

(51) International Patent Classification:
G06Q 30/00 (2012.01)(21) International Application Number:
PCT/US2017/056116(22) International Filing Date:
11 October 2017 (11.10.2017)

(25) Filing Language: English

(26) Publication Language: English


(30) Priority Data:
15/294,765 16 October 2016 (16.10.2016) US(63) Related by continuation (CON) or continuation-in-part (CIP) to earlier application:
US 15/294,765 (CON)
Filed on 16 October 2016 (16.10.2016)(71) Applicant: **EBAY INC.** [US/US]; 2145 Hamilton Avenue,
San Jose, California 95125 (US).(72) Inventors: **KALE, Ajinkya Gorakhnath**; 3190 Rubino Drive, Apt. 210, San Jose, California 95125 (US). **YANG, Fan**; 310 Crescent Village Circle, Unit 1316, San Jose, California 95134 (US). **WANG, Qiaosong**; 540 Crestlake Dr, San Francisco, California 94132 (US). **KIAPOUR, Mohammadhadi**; 2190 Beach St., Apt #106, San Francisco,California 94123 (US). **PIRAMUTHU, Robinson**; 7797 Hanson Dr., Oakland, California 94605 (US).(74) Agent: **SCHEER, Bradley** et al.; SCHWEGMAN, LUNDBERG & WOESSNER, P.A., P.O. Box 2938, Minneapolis, Minnesota 55402 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM,

(54) Title: PERSONAL ASSISTANT WITH VISUAL MULTI-TURN DIALOG

FIG. 1



(57) Abstract: Systems, methods, and computer program products for identifying a relevant candidate product in an electronic marketplace. Embodiments perform a visual similarity comparison between candidate product image visual content and input query image visual content, process formal and informal natural language user inputs, and coordinate aggregated past user interactions with the marketplace stored in a knowledge graph. Visually similar items and their corresponding product categories, aspects, and aspect values can determine suggested candidate products without discernible delay during a multi-turn user dialog. The user can then refine the search for the most relevant items available for purchase by providing responses to machine-generated prompts that are based on the initial search results from visual, voice, and/or text inputs. An intelligent online personal assistant can thus guide a user to the most relevant candidate product more efficiently than existing search tools.

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

Published:

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

PERSONAL ASSISTANT WITH VISUAL MULTI-TURN DIALOG

RELATED APPLICATIONS

[0001] This international application claims the benefit of priority to U.S. Patent Application Serial No. 15/294,765 filed on October 16, 2016, entitled “Intelligent Online Personal Assistant With Multi-Turn Dialog Based On Visual Search,” which is incorporated herein by reference in its entirety.

BACKGROUND

[0002] Traditional searching is text-based rather than image-based or voice-based. Searching is overly time-consuming when too many irrelevant results must be presented, browsed, and rejected by a user. The technical limitations of conventional search tools make it difficult for a user to communicate search intent, for example by sharing photos of interesting products, to help start a search that may be refined by further user input, such as in a multi-turn dialog. As online searches balloon to billions of possible selectable products, comparison searching has become more important than ever, but current text-based solutions were not designed for this scale. Irrelevant results are often shown and do not bring out the best results. Traditional forms of comparison searching (search+refinement+browse) are no longer useful.

BRIEF SUMMARY

[0003] In one example, an intelligent personal assistant system includes scalable artificial intelligence (AI) that permeates the fabric of existing messaging platforms to provide an intelligent online personal assistant (or “bot”). The system may leverage existing inventories and curated databases to provide intelligent, personalized answers in predictive turns of communication between a human user and an intelligent online personal assistant. One example of an intelligent personal assistant system includes a knowledge graph. Machine learning components may continuously identify

and learn from user intents so that user identity and understanding is enhanced over time. The user experience thus provided is inspiring, intuitive, unique, and may be focused on the usage and behavioral patterns of certain age groups, such as millennials, for example.

BRIEF DESCRIPTION OF THE SEVERAL VIEWS OF THE DRAWINGS

[0004] The drawings illustrate generally, by way of example, but not by way of limitation, various embodiments discussed in the present document. In order more easily to identify the discussion of any particular element or act, the most significant digit or digits in a reference number refer to the figure number in which that element is first introduced.

[0005] FIG. 1 shows a networked system, according to some example embodiments.

[0006] FIG. 2 shows a general architecture of an intelligent personal assistant system, according to some example embodiments.

[0007] FIGS. 3A and 3B show components of a speech recognition component, according to some example embodiments.

[0008] FIG. 4 shows a representative software architecture software architecture, which may be used in conjunction with various hardware architectures described herein.

[0009] FIG. 5 shows components of a machine, according to some example embodiments, able to read instructions from a machine-readable medium (e.g., a computer-readable storage medium) and perform any one or more of the methodologies discussed herein.

[0010] FIG. 6 shows an example environment into which an intelligent online personal assistant can be deployed, according to some example embodiments.

[0011] FIG. 7 shows an overview of the intelligent personal assistant system processing natural language user inputs to generate an item recommendation in an electronic marketplace, according to some example embodiments.

[0012] FIG. 8 shows a visual search service that interacts with a knowledge graph service, according to some example embodiments.

[0013] FIG. 9 shows a visual product search that generates product suggestions, according to some example embodiments.

[0014] FIG. 10 shows a visual product search that generates product suggestions in response to user non-image input, according to some example embodiments.

[0015] FIG. 11 shows a visual product search that generates product suggestions in response to user category input, according to some example embodiments.

[0016] FIG. 12 shows a flowchart of a methodology for visual search service and knowledge graph service interaction, according to some example embodiments.

[0017] FIG. 13 shows an offline index generation service that interacts with an on-demand visual search service, according to some example embodiments.

[0018] FIG. 14 shows a flowchart of a methodology for offline index generation for visual search, according to some example embodiments.

DETAILED DESCRIPTION

[0019] “CARRIER SIGNAL” in this context refers to any intangible medium that is capable of storing, encoding, or carrying instructions for execution by the machine, and includes digital or analog communications signals or other intangible medium to facilitate communication of such instructions. Instructions may be transmitted or received over the network using a transmission medium via a network interface device and using any one of a number of well-known transfer protocols.

