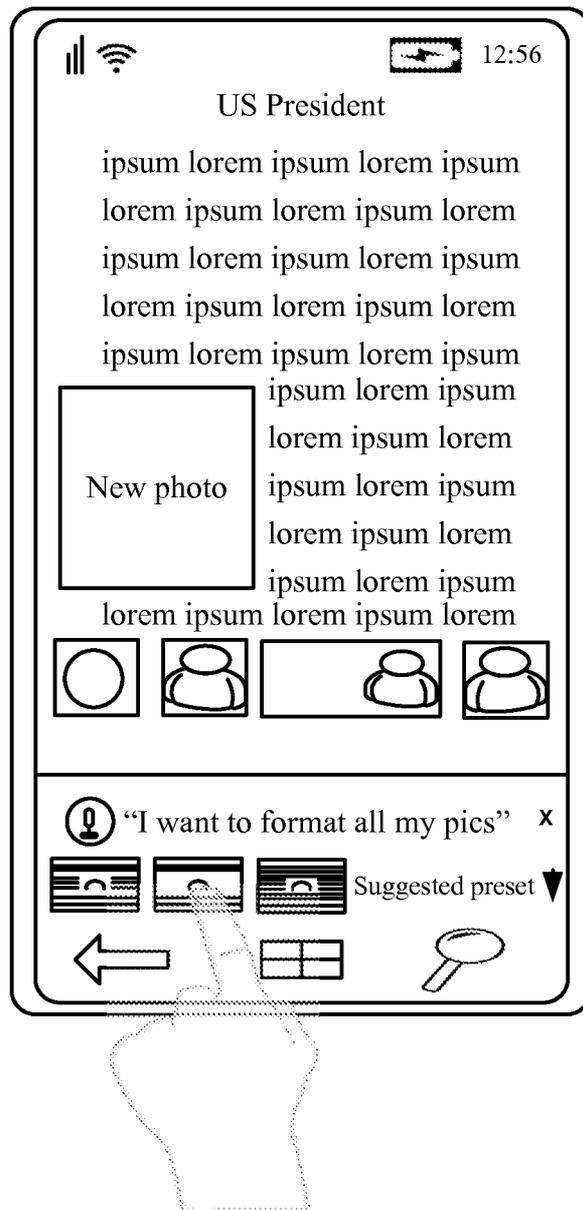


FIG. 5V

44/58

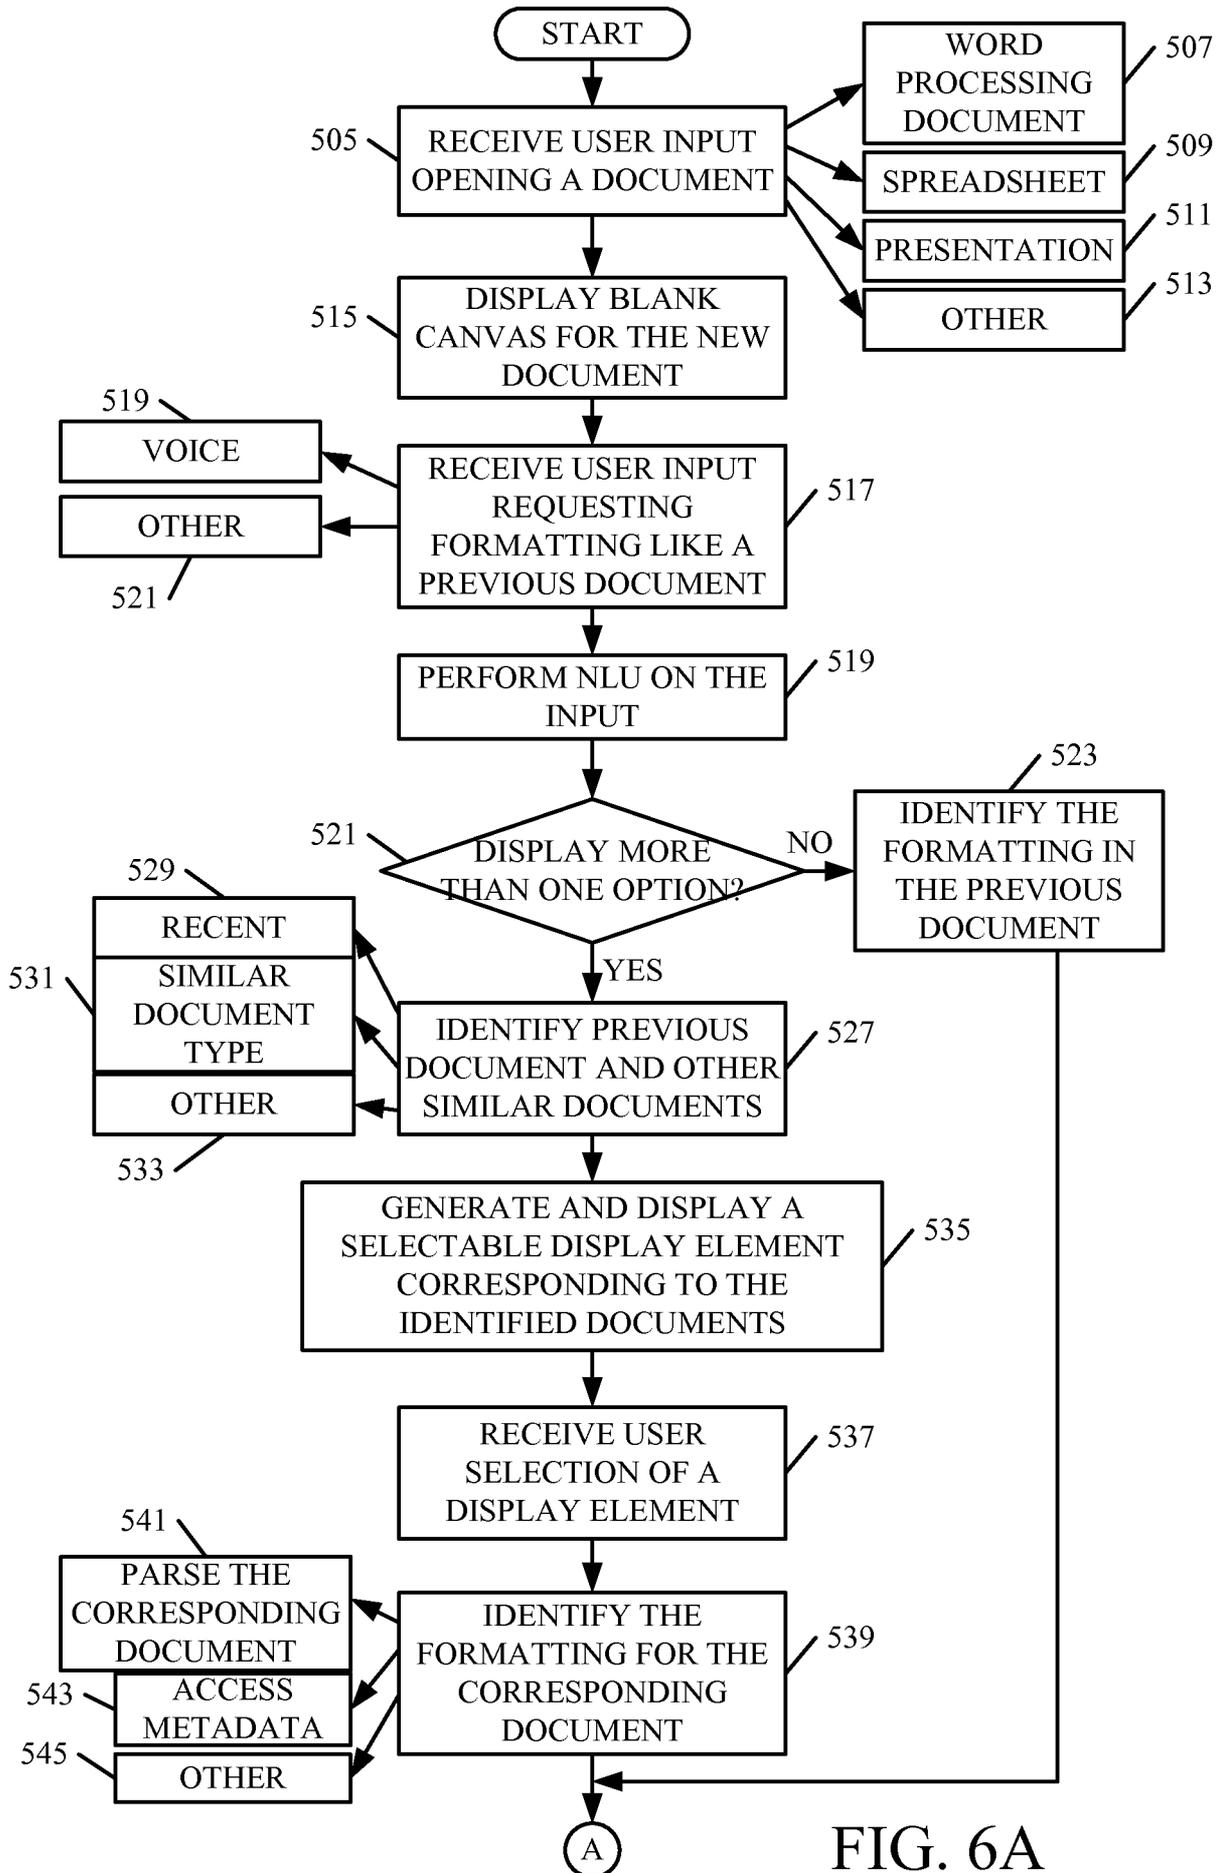


FIG. 6A

45/58

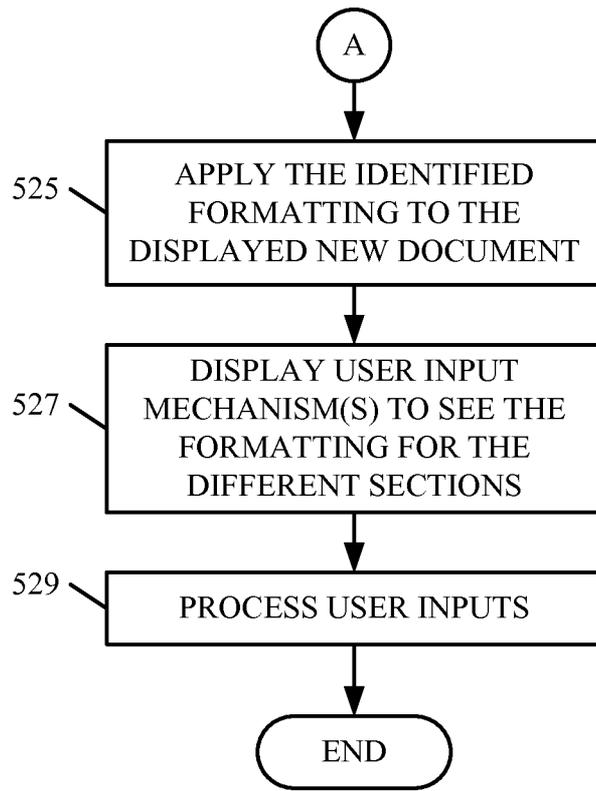


