September 29, 2017

DISAMBIGUATING ASSISTANT TASKS USING SUPPLEMENTAL DATA

Google Inc.

DISAMBIGUATING ASSISTANT TASKS USING SUPPLEMENTAL DATA

ABSTRACT

A virtual, intelligent, or computational assistant (e.g., also referred to simply as an "assistant") is described that relies on supplemental data (e.g., contextual information, user information, etc.) to decipher, and determine actions to perform based on, otherwise ambiguous user commands. With explicit permission from a user, the assistant may access a user's location history, calendar, e-mail, messages, past assistant interactions, contacts, photos, search history, sensor data, and other contextual or user information to interpret user statements. The supplemental data can be stored locally on a device that is executing the assistant or in a cloud computing environment that is accessible to the assistant from the device. This way, the assistant can understand, and perform actions in response to, more complex, compound, and often ambiguous commands even if the commands are generalized, not specific, and/or could be interpreted in different ways, depending on the context or location in which the user provides the command.

DESCRIPTION

Virtual, intelligent, or computational assistants (e.g., also referred to simply "assistants") execute on counter-top computing devices, mobile phones, automobiles, and many other types of computing devices. Assistants output useful information, respond to user queries, or otherwise perform certain operations to help users complete real-world and/or virtual tasks. The usefulness of an assistant may depend on what information the assistant already knows about its users or what information the assistant has access to.

The example system shown in FIG. 1 provides an assistant architecture that that relies on supplemental data, including contextual information and user information, when interpreting and performing actions in response to user commands. That is, with explicit permission from a user, the assistant may access a user's location history, calendar, e-mail, messages, past assistant interactions, contacts, photos, and other contextual or user information that is outside the assistant's typical control, to perform an action in response to a user command. In other words, the assistant may rely on supplemental data from external data sources to decipher, and take appropriate actions, in response to complex, compound and/or ambiguous user commands. The contextual information can be stored locally on a device that is executing the assistant or in a computing environment that is accessible to the assistant from the device. This way, the assistant can understand, and perform actions in response to more complex, compound, and often ambiguous commands even if the commands are generalized, not specific, and/or could be interpreted in different ways, depending on the context or location in which the user provides the command.

Further to the descriptions below, a user may be provided with controls allowing the user to make an election as to both if and when the assistant, the computing device, or the computing systems described herein can collect or make use of supplemental data (e.g., user information or contextual information about a user's social network, social actions or activities, profession, a user's preferences, or a user's current location), and if and when the user is sent content or communications from a server. In addition, certain data may be treated in one or more ways before it is stored or used, so that personally identifiable information is removed. For example, a user's identity may be treated so that no personally identifiable information can be determined for the user, or a user's geographic location may be generalized where location information is

obtained (such as to a city, ZIP code, or state level), so that a particular location of a user cannot be determined. Thus, the user may have control over what supplemental data is collected about the user, how that supplemental data is used, and what supplemental data is provided to the user.