[0020] “CLIENT DEVICE” in this context refers to any machine that interfaces to a communications network to obtain resources from one or more server systems or other client devices. A client device may be, but is not limited to, a mobile phone, desktop computer, laptop, portable digital assistants (PDAs), smart phones, tablets, ultra books, netbooks, laptops, multi-processor systems, microprocessor-based or programmable consumer electronics, game consoles, set-top boxes, or any other communication device that a user may use to access a network.

[0021] “COMMUNICATIONS NETWORK” in this context refers to one or more portions of a network that may be an ad hoc network, an intranet, an extranet, a virtual private network (VPN), a local area network (LAN), a wireless LAN (WLAN), a wide area network (WAN), a wireless WAN (WWAN), a metropolitan area network (MAN), the Internet, a portion of the Internet, a portion of the Public Switched Telephone Network (PSTN), a plain old telephone service (POTS) network, a cellular telephone network, a wireless network, a Wi-Fi® network, another type of network, or a combination of two or more such networks. For example, a network or a portion of a network may include a wireless or cellular network and the coupling may be a Code Division Multiple Access (CDMA) connection, a Global System for Mobile communications (GSM) connection, or other type of cellular or wireless coupling. In this example, the coupling may implement any of a variety of types of data transfer technology, such as Single Carrier Radio Transmission Technology (1xRTT), Evolution-Data Optimized (EVDO) technology, General Packet Radio Service (GPRS) technology, Enhanced Data rates for GSM Evolution (EDGE) technology,

third Generation Partnership Project (3GPP) including 3G, fourth generation wireless (4G) networks, Universal Mobile Telecommunications System (UMTS), High Speed Packet Access (HSPA), Worldwide Interoperability for Microwave Access (WiMAX), Long Term Evolution (LTE) standard, others defined by various standard setting organizations, other long range protocols, or other data transfer technology.

[0022] “COMPONENT” in this context refers to a device, physical entity or logic having boundaries defined by function or subroutine calls, branch points, application program interfaces (APIs), or other technologies that provide for the partitioning or modularization of particular processing or control functions. Components may be combined via their interfaces with other components to carry out a machine process. A component may be a packaged functional hardware unit designed for use with other components and a part of a program that usually performs a particular function of related functions. Components may constitute either software components (e.g., code embodied on a machine-readable medium) or hardware components. A “hardware component” is a tangible unit capable of performing certain operations and may be configured or arranged in a certain physical manner. In various example embodiments, one or more computer systems (e.g., a standalone computer system, a client computer system, or a server computer system) or one or more hardware components of a computer system (e.g., a processor or a group of processors) may be configured by software (e.g., an application or application portion) as a hardware component that operates to perform certain operations as described herein. A hardware component may also be implemented mechanically, electronically, or any suitable combination thereof. For example, a hardware component may include dedicated circuitry or logic that is permanently configured to perform certain operations. A hardware component may be a special-purpose processor, such as a Field-Programmable Gate Array (FPGA) or an Application Specific Integrated Circuit (ASIC). A hardware component may also include programmable logic or circuitry that is temporarily configured by software to perform certain operations. For example, a hardware component may include software executed by a general-purpose processor or other

programmable processor. Once configured by such software, hardware components become specific machines (or specific components of a machine) uniquely tailored to perform the configured functions and are no longer general-purpose processors. It will be appreciated that the decision to implement a hardware component mechanically, in dedicated and permanently configured circuitry, or in temporarily configured circuitry (e.g., configured by software) may be driven by cost and time considerations. Accordingly, the phrase “hardware component” (or “hardware-implemented component”) should be understood to encompass a tangible entity, be that an entity that is physically constructed, permanently configured (e.g., hardwired), or temporarily configured (e.g., programmed) to operate in a certain manner or to perform certain operations described herein. Considering embodiments in which hardware components are temporarily configured (e.g., programmed), each of the hardware components need not be configured or instantiated at any one instance in time. For example, where a hardware component comprises a general-purpose processor configured by software to become a special-purpose processor, the general-purpose processor may be configured as respectively different special-purpose processors (e.g., comprising different hardware components) at different times. Software accordingly configures a particular processor or processors, for example, to constitute a particular hardware component at one instance of time and to constitute a different hardware component at a different instance of time. Hardware components can provide information to, and receive information from, other hardware components. Accordingly, the described hardware components may be regarded as being communicatively coupled. Where multiple hardware components exist contemporaneously, communications may be achieved through signal transmission (e.g., over appropriate circuits and buses) between or among two or more of the hardware components. In embodiments in which multiple hardware components are configured or instantiated at different times, communications between such hardware components may be achieved, for example, through the storage and retrieval of information in memory structures to which the multiple hardware

components have access. For example, one hardware component may perform an operation and store the output of that operation in a memory device to which it is communicatively coupled. A further hardware component may then, at a later time, access the memory device to retrieve and process the stored output. Hardware components may also initiate communications with input or output devices, and can operate on a resource (e.g., a collection of information). The various operations of example methods described herein may be performed, at least partially, by one or more processors that are temporarily configured (e.g., by software) or permanently configured to perform the relevant operations. Whether temporarily or permanently configured, such processors may constitute processor-implemented components that operate to perform one or more operations or functions described herein. As used herein, “processor-implemented component” refers to a hardware component implemented using one or more processors. Similarly, the methods described herein may be at least partially processor-implemented, with a particular processor or processors being an example of hardware. For example, at least some of the operations of a method may be performed by one or more processors or processor-implemented components. Moreover, the one or more processors may also operate to support performance of the relevant operations in a “cloud computing” environment or as a “software as a service” (SaaS). For example, at least some of the operations may be performed by a group of computers (as examples of machines including processors), with these operations being accessible via a network (e.g., the Internet) and via one or more appropriate interfaces (e.g., an Application Program Interface (API)). The performance of certain of the operations may be distributed among the processors, not only residing within a single machine, but deployed across a number of machines. In some example embodiments, the processors or processor-implemented components may be located in a single geographic location (e.g., within a home environment, an office environment, or a server farm). In other example embodiments, the processors or processor-implemented components may be distributed across a number of geographic locations.