FIG. 6B

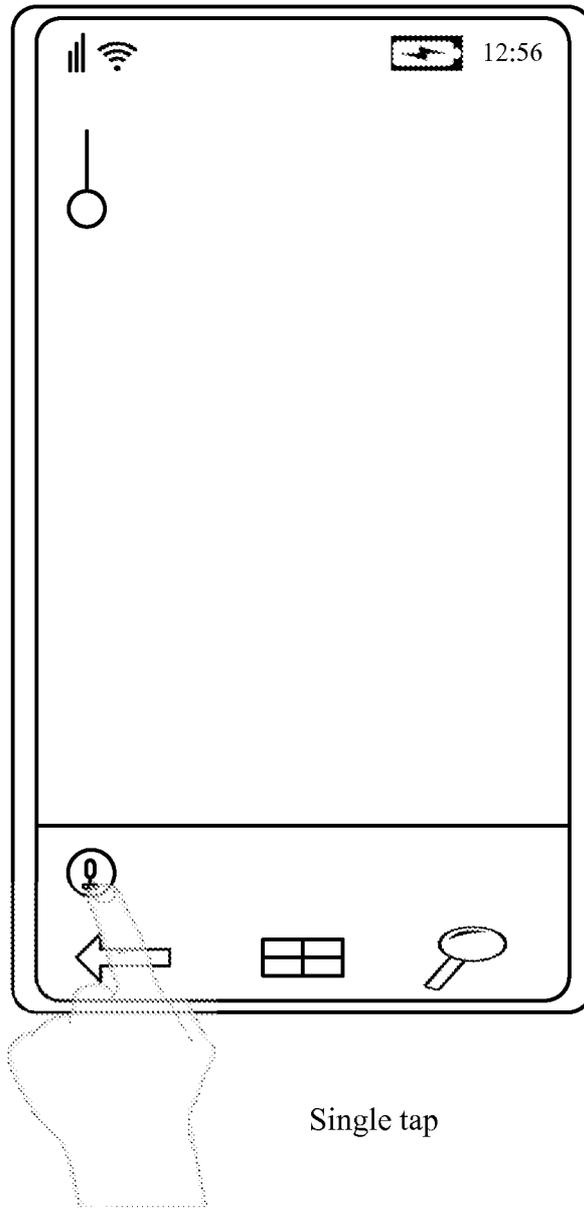
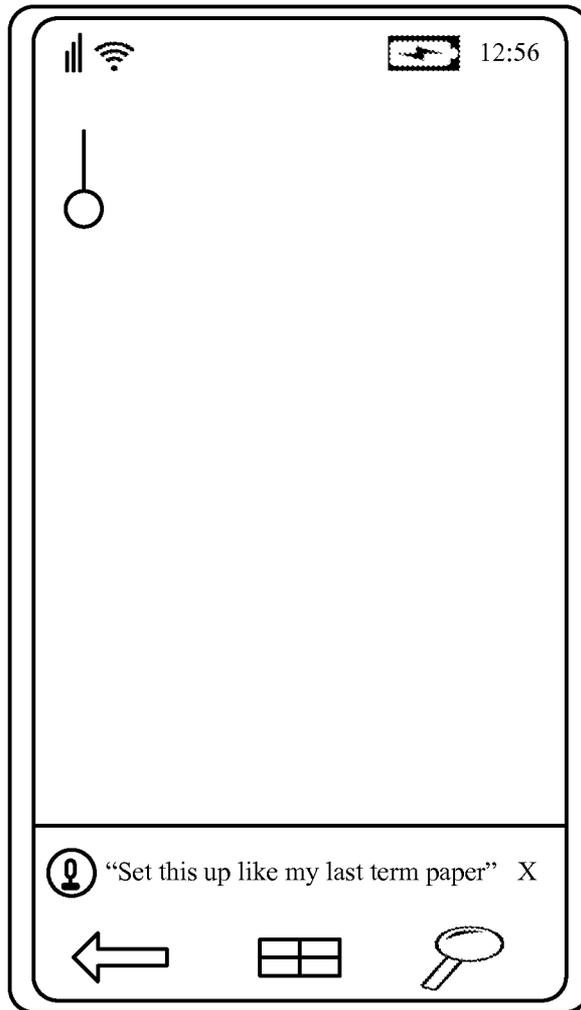
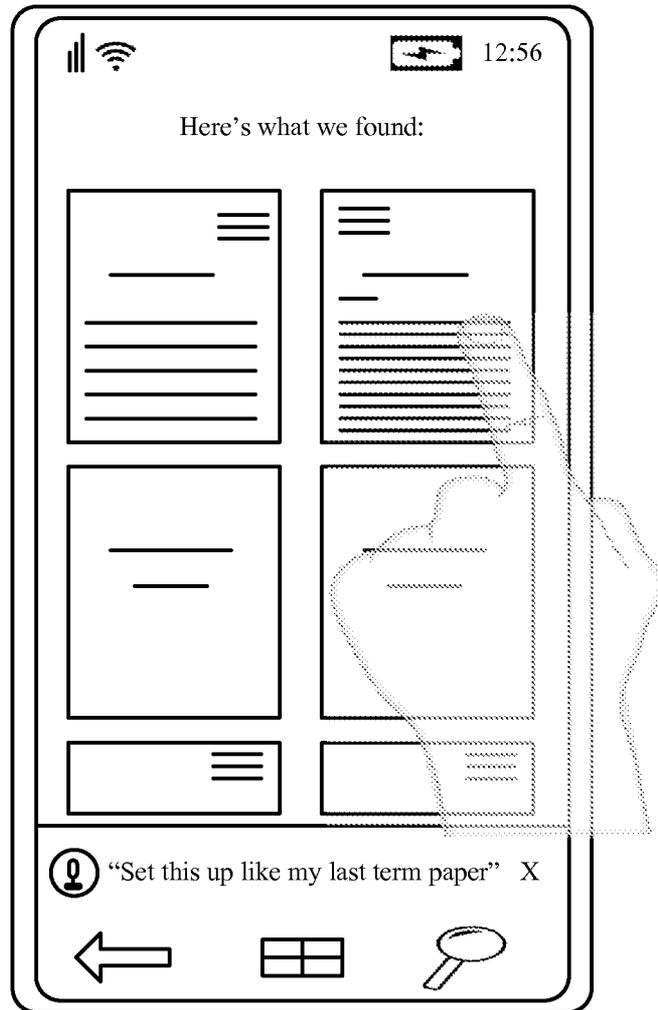