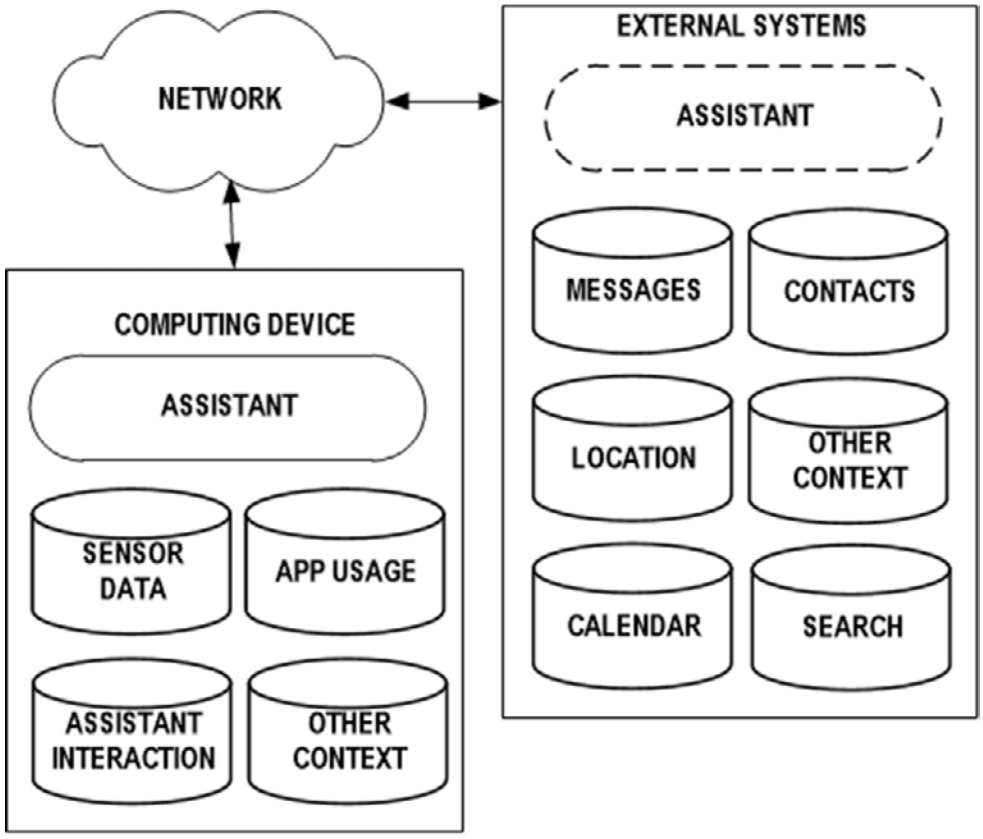

FIG. 1

The system of FIG. 1 includes one or more external systems and a computing device communicating across a network to provide an assistant service that maintains and has access to user information. The network of FIG. 1 represents a combination of any one or more public or private communication networks, for instance, television broadcast networks, cable or satellite networks, cellular networks, Wi-Fi networks, broadband networks, and/or other type of network for transmitting data (e.g., telecommunications and/or media data) between various computing devices, systems, and other communications and media equipment.

The computing device represents any type of computing device that is configured to execute an assistant and communicate on a network. The external systems and computing device can be one or more personal computing devices. In some examples, the computing device and external systems may be shared assets of multiple users. Examples of the computing device and the external systems include cloud computing environments, mobile phones, tablet computers, wearable computing devices, countertop computing devices, home automation computing devices, laptop computers, desktop computers, televisions, stereos, automobiles, and all other type of mobile and non-mobile computing device that is configured to execute an assistant.

The computing device and external systems may store or provide access to supplemental data including personal information about users. Examples of personal information include: sensor data, calendars, location histories, search histories, messages, e-mails, preferences, notes, lists, contacts, other communications, interests, application usage data, past assistant interactions, etc. After receiving explicit permission from a user, the computing device and external systems may store the supplemental data and enable an assistant, or other applications, executing at the computing device and external systems to access the supplemental data.

The external systems and the computing device treat the supplemental data so the supplemental data is protected, encrypted, or otherwise not susceptible to unauthorized access or use. The supplemental data may be stored locally at the computing device and/or remotely (e.g., in a cloud computing environment provided by the external systems and which is accessible via the network of FIG. 1).

The computing device includes an assistant that executes across the external systems and the computing device to provide assistant services to users of the computing device. Examples of assistant services include: setting up reminders, creating calendar entries, booking travel,

online ordering, sending messages or other communications, controlling televisions, lights, thermostats, appliances, or other computing devices, providing navigational instructions, or any other conceivable task or operation that may be performed by an assistant. The assistant relies on the supplemental data stored on the computing device or the external systems when interpreting, and determining answers to, user queries.

As one example, the assistant may use location to disambiguate a shopping request. For example, when a user of the computing device is in their laundry room, the user may tell the assistant "get more soap and dryer things." The assistant may determine the command to be a shopping request but based only on the request, the assistant may not know what type of soap to buy or what "dryer things" are. For instance, the soap could be dish soap, shower soap, or any other variety of soap, and "dryer things" could easily be towels, napkins, or dryer sheets.