[0023] “MACHINE-READABLE MEDIUM” in this context refers to a component, device or other tangible media able to store instructions and data temporarily or permanently and may include, but is not be limited to, random-access memory (RAM), read-only memory (ROM), buffer memory, flash memory, optical media, magnetic media, cache memory, other types of storage (e.g., Erasable Programmable Read-Only Memory (EEPROM)) and/or any suitable combination thereof. The term “machine-readable medium” should be taken to include a single medium or multiple media (e.g., a centralized or distributed database, or associated caches and servers) able to store instructions. The term “machine-readable medium” will also be taken to include any medium, or combination of multiple media, that is capable of storing instructions (e.g., code) for execution by a machine, such that the instructions, when executed by one or more processors of the machine, cause the machine to perform any one or more of the methodologies described herein. Accordingly, a “machine-readable medium” refers to a single storage apparatus or device, as well as “cloud-based” storage systems or storage networks that include multiple storage apparatus or devices. The term “machine-readable medium” excludes signals per se.

[0024] “PROCESSOR” in this context refers to any circuit or virtual circuit (a physical circuit emulated by logic executing on an actual processor) that manipulates data values according to control signals (e.g., “commands”, “op codes”, “machine code”, etc.) and which produces corresponding output signals that are applied to operate a machine. A processor may, for example, be a Central Processing Unit (CPU), a Reduced Instruction Set Computing (RISC) processor, a Complex Instruction Set Computing (CISC) processor, a Graphics Processing Unit (GPU), a Digital Signal Processor (DSP), an Application Specific Integrated Circuit (ASIC), a Radio-Frequency Integrated Circuit (RFIC) or any combination thereof. A processor may further be a multi-core processor having two or more independent processors (sometimes referred to as “cores”) that may execute instructions contemporaneously.

[0025] A portion of the disclosure of this patent document contains material that is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction by anyone of the patent document or the patent disclosure, as it appears in the Patent and Trademark Office patent files or records, but otherwise reserves all copyright rights whatsoever. The following notice applies to the software and data as described below and in the drawings that form a part of this document: Copyright 2016, eBay Inc, All Rights Reserved.

[0026] The description that follows includes systems, methods, techniques, instruction sequences, and computing machine program products that embody illustrative embodiments of the disclosure. In the following description, for the purposes of explanation, numerous specific details are set forth in order to provide an understanding of various embodiments of the inventive subject matter. It will be evident, however, to those skilled in the art, that embodiments of the inventive subject matter may be practiced without these specific details. In general, well-known instruction instances, protocols, structures, and techniques are not necessarily shown in detail.

[0027] With reference to Figure 1, an example embodiment of a high-level SaaS network architecture 100 is shown. A networked system 116 provides server-side functionality via a network 110 (e.g., the Internet or wide area network (WAN)) to a client device 108. A web client 102 and a programmatic client, in the example form of an application 104 are hosted and execute on the client device 108. The networked system 116 includes an application server 122, which in turn hosts an intelligent personal assistant system 106 that provides a number of functions and services to the application 104 that accesses the networked system 116. The application 104 also provides a number of interfaces described herein, which present output of the tracking and analysis operations to a user of the client device 108.

[0028] The client device 108 enables a user to access and interact with the networked system 116. For instance, the user provides input (e.g., touch screen input or alphanumeric input) to the client device 108, and the input is communicated to the networked system 116 via the network 110. In this

instance, the networked system 116, in response to receiving the input from the user, communicates information back to the client device 108 via the network 110 to be presented to the user.

[0029] An Application Program Interface (API) server 118 and a web server 120 are coupled to, and provide programmatic and web interfaces respectively, to the application server 122. The application server 122 hosts an intelligent personal assistant system 106, which includes components or applications. The application server 122 is, in turn, shown to be coupled to a database server 124 that facilitates access to information storage repositories (e.g., a database/cloud 126). In an example embodiment, the database/cloud 126 includes storage devices that store information accessed and generated by the intelligent personal assistant system 106.

[0030] Additionally, a third party application 114, executing on a third party server 112, is shown as having programmatic access to the networked system 116 via the programmatic interface provided by the Application Program Interface (API) server 118. For example, the third party application 114, using information retrieved from the networked system 116, may support one or more features or functions on a website hosted by the third party.

[0031] Turning now specifically to the applications hosted by the client device 108, the web client 102 may access the various systems (e.g., intelligent personal assistant system 106) via the web interface supported by the web server 120. Similarly, the application 104 (e.g., an “app”) accesses the various services and functions provided by the intelligent personal assistant system 106 via the programmatic interface provided by the Application Program Interface (API) server 118. The application 104 may, for example, an “app” executing on a client device 108, such as an iOS or Android OS application to enable user to access and input data on the networked system 116 in an off-line manner, and to perform batch-mode communications between the programmatic client application 104 and the networked system 116.

[0032] Further, while the SaaS network architecture 100 shown in Figure 1 employs a client-server architecture, the present inventive subject matter is of course not limited to such an architecture, and could equally well find application in a distributed, or peer-to-peer, architecture system, for example. The intelligent personal assistant system 106 could also be implemented as a standalone software program, which does not necessarily have networking capabilities.

[0033] Figure 2 is a block diagram showing the general architecture of an intelligent personal assistant system 106, according to some example embodiments. Specifically, the intelligent personal assistant system 106 is shown to include a front end component 202 (FE) by which the intelligent personal assistant system 106 communicates (e.g., over the network 110) with other systems within the SaaS network architecture 100. The front end component 202 can communicate with the fabric of existing messaging systems. As used herein, the term messaging fabric refers to a collection of APIs and services that can power third party platforms such as Facebook messenger, Microsoft Cortana and other “bots”. In one example, a messaging fabric can support an online commerce ecosystem that allows users to interact with commercial intent. Output of the front end component 202 can be rendered in a display of a client device, such as the client device 108 in Figure 1 as part of an interface with an intelligent personal assistant, or “bot”.