FIG. 6C



 "Set this up like my last term paper"

FIG. 6D



Ⓔ “Set this up like my last term paper”

FIG. 6E

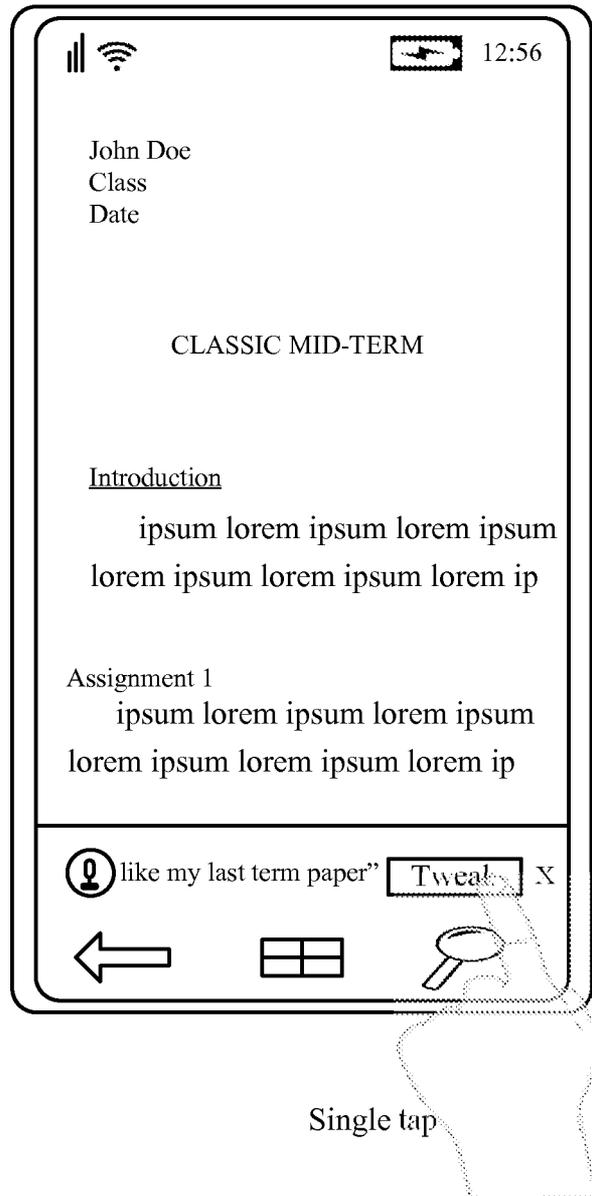
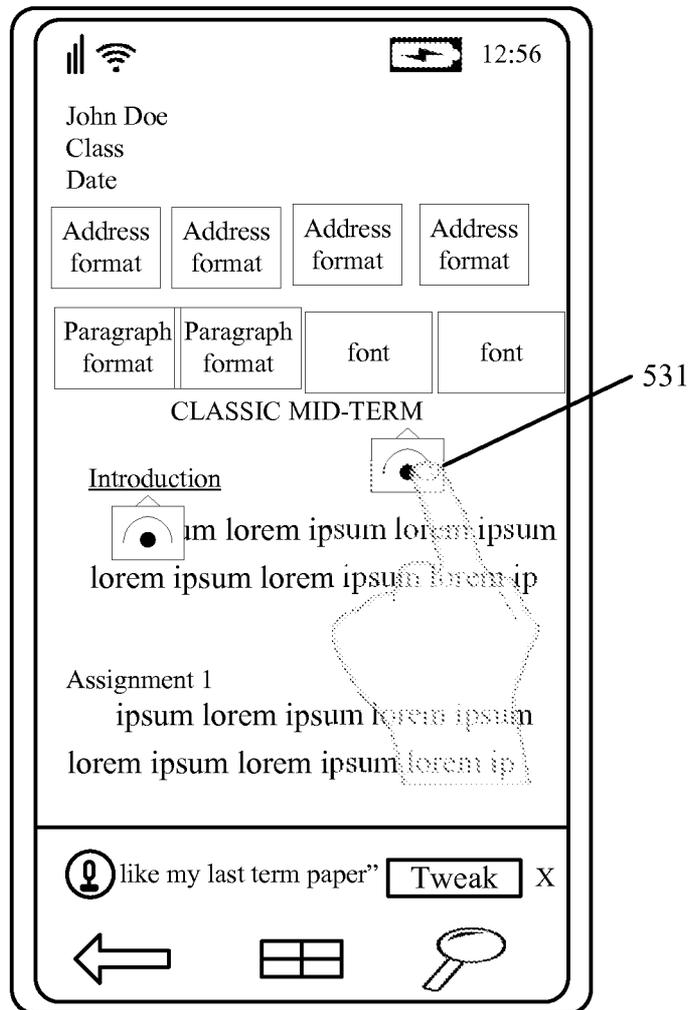
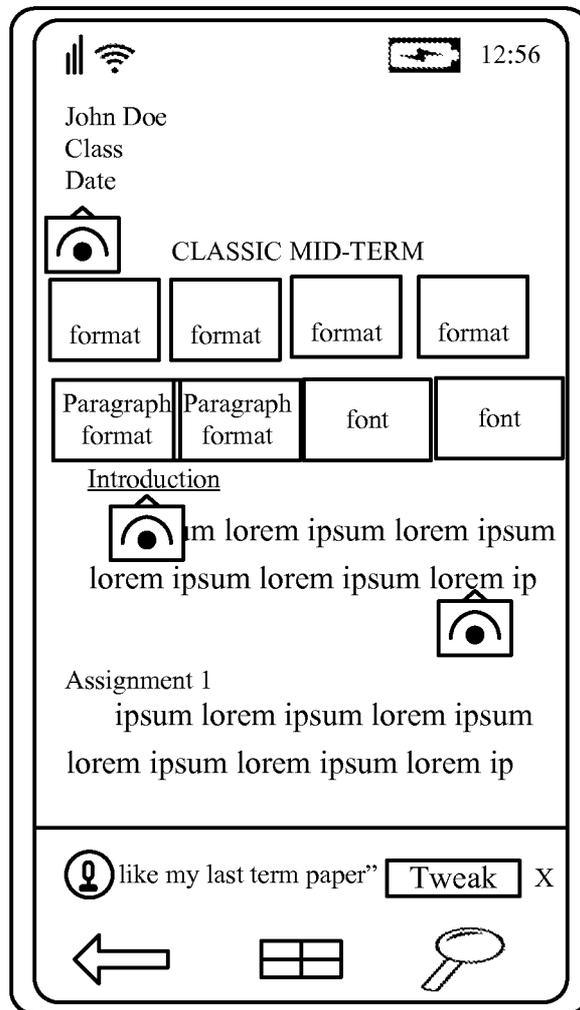