To disambiguate the request, the assistant relies on supplemental data. Using, as one example, location information, the assistant may first determine that the user is in his or her laundry room. In some cases, using information obtained from the clothes dryer (e.g., one of the external systems, in this case a network connected appliance), the assistant may determine that the door to the dryer was recently opened or closed. The assistant may infer based on the location of the user and/or the opening and closing of the dryer door, that the user is doing laundry. In response, the assistant may interpret the command to procure "soap" as likely being a command to procure laundry detergent as opposed to other types of soap. Likewise, given that the user is likely doing his or her laundry in the laundry room, the assistant may infer that "dryer things" must refer to "dryer sheets" as opposed to other items or articles that could be used for drying. These likelihoods may be reflected in ranking or otherwise comparing the various options to arrive at a single interpretation of user intent. The system may check with the user to

determine if the single interpretation is correct. For example, the system may respond to the user's command with "I think you mean laundry soap and dryer sheets. I'll add those to the shopping list." The user is then presented with an opportunity to correct the assistant's interpretation (and such correction may be referenced in future cases to further aid the assistant in interpreting ambiguous user commands). The assistant may communicate with an online retailer to place an order for laundry soap and dryer sheets. In some implementations, the assistant may add the items to a shopping list rather than necessarily purchasing the items immediately.

As another example, a user of the computing device may be at a work location and command the assistant to "order pizza for lunch". The assistant may determine the location of the computing device to be at the user's work location and/or determine that the time of day of during the user's regular work hours. Based on the location and/or time, the assistant may infer that the user most likely wants the assistant to order the user's regular pizza order from the pizza place nearest the user's work location and may place the order.

Alternatively, the user may be at a home location and provide the assistant with the same command to "order pizza for lunch". The assistant may determine the location of the computing device to be at the user's home location and/or determine that the time of day is outside the user's work hours. Based on the location and/or time, the assistant may infer that the user most likely wants the assistant to order the user's typical family pizza order from the pizza place nearest the user's home and may place the order.

As yet another example, a user of the computing device may be driving in her automobile when a check engine light comes on the dash that indicates a front headlight bulb needs replacing. The user may tell the assistant to "navigate me to a store to buy a replacement light

bulb." Using movement information obtained from the computing device, or information obtained from a Bluetooth receiver in the automobile (e.g., an external system), the assistant may first determine that the user is driving in her automobile. In some cases, using further information obtained from the automobile itself (e.g., an infotainment system or other component of the vehicle) the assistant may determine that error code for a bad lightbulb has been thrown by the automobile. The assistant may infer based on the movement information, Bluetooth information, and/or information from the automobile, that the user is commanding the assistant to navigate to a store to buy a lightbulb for her vehicle and not some generic home bulb. The assistant may execute a search for lightbulbs for her vehicle and identify a nearby auto parts store that carries the bulb. The assistant may immediately begin navigating the user to the nearby auto parts store and may even place an order for the bulb to be held at customer pickup.

In any of the above examples, the assistant may at some point request that the user confirm an assumption. That is, before placing an online order for an item or taking some other action, the assistant may request that the user confirm he or she wants the assistant to perform that particular action. For example, before navigating to the auto parts store, the assistant may tell the user "You want me to navigate you to an auto parts store to buy a light bulb for the front headlight that just went out, correct?"

In other examples, the assistant may disambiguate contacts that share the same or similar name when performing a task. For example, the user may tell the assistant "send flowers to Kelly". The user may be traveling on a long business trip, his boss may be a person named Kelly, and coincidentally, the user's wife may also be named Kelly. The assistant may have access to the user's calendar and recognize that the user's wedding anniversary is the next day (when the user is still traveling). Given the potentially personal or sensitive nature of the

command, the assistant may infer that the user likely wants to send flowers to his wife Kelly for their anniversary and not to his boss Kelly. Again, this likelihood may manifest in ranking the various contacts named Kelly in the user's contacts list by the identified attribute, such as family relationship (or other attributes as might for example be associated with entities in an entity graph), and selecting the highest ranked Kelly. Of course, the assistant may request that the user confirm the assistant's assumption of which Kelly he is referring to.

In another example, (at a different date and time that does not coincide with the user's anniversary) the user may again tell the assistant "send flowers to Kelly". Again, the user may be traveling on a long business trip, his boss may be a person named Kelly, and coincidentally, the user's wife may also be named Kelly. The assistant may have access to the user's messages or e-mail and recognize that his boss Kelly recently sent an e-mail to all her reports that she just had a baby. The assistant may infer that the user likely wants to send flowers to his boss Kelly as congratulations for the baby. Of course, the assistant may request that the user confirm the assistant's assumption of which Kelly he is referring to.

By relying on supplemental data, the assistant can understand, and perform actions in response to, more complex, compound, and often ambiguous commands even if the commands are generalized, not specific, and could be interpreted in different ways, depending on the context or location in which the user provides the command. The above examples are just some use cases for the assistant architecture shown in FIG. 1, the assistant architecture has many other applications and use cases.