[0034] The front end component 202 of the intelligent personal assistant system 106 is coupled to a back end component 204 for the front end (BFF) that operates to link the front end component 202 with an artificial intelligence framework 128. The artificial intelligence framework 128 may include several components as discussed below. The data exchanged between various components and the function of each component may vary to some extent, depending on the particular implementation.

[0035] In one example of an intelligent personal assistant system 106, an AI orchestrator 206 orchestrates communication between components inside and outside the artificial intelligence framework 128. Input modalities for the AI orchestrator 206 may be derived from a computer vision component

208, a speech recognition component 210, and a text normalization component which may form part of the speech recognition component 210, for example. The computer vision component 208 may identify objects and attributes from visual input (e.g., a photo). The speech recognition component 210 may convert audio signals (e.g., spoken utterances) into text. A text normalization component may operate to make input normalization, such as language normalization by rendering emoticons into text, for example. Other normalization is possible such as orthographic normalization, foreign language normalization, conversational text normalization, and so forth.

[0036] The artificial intelligence framework 128 further includes a natural language understanding or NLU component 214 that operates to extract user intent and various intent parameters. The NLU component 214 is described in further detail beginning with Figure 8.

[0037] The artificial intelligence framework 128 further includes a dialog manager 216 that operates to understand a “completeness of specificity” (for example of an input, such as a search query or utterance) and decide on a next action type and a related parameter (e.g., “search” or “request further information from user”). For convenience, all user inputs in this description may be referred to as “utterances”, whether in text, voice, or image-related formats.

[0038] In one example, the dialog manager 216 operates in association with a context manager 218 and a Natural Language Generation (NLG) component 212. The context manager 218 manages the context and communication of a user with respect to the intelligent online personal assistant (or “bot”) and the assistant's associated artificial intelligence. The context manager 218 retains a short term history of user interactions. A longer term history of user preferences may be retained in an identity service 222, described below. Data entries in one or both of these histories may include the relevant intent and all parameters and all related results of a given input, bot interaction, or turn of communication, for example. The NLG component 212 operates to compose a natural language utterance out of an AI message to present to a user interacting with the intelligent bot.

[0039] A search component 220 is also included within the artificial intelligence framework 128. The search component 220 may have front and back end units. The back end unit may operate to manage item or product inventory and provide functions of searching against the inventory, optimizing towards a specific tuple of user intent and intent parameters. The search component 220 is designed to serve several billion queries per day globally against very large high quality inventories. The search component 220 can accommodate text, or Artificial Intelligence (AI) encoded voice and image inputs, and identify relevant inventory items to users based on explicit and derived query intents.

[0040] An identity service 222 component operates to manage user profiles, for example explicit information in the form of user attributes, e.g., “name”, “age”, “gender”, “geolocation”, but also implicit information in forms such as “information distillates” such as “user interest”, or “similar persona”, and so forth. The artificial intelligence framework 128 may comprise part of or operate in association with, the identity service 222. The identity service 222 includes a set of policies, APIs, and services that elegantly centralizes all user information, helping the artificial intelligence framework 128 to have “intelligent” insights into user intent. The identity service 222 can protect online retailers and users from fraud or malicious use of private information.

[0041] The identity service 222 of the present disclosure provides many advantages. The identity service 222 is a single central repository containing user identity and profile data. It may continuously enrich the user profile with new insights and updates. It uses account linking and identity federation to map relationships of a user with a company, household, other accounts (e.g., core account), as well as a user’s social graph of people and relationships. The identity service 222 evolves a rich notification system that communicates all and only the information the user wants at the times and media they choose.

[0042] In one example, the identity service 222 concentrates on unifying as much user information as possible in a central clearinghouse for search, AI, merchandising, and machine learning models to maximize each component’s

capability to deliver insights to each user. A single central repository contains user identity and profile data in a meticulously detailed schema. In an onboarding phase, the identity service 222 primes a user profile and understanding by mandatory authentication in a bot application. Any public information available from the source of authentication (e.g., social media) may be loaded. In sideboarding phases, the identity service 222 may augment the profile with information about the user that is gathered from public sources, user behaviors, interactions, and the explicit set of purposes the user tells the AI (e.g., shopping missions, inspirations, preferences). As the user interacts with the artificial intelligence framework 128, the identity service 222 gathers and infers more about the user and stores the explicit data, derived information, and updates probabilities and estimations of other statistical inferences. Over time, in profile enrichment phases, the identity service 222 also mines behavioral data such as clicks, impressions, and browse activities for derived information such as tastes, preferences, and shopping verticals. In identity federation and account linking phases, when communicated or inferred, the identity service 222 updates the user's household, employer, groups, affiliations, social graph, and other accounts, including shared accounts.

[0043] The functionalities of the artificial intelligence framework 128 can be grouped into multiple parts, for example decisioning and context parts. In one example, the decisioning part includes operations by the AI orchestrator 206, the NLU component 214, the dialog manager 216, the NLG component 212, the computer vision component 208 and speech recognition component 210. The context part of the AI functionality relates to the parameters (implicit and explicit) around a user and the communicated intent (for example, towards a given inventory, or otherwise). In order to measure and improve AI quality over time, the artificial intelligence framework 128 may be trained using sample queries (e.g., a dev set) and tested on a different set of queries (e.g., an eval set), where both sets may be developed by human curation. Also, the artificial intelligence framework 128 may be trained on transaction and interaction flows defined by experienced curation specialists, or human tastemaker override rules 224. The flows and the logic encoded

within the various components of the artificial intelligence framework 128 define what follow-up utterance or presentation (e.g., question, result set) is made by the intelligent assistant based on an identified user intent.