FIG. 6F



Single tap

FIG. 6G



Single tap

FIG. 6H

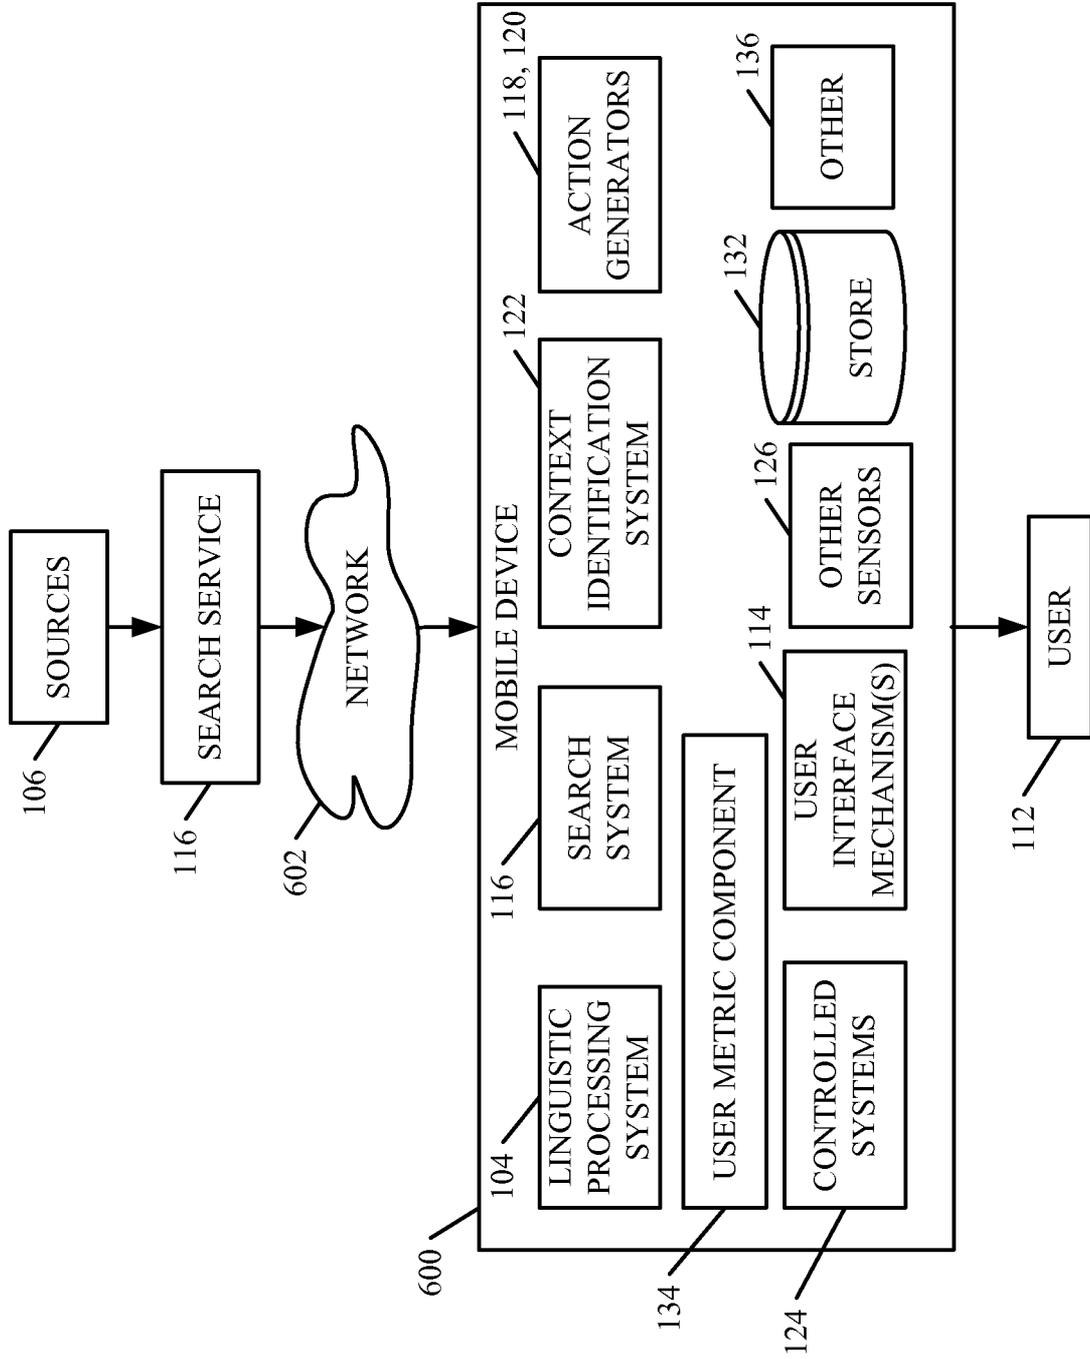


FIG. 7

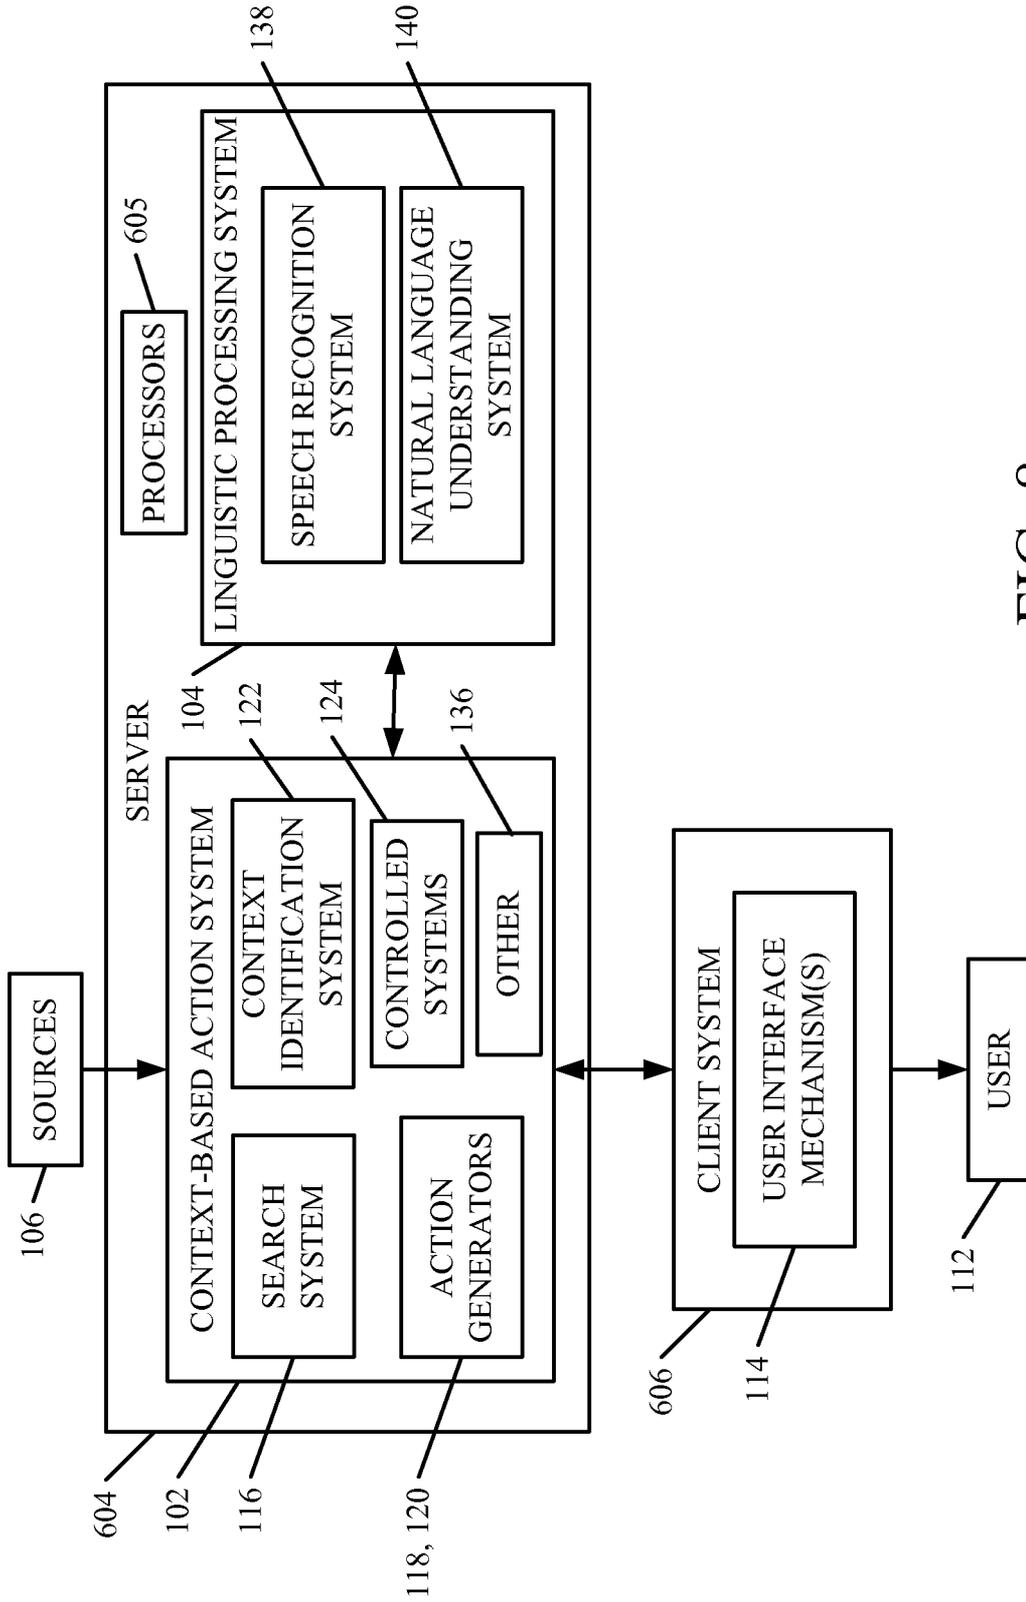


FIG. 8

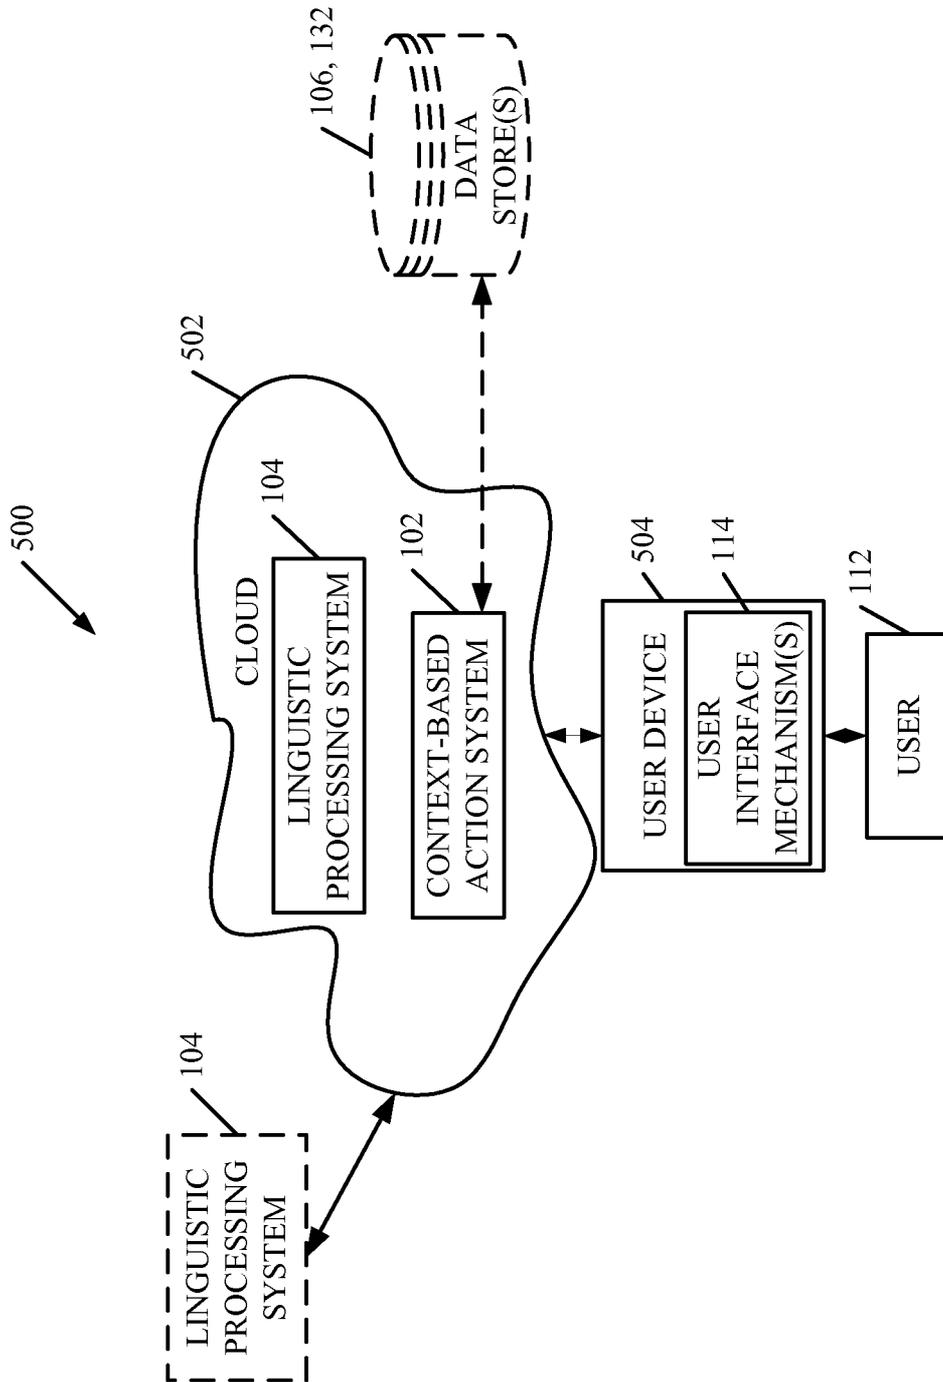


FIG. 9

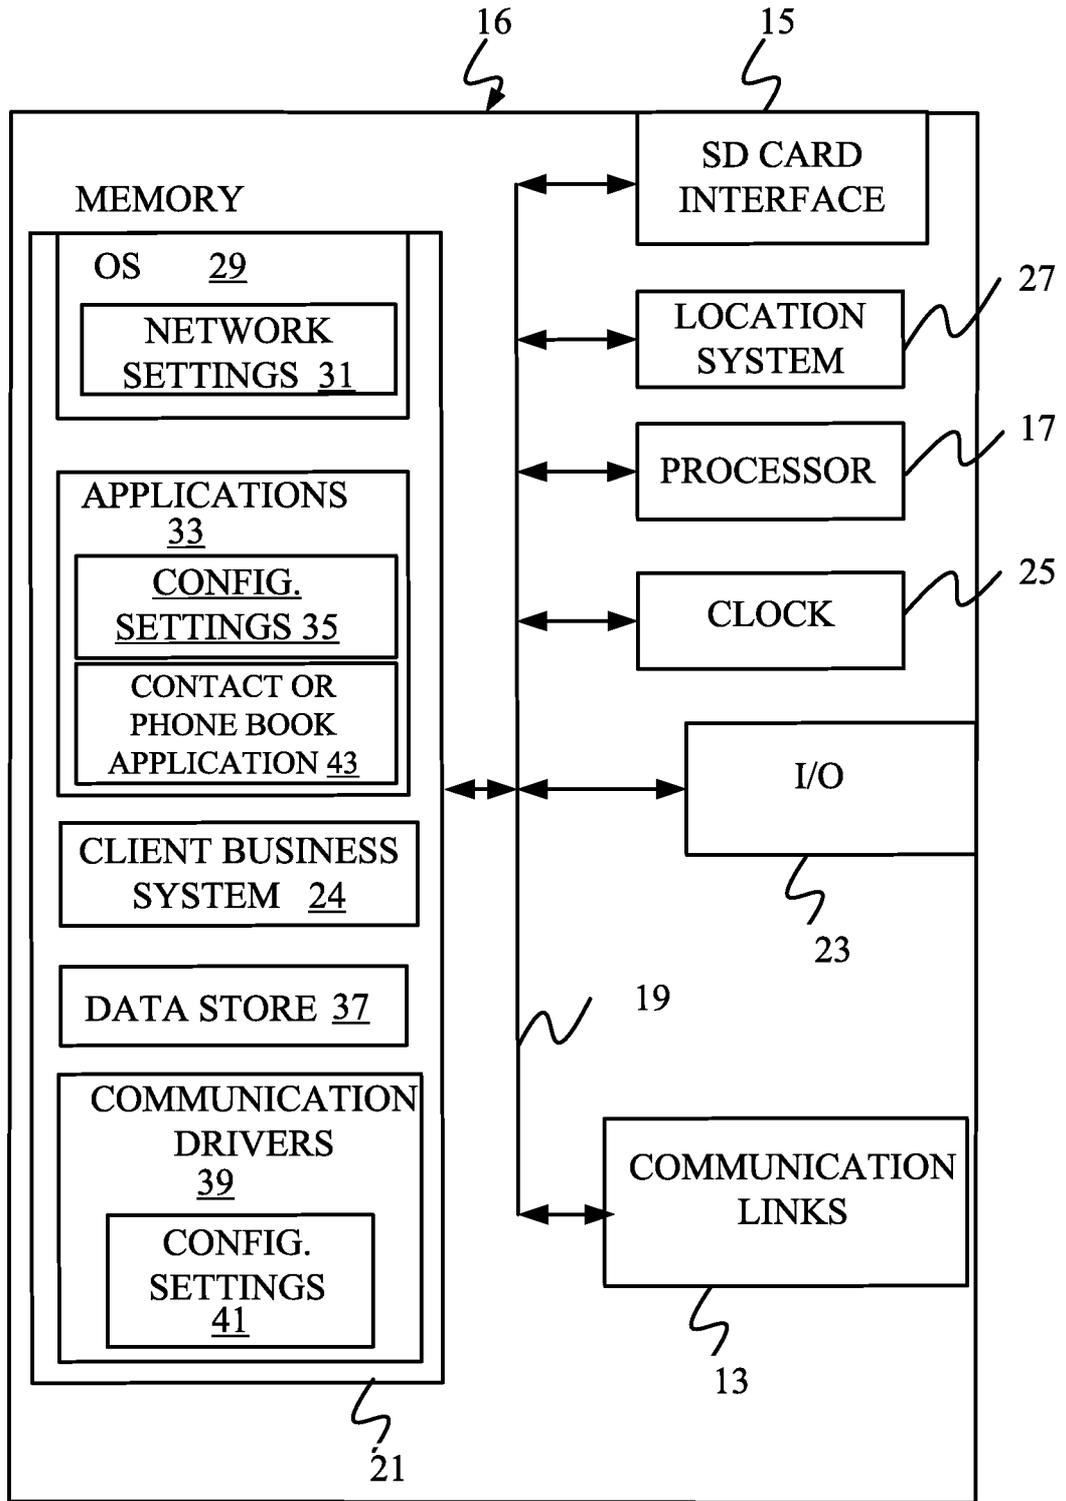


FIG. 10

56/58

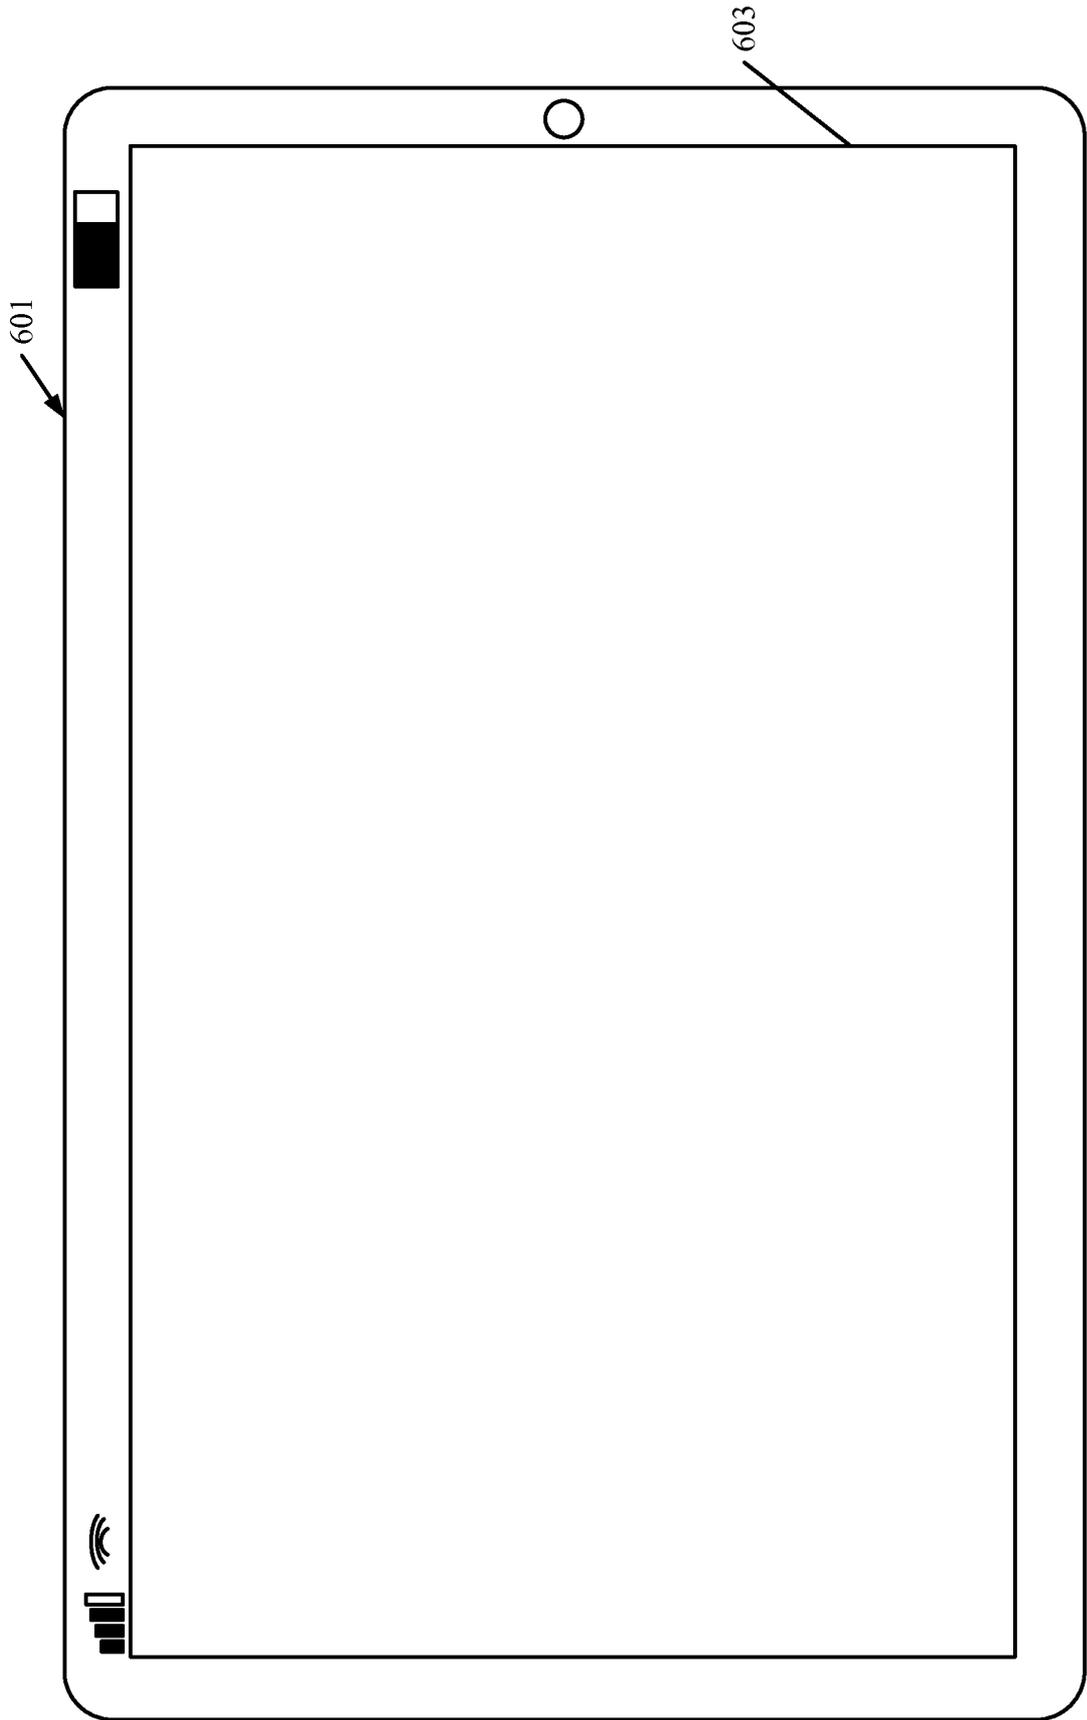


FIG. 11

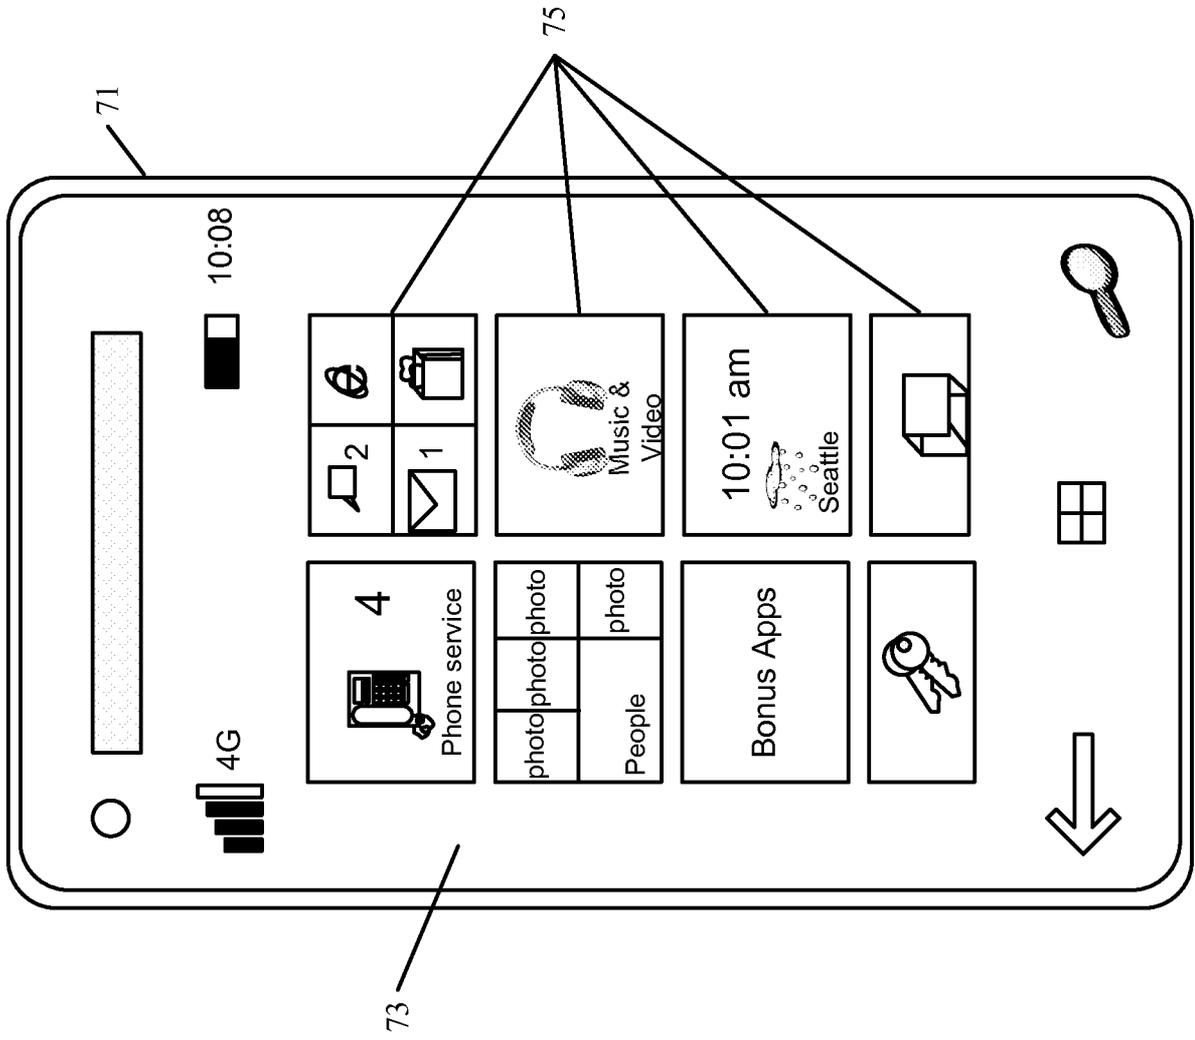


FIG. 12

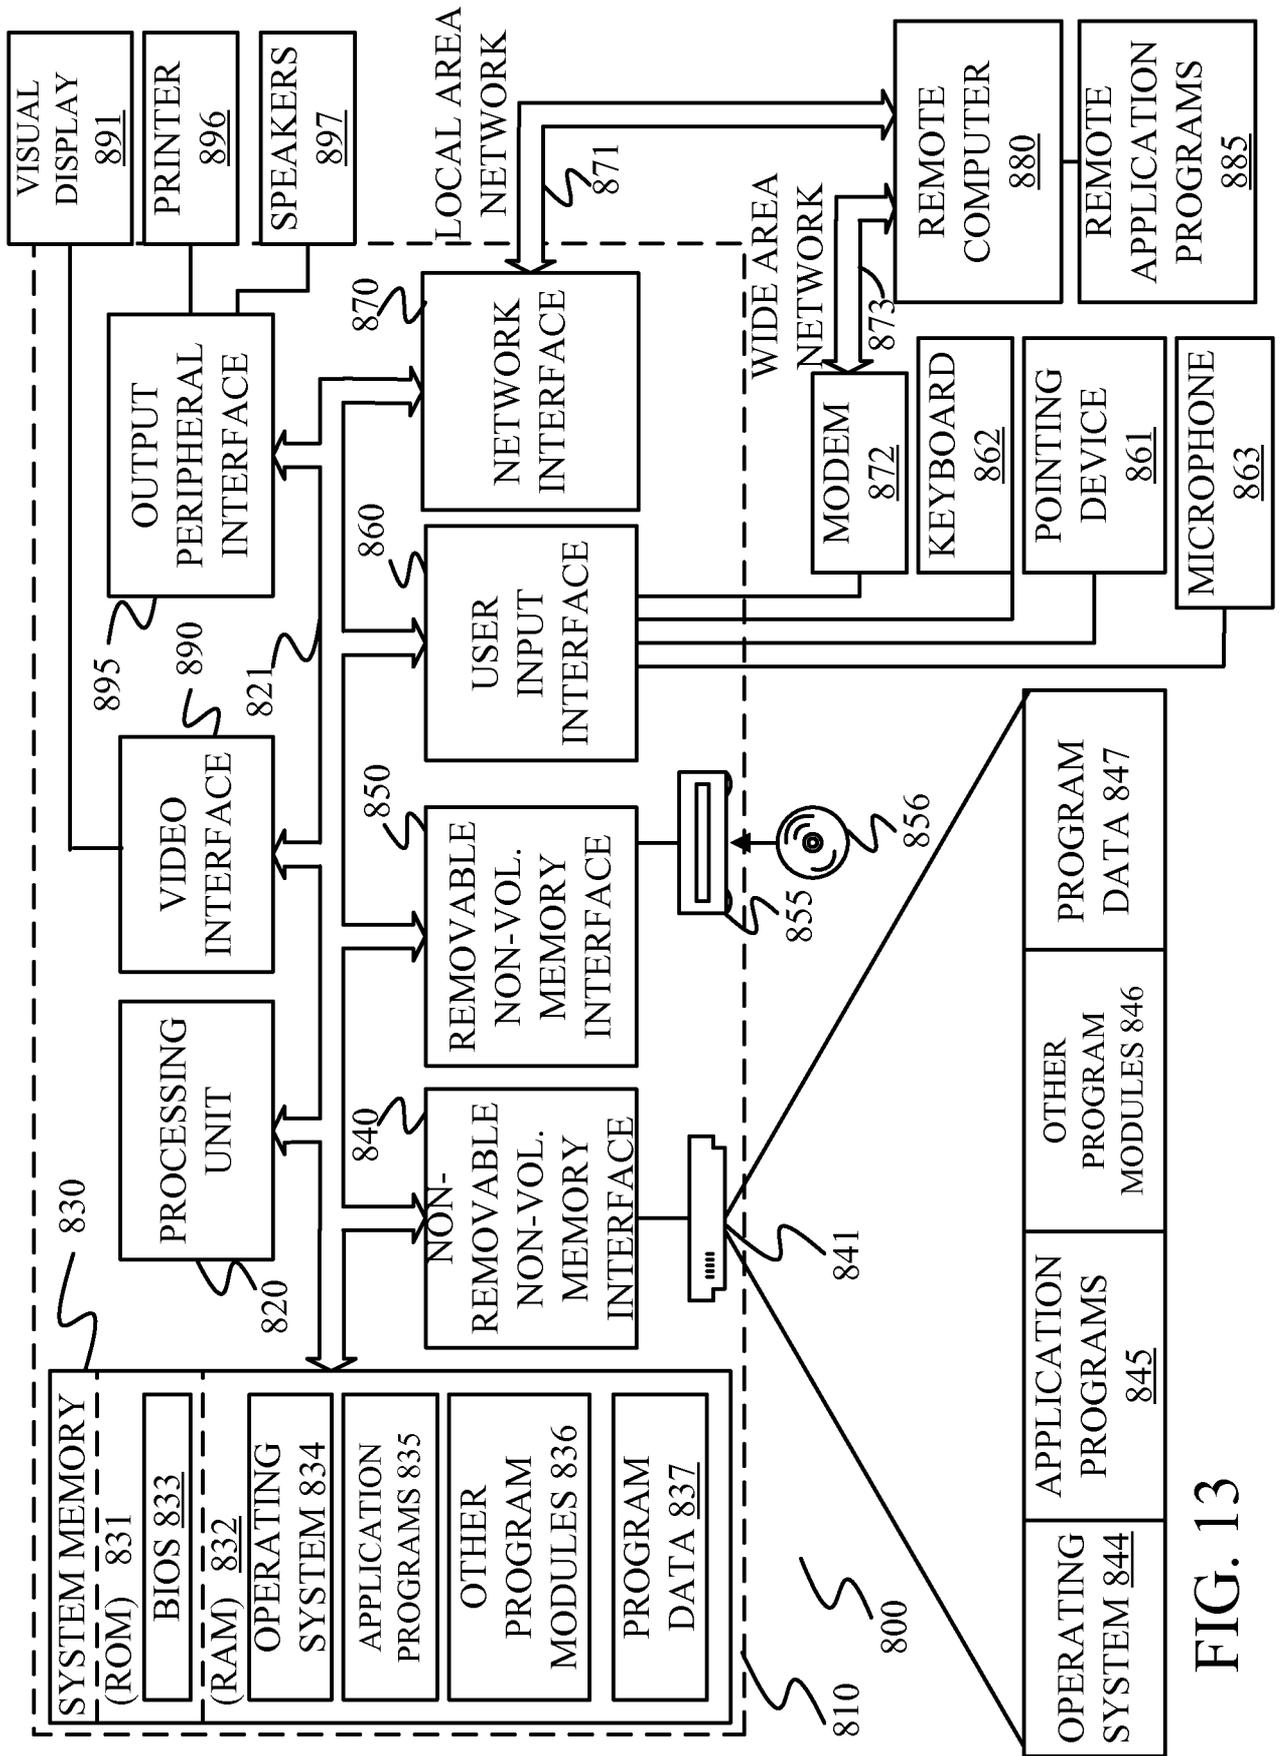


FIG. 13

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2015/059355

A. CLASSIFICATION OF SUBJECT MATTER
 INV. G06F9/44 G10L15/22 G06F17/30 G06F3/16
 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)
 G06F G10L