[0044] Reference is made further above to example input modalities of the intelligent online personal assistant or bot in an intelligent personal assistant system 106. The intelligent personal assistant system 106 seeks to understand a user's intent (e.g., targeted search, compare, shop/browse, and so forth) and any mandatory parameters (e.g., product, product category, item, and so forth) and/or optional parameters (e.g., explicit information such as attributes of item/product, occasion, and so forth) as well as implicit information (e.g., geolocation, personal preferences, age, and gender, and so forth) and respond to the user with a well thought out or "intelligent" response. Explicit input modalities may include text, speech, and visual input and can be enriched with implicit knowledge of user (e.g., geolocation, previous browse history, and so forth). Output modalities can include text (such as speech, or natural language sentences, or product-relevant information, and images on the screen of a smart device, e.g., client device 108. Input modalities thus refer to the different ways users can communicate with the bot. Input modalities can also include keyboard or mouse navigation, touch-sensitive gestures, and so forth.

[0045] In relation to a modality for the computer vision component 208, a photograph can often represent what a user is looking for better than text. The user may not know what an item is called, or it may be hard or even impossible to use text for fine detailed information that only an expert may know, for example a complicated pattern in apparel or a certain style in furniture. Moreover, it is inconvenient to type complex text queries on mobile phones, and long text queries typically have poor recall. Thus, key functionalities of the computer vision component 208 may include object localization, object recognition, optical character recognition (OCR) and matching against inventory based on visual cues from an image or video. A bot enabled with computer vision is advantageous when running on a mobile device which has a built-in camera. Powerful deep neural networks can be used to enable computer vision applications.

[0046] In one example, the dialog manager 216 has as sub-components the context manager 218 and the NLG component 212. As mentioned above, the dialog manager 216 operates to understand the “completeness of specificity” and deciding on a next action type and parameter (e.g., “search” or “request further information from user”). The context manager 218 operates to manage the context and communication of a given user towards the bot and its AI. The context manager 218 comprises two parts: a long term history and a short term memory. Each context manager entry may describe the relevant intent and all parameters and all related results. The context is towards the inventory, as well as towards other, future sources of knowledge. The NLG component 212 operates to compose a natural language utterance out of an AI message to present to a user interacting with the intelligent bot.

[0047] Fluent, natural, informative, and even entertaining dialog between man and machine is a difficult technical problem that has been studied for much of the past century, yet is still considered unsolved. However, recent developments in AI have produced useful dialog systems such as Siri™ and Alexa™.

[0048] In an ecommerce example of an intelligent bot, an initial very helpful element in seeking to solve this problem is to leverage enormous sets of e-commerce data. Some of this data may be retained in proprietary databases or in the cloud e.g., database/cloud 126. Statistics about this data may be communicated to dialog manager 216 from the search component 220 as context. The artificial intelligence framework 128 may act directly upon utterances from the user, which may be run through speech recognition component 210, then the NLU component 214, and then passed to context manager 218 as semi-parsed data. The NLG component 212 may thus help the dialog manager 216 generate human-like questions and responses in text or speech to the user. The context manager 218 maintains the coherency of multi-turn and long term discourse between the user and the artificial intelligence framework 128.

[0049] Discrimination may be recommended to poll a vast e-commerce dataset for only relevant, useful information. In one example, the artificial

intelligence framework 128 uses results from the search component 220 and intelligence within the search component 220 to provide this information. This information may be combined with the history of interaction from the context manager 218. The artificial intelligence framework 128 then may decide on the next turn of dialog, e.g., whether it should be a question, or a “grounding statement” to validate, for example, an existing understanding or user intent, or an item recommendation (or, for example, any combination of all three). These decisions may be made by a combination of the dataset, the chat history of the user, and a model of the user’s understanding. The NLG component 212 may generate language for a textual or spoken reply to the user based on these decisions.

[0050] Technical solutions provided by the present inventive subject matter allow users to communicate with an intelligent online personal assistant in a natural conversation. The assistant is efficient as over time it increasingly understands specific user preferences and is knowledgeable about a wide range of products. Though a variety of convenient input modalities, a user can share photos, or use voice or text, and the assisted user experience may be akin to talking to a trusted, knowledgeable human shopping assistant in a high-end store, for example.

[0051] Conventionally, the approach and data used by online shopping systems aim at a faceless demographic group of buyers with blunt, simplified assumptions to maximize short-term revenue. Conventional sites and apps do not understand how, why, and when users want to be notified. Notifications may be annoying, inappropriate, and impersonal, oblivious to each user’s preferences. One person is not the same as a single account. People share accounts and devices. Passwords make platforms neither safe nor easy to use. Problems of weak online identity and the ignoring of environmental signals (such as device, location, notification after anomalous behavior) make it easy to conduct fraud in the marketplace.

[0052] With reference to Figure 3A, the illustrated components of the speech recognition component 210 are now described. A feature extraction component operates to convert raw audio waveform to some-dimensional vector of numbers that represents the sound. This component uses deep

learning to project the raw signal into a high-dimensional semantic space. An acoustic model component operates to host a statistical model of speech units, such as phonemes and allophones. These can include Gaussian Mixture Models (GMM) although the use of Deep Neural Networks is possible. A language model component uses statistical models of grammar to define how words are put together in a sentence. Such models can include n-gram-based models or Deep Neural Networks built on top of word embeddings. A speech-to-text (STT) decoder component may convert a speech utterance into a sequence of words typically leveraging features derived from a raw signal using the feature extraction component, the acoustic model component, and the language model component in a Hidden Markov Model (HMM) framework to derive word sequences from feature sequences. In one example, a speech-to-text service in the cloud (e.g., database/cloud 126) has these components deployed in a cloud framework with an API that allows audio samples to be posted for speech utterances and to retrieve the corresponding word sequence. Control parameters are available to customize or influence the speech-to-text process.

[0053] In one example of an artificial intelligence framework 128, two additional parts for the speech recognition component 210 are provided, a speaker adaptation component and a Language Model (LM) adaptation component. The speaker adaptation component allows clients of an STT system (e.g., speech recognition component 210) to customize the feature extraction component and/or the acoustic model component for each speaker/user. This can be important because most speech-to-text systems are trained on data from a representative set of speakers from a target region and typically the accuracy of the system depends heavily on how well the target speaker matches the speakers in the training pool. The speaker adaptation component allows the speech recognition component 210 (and consequently the artificial intelligence framework 128) to be robust to speaker variations by continuously learning the idiosyncrasies of a user's intonation, pronunciation, accent, and other speech factors, and apply these to the speech-dependent components, e.g., the feature extraction component, and the acoustic model component. While this approach may require a small

voice profile to be created and persisted for each speaker, the potential benefits of accuracy generally far outweigh the storage drawbacks.

[0054] The LM adaptation component operates to customize the language model component and the speech-to-text vocabulary with new words and representative sentences from a target domain, for example, inventory categories or user personas. This capability allows the artificial intelligence framework 128 to be scalable as new categories and personas are supported.

[0055] Figure 3B shows a flow sequence 302 for text normalization in an artificial intelligence framework 128. A text normalization component performing the flow sequence 302 is included in the speech recognition component 210 in one example. Key functionalities in the flow sequence 302 include orthographic normalization (to handle punctuation, numbers, case, and so forth), conversational text normalization (to handle informal chat-type text with acronyms, abbreviations, incomplete fragments, slang, and so forth), and machine translation (to convert a normalized sequence of foreign-language words into a sequence of words in an operating language, including but not limited to English for example).

[0056] The artificial intelligence framework 128 facilitates modern communications. Millennials for example often want to communicate via photos, voice, and text. The technical ability of the artificial intelligence framework 128 to use multiple modalities allows the communication of intent instead of just text. The artificial intelligence framework 128 provides technical solutions and is efficient. It is faster to interact with a smart personal assistant using voice commands or photos than text in many instances.

[0057] Figure 4 is a block diagram illustrating an example software architecture 406, which may be used in conjunction with various hardware architectures described herein. Figure 4 is a non-limiting example of a software architecture and it will be appreciated that many other architectures may be implemented to facilitate the functionality described herein. The software architecture 406 may execute on hardware such as machine 500 of Figure 5 that includes, among other things, processors 504, memory 514, and

input/output (I/O) components 518. A representative hardware layer 452 is illustrated and can represent, for example, the machine 500 of Figure 5. The representative hardware layer 452 includes a processing unit 454 having associated executable instructions 404. Executable instructions 404 represent the executable instructions of the software architecture 406, including implementation of the methods, components and so forth described herein. The hardware layer 452 also includes memory and/or storage modules memory/storage 456, which also have executable instructions 404. The hardware layer 452 may also comprise other hardware 458.

[0058] In the example architecture of Figure 4, the software architecture 406 may be conceptualized as a stack of layers where each layer provides particular functionality. For example, the software architecture 406 may include layers such as an operating system 402, libraries 420, applications 416 and a presentation layer 414. Operationally, the applications 416 and/or other components within the layers may invoke application programming interface (API) calls 408 through the software stack and receive a response as in response to the API calls 408. The layers illustrated are representative in nature and not all software architectures have all layers. For example, some mobile or special purpose operating systems may not provide a frameworks/middleware 418, while others may provide such a layer. Other software architectures may include additional or different layers.

[0059] The operating system 402 may manage hardware resources and provide common services. The operating system 402 may include, for example, a kernel 422, services 424 and drivers 426. The kernel 422 may act as an abstraction layer between the hardware and the other software layers. For example, the kernel 422 may be responsible for memory management, processor management (e.g., scheduling), component management, networking, security settings, and so on. The services 424 may provide other common services for the other software layers. The drivers 426 are responsible for controlling or interfacing with the underlying hardware. For instance, the drivers 426 may include display drivers, camera drivers, Bluetooth® drivers, flash memory drivers, serial communication drivers (e.g., Universal Serial Bus (USB) drivers), Wi-Fi® drivers, audio

drivers, power management drivers, and so forth depending on the hardware configuration.

[0060] The libraries 420 provide a common infrastructure that is used by the applications 416 and/or other components and/or layers. The libraries 420 may provide functionality that allows other software components to perform tasks in an easier fashion than to interface directly with the underlying operating system 402 functionality (e.g., kernel 422, services 424, and/or drivers 426). The libraries 420 may include system libraries 444 (e.g., C standard library) that may provide functions such as memory allocation functions, string manipulation functions, mathematical functions, and the like. In addition, the libraries 420 may include API libraries 446 such as media libraries (e.g., libraries to support presentation and manipulation of various known media formats such as MPREG4, H.264, MP3, AAC, AMR, JPG, and PNG), graphics libraries (e.g., an OpenGL framework that may be used to render 2D and 3D graphic content on a display), database libraries (e.g., SQLite that may provide various relational database functions), web libraries (e.g., WebKit that may provide web browsing functionality), and the like. The libraries 420 may also include a wide variety of other libraries 448 to provide many other APIs to the applications 416 and other software components/modules.

[0061] The frameworks frameworks/middleware 418 (also sometimes referred to as middleware) may provide a higher-level common infrastructure that may be used by the applications 416 and/or other software components/modules. For example, the frameworks/middleware 418 may provide various graphic user interface (GUI) functions, high-level resource management, high-level location services, and so forth. The frameworks/middleware 418 may provide a broad spectrum of other APIs that may be utilized by the applications 416 and/or other software components/modules, some of which may be specific to a particular operating system or platform.

[0062] The applications 416 include built-in applications 438 and/or third-party applications 440. Examples of representative built-in applications 438 may include, but are not limited to, a contacts application, a browser

application, a book reader application, a location application, a media application, a messaging application, and/or a game application. Third-party applications 440 may include any an application developed using the ANDROID™ or IOS™ software development kit (SDK) by an entity other than the vendor of the particular platform, and may be mobile software running on a mobile operating system such as IOS™, ANDROID™, WINDOWS® Phone, or other mobile operating systems. The third-party applications 440 may invoke the API calls 408 provided by the mobile operating system (such as operating system 402) to facilitate functionality described herein.

[0063] The applications 416 may use built in operating system functions (e.g., kernel 422, services 424 and/or drivers 426), libraries 420, and frameworks/middleware 418 to create user interfaces to interact with users of the system. Alternatively, or additionally, in some systems interactions with a user may occur through a presentation layer, such as presentation layer 414. In these systems, the application/component “logic” can be separated from the aspects of the application/component that interact with a user.