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, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2012/016678 A1 (GRUBER THOMAS ROBERT [US] ET AL) 19 January 2012 (2012-01-19) paragraphs [0127] - [0133] paragraphs [0139] - [0146] paragraph [0412] paragraphs [0598] - [0604] figures 29,35,40 paragraphs [0086], [0087], [0091] -----	1-15
X	MARTI A HEARST: "'Natural' search user interfaces", COMMUNICATIONS OF THE ACM, vol. 54, no. 11, November 2011 (2011-11), pages 60-67, XP058004857, ISSN: 0001-0782, DOI: 10.1145/2018396.2018414 section titled "Speech Input" and Figures 1 and 2 at page 62 -----	1-15

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 13 January 2016	Date of mailing of the international search report 21/01/2016
--	--

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 Del Castillo, G
--	---

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2015/059355

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2012016678	A1	19-01-2012	
		AU 2011205426	A1 23-08-2012
		CA 2787351	A1 21-07-2011
		CA 2791791	A1 21-07-2011
		CA 2792412	A1 21-07-2011
		CA 2792442	A1 21-07-2011
		CA 2792570	A1 21-07-2011
		CA 2793002	A1 21-07-2011
		CA 2793118	A1 21-07-2011
		CA 2793248	A1 21-07-2011
		CA 2793741	A1 21-07-2011
		CA 2793743	A1 21-07-2011
		CN 102792320	A 21-11-2012
		EP 2526511	A2 28-11-2012