[0064] Some software architectures use virtual machines. In the example of Figure 4, this is illustrated by a virtual machine 410. The virtual machine 410 creates a software environment where applications/components can execute as if they were executing on a hardware machine (such as the machine 500 of Figure 5, for example). The virtual machine 410 is hosted by a host operating system (operating system (OS) 436 in Figure 4) and typically, although not always, has a virtual machine monitor 460, which manages the operation of the virtual machine as well as the interface with the host operating system (e.g., operating system 402). A software architecture executes within the virtual machine 410 such as an operating system operating system (OS) 436, libraries 434, frameworks 432, applications 430 and/or presentation layer 428. These layers of software architecture executing within the virtual machine 410 can be the same as corresponding layers previously described or may be different.

[0065] Figure 5 is a block diagram illustrating components of a machine 500, according to some example embodiments, which is able to read instructions from a machine-readable medium (e.g., a machine-readable storage medium) and perform any one or more of the methodologies discussed herein. Specifically, Figure 5 shows a diagrammatic representation of the machine 500 in the example form of a computer system, within which instructions 510 (e.g., software, a program, an application, an applet, an app, or other executable code) for causing the machine 500 to perform any one or more of the methodologies discussed herein may be executed. As such, the instructions may be used to implement modules or components described herein. The instructions transform the general, non-programmed machine into a particular machine programmed to carry out the described and illustrated functions in the manner described. In alternative embodiments, the machine 500 operates as a standalone device or may be coupled (e.g., networked) to other machines. In a networked deployment, the machine 500 may operate in the capacity of a server machine or a client machine in a server-client network environment, or as a peer machine in a peer-to-peer (or distributed) network environment. The machine 500 may comprise, but is not limited to, a server computer, a client computer, a personal computer (PC), a tablet computer, a laptop computer, a netbook, a set-top box (STB), a personal digital assistant (PDA), an entertainment media system, a cellular telephone, a smart phone, a mobile device, a wearable device (e.g., a smart watch), a smart home device (e.g., a smart appliance), other smart devices, a web appliance, a network router, a network switch, a network bridge, or any machine capable of executing the instructions 510, sequentially or otherwise, that specify actions to be taken by machine 500. Further, while only a single machine 500 is illustrated, the term “machine” will also be taken to include a collection of machines that individually or jointly execute the instructions 510 to perform any one or more of the methodologies discussed herein.

[0066] The machine 500 may include processors 504, memory memory/storage 506, and I/O components 518, which may be configured to communicate with each other such as via a bus 502. The memory/storage

506 may include a memory 514, such as a main memory, or other memory storage, and a storage unit 516, both accessible to the processors 504 such as via the bus 502. The storage unit 516 and memory 514 store the instructions 510 embodying any one or more of the methodologies or functions described herein. The instructions 510 may also reside, completely or partially, within the memory 514, within the storage unit 516, within at least one of the processors 504 (e.g., within the processor's cache memory), or any suitable combination thereof, during execution thereof by the machine 500.

Accordingly, the memory 514, the storage unit 516, and the memory of processors 504 are examples of machine-readable media.

[0067] The I/O components 518 may include a wide variety of components to receive input, provide output, produce output, transmit information, exchange information, capture measurements, and so on. The specific I/O components 518 that are included in a particular machine will depend on the type of machine. For example, portable machines such as mobile phones will likely include a touch input device or other such input mechanisms, while a headless server machine will likely not include such a touch input device. It will be appreciated that the I/O components 518 may include many other components that are not shown in Figure 5. The I/O components 518 are grouped according to functionality merely for simplifying the following discussion and the grouping is in no way limiting. In various example embodiments, the I/O components 518 may include output components 526 and input components 528. The output components 526 may include visual components (e.g., a display such as a plasma display panel (PDP), a light emitting diode (LED) display, a liquid crystal display (LCD), a projector, or a cathode ray tube (CRT)), acoustic components (e.g., speakers), haptic components (e.g., a vibratory motor, resistance mechanisms), other signal generators, and so forth. The input components 528 may include alphanumeric input components (e.g., a keyboard, a touch screen configured to receive alphanumeric input, a photo-optical keyboard, or other alphanumeric input components), point based input components (e.g., a mouse, a touchpad, a trackball, a joystick, a motion sensor, or other pointing instrument), tactile input components (e.g., a physical button, a

touch screen that provides location and/or force of touches or touch gestures, or other tactile input components), audio input components (e.g., a microphone), and the like.

[0068] In further example embodiments, the I/O components 518 may include biometric components 530, motion components 534, environmental environment components 536, or position components 538 among a wide array of other components. For example, the biometric components 530 may include components to detect expressions (e.g., hand expressions, facial expressions, vocal expressions, body gestures, or eye tracking), measure biosignals (e.g., blood pressure, heart rate, body temperature, perspiration, or brain waves), identify a person (e.g., voice identification, retinal identification, facial identification, fingerprint identification, or electroencephalogram based identification), and the like. The motion components 534 may include acceleration sensor components (e.g., accelerometer), gravitation sensor components, rotation sensor components (e.g., gyroscope), and so forth. The environment components 536 may include, for example, illumination sensor components (e.g., photometer), temperature sensor components (e.g., one or more thermometer that detect ambient temperature), humidity sensor components, pressure sensor components (e.g., barometer), acoustic sensor components (e.g., one or more microphones that detect background noise), proximity sensor components (e.g., infrared sensors that detect nearby objects), gas sensors (e.g., gas detection sensors to detection concentrations of hazardous gases for safety or to measure pollutants in the atmosphere), or other components that may provide indications, measurements, or signals corresponding to a surrounding physical environment. The position components 538 may include location sensor components (e.g., a Global Position System (GPS) receiver component), altitude sensor components (e.g., altimeters or barometers that detect air pressure from which altitude may be derived), orientation sensor components (e.g., magnetometers), and the like.