		GB 2490444	A 31-10-2012
		JP 2013517566	A 16-05-2013
		JP 2014222509	A 27-11-2014
		JP 2014222510	A 27-11-2014
		JP 2014222511	A 27-11-2014
		JP 2014222512	A 27-11-2014
		JP 2014222513	A 27-11-2014
		JP 2014222514	A 27-11-2014
		JP 2014222515	A 27-11-2014
		JP 2014222516	A 27-11-2014
		JP 2014222517	A 27-11-2014
		KR 20120120316	A 01-11-2012
		KR 20120136417	A 18-12-2012
		KR 20120137424	A 20-12-2012
		KR 20120137425	A 20-12-2012
		KR 20120137434	A 20-12-2012
		KR 20120137435	A 20-12-2012
		KR 20120137440	A 20-12-2012
		KR 20120138826	A 26-12-2012
		KR 20120138827	A 26-12-2012
		KR 20130000423	A 02-01-2013
		RU 2012135502	A 27-02-2014
		RU 2012144605	A 27-04-2014
		RU 2012144606	A 10-05-2014
		RU 2012144637	A 10-05-2014
		RU 2012144639	A 10-05-2014
		RU 2012144640	A 10-05-2014
		RU 2012144643	A 10-05-2014
		RU 2012144644	A 10-05-2014
		RU 2012144647	A 10-05-2014
		RU 2012144648	A 10-05-2014
		US 2012016678	A1 19-01-2012
		US 2012245944	A1 27-09-2012
		US 2013110505	A1 02-05-2013
		US 2013110515	A1 02-05-2013
		US 2013110518	A1 02-05-2013
		US 2013110519	A1 02-05-2013
		US 2013110520	A1 02-05-2013
		US 2013111348	A1 02-05-2013
		US 2013111487	A1 02-05-2013
		US 2013117022	A1 09-05-2013
		US 2013185074	A1 18-07-2013
		US 2013185081	A1 18-07-2013
		WO 2011088053	A2 21-07-2011