[0069] Communication may be implemented using a wide variety of technologies. The I/O components 518 may include communication components 540 operable to couple the machine 500 to a network 532 or

devices 520 via coupling 522 and coupling 524 respectively. For example, the communication components 540 may include a network interface component or other suitable device to interface with the network 532. In further examples, communication components 540 may include wired communication components, wireless communication components, cellular communication components, Near Field Communication (NFC) components, Bluetooth® components (e.g., Bluetooth® Low Energy), Wi-Fi® components, and other communication components to provide communication via other modalities. The devices 520 may be another machine or any of a wide variety of peripheral devices (e.g., a peripheral device coupled via a Universal Serial Bus (USB)).

[0070] Moreover, the communication components 540 may detect identifiers or include components operable to detect identifiers. For example, the communication components processors communication components 540 may include Radio Frequency Identification (RFID) tag reader components, NFC smart tag detection components, optical reader components (e.g., an optical sensor to detect one-dimensional bar codes such as Universal Product Code (UPC) bar code, multi-dimensional bar codes such as Quick Response (QR) code, Aztec code, Data Matrix, Dataglyph, MaxiCode, PDF417, Ultra Code, UCC RSS-2D bar code, and other optical codes), or acoustic detection components (e.g., microphones to identify tagged audio signals). In addition, a variety of information may be derived via the communication components 540, such as, location via Internet Protocol (IP) geo-location, location via Wi-Fi® signal triangulation, location via detecting a NFC beacon signal that may indicate a particular location, and so forth.

[0071] With reference now to Figure 6, an example environment 600 into which an intelligent online personal assistant provided by the intelligent personal assistant system 106 can be deployed is shown. At the center of the environment 600, the intelligent bot 602 with AI appears. The bot leverages the computer vision component 208, the speech recognition component 210, the NLU component 214, the dialog manager 216, the NLG component 212, the search component 220, and identity service 222 to engage users in

efficient, interesting, and effective dialog to decode their intent and deliver personalized results.

[0072] An associated application 604 can showcase the bot 602's full power and intelligence with compelling mobile design capabilities and elements. The fabric 606 integrates with Facebook Messenger™, Skype™, and Cortana™ (for example) to enable users to transact where they are already spending time. A smart notifications 610 platform delivers the right information at the right time via any number of channels (e.g., SMS, push notification, email, messaging) to users to encourage them to engage with the bot 602 and associated marketplaces. Communities 608 features enable users to connect, engage, and interact with their friends, tastemakers, and brands using the same messaging systems in which they already spend most of their time. Other features include group buying and gift buying. A rewards 612 platform incentivizes users to engage more deeply with the bot 602. Rewards can include deep discounts on products, access to unique inventory, and recognition in the app through scores, levels, etc. At marketing 614, a combination of traditional, social and other marketing is performed to win the attention of some populations (e.g., millennials) in more personal ways. Conventional techniques can include merchandising, email, search engine optimization (SEO), and search engine marketing (SEM) as well as experimental techniques such as social ads, viral coupons, and more to target new and existing users.

[0073] FIG. 7 shows an overview of the intelligent personal assistant system 106 processing natural language user inputs to generate an item recommendation in an electronic marketplace. Although the intelligent personal assistant system 106 is not limited to this use scenario, it may be of particular utility in this situation. As previously described, any combination of text, image, and voice data may be received by the artificial intelligence framework 128. Image data may be processed by the computer vision component 208 to provide image attribute data. Voice data may be processed by the speech recognition component 210 into text.

[0074] All of these inputs and others may be provided to the NLU component 214 for analysis. The NLU component 214 may operate to parse

user inputs and help determine the user intent and intent-related parameters. For example, the NLU component 214 may discern the dominant object of user interest, and a variety of attributes and attribute values related to that dominant object. The NLU component 214 may also determine other parameters such as the user input type (e.g., a question or a statement) and targeted item recipients. The NLU component 214 may provide extracted data to the dialog manager 216, as well as the AI orchestrator 206 previously shown.

[0075] The NLU component 214 may generally transform formal and informal natural language user inputs into a more formal, machine-readable, structured representation of a user's query. That formalized query may be enhanced further by the dialog manager 216. In one scenario, the NLU component 214 processes a sequence of user inputs including an original query and further data provided by a user in response to machine-generated prompts from the dialog manager 216 in a multi-turn interactive dialog. This user-machine interaction may improve the efficiency and accuracy of one or more automated searches for the most relevant items available for purchase in an electronic marketplace. The searches may be performed by the search component 220.

[0076] Extracting user intent is very helpful for the AI bot in determining what further action is needed. In one ecommerce-related example, at the very highest level, user intent could be shopping, chit-chat, jokes, weather, etc. If the user intent is shopping, it could relate to the pursuit of a specific shopping mission, gifting an item for a target recipient other than the user, or just to browse an inventory of items available for purchase. Once the high level intent is identified, the artificial intelligence framework 128 is tasked with determining what the user is looking for; that is, is the need broad (e.g., shoes, dresses) or more specific (e.g., two pairs of new black NikeTM size 10 sneakers) or somewhere in between (e.g., black sneakers)?

[0077] In a novel and distinct improvement over the prior art in this field, the artificial intelligence framework 128 may map the user request to certain primary dimensions, such as categories, attributes, and attribute values, that best characterize the available items desired. This gives the bot the ability to

engage with the user to further refine the search constraints if necessary. For example, if a user asks the bot for information relating to dresses, the top attributes that need specification might be color, material, and style.

Further, over time, machine learning may add deeper semantics and wider “world knowledge” to the system, to better understand the user intent. For example the input “I am looking for a dress for a wedding in June in Italy” means the dress should be appropriate for particular weather conditions at a given time and place, and should be appropriate for a formal occasion.

Another example might include a user asking the bot for “gifts for my nephew”. The artificial intelligence framework 128 when trained will understand that gifting is a special type of intent, that the target recipient is male based on the meaning of “nephew”, and that attributes such as age, occasion, and hobbies/likes of the target recipient should be clarified.

[0078] Figure 8 shows a visual search service 800 that interacts with a knowledge graph service 822, according to some example embodiments. The present inventors have recognized, among other things, that product searches in an electronic marketplace may be improved by combining a visual search with a knowledge graph based search. An electronic marketplace that processes millions of products per week may collect a vast amount of data about its inventory, which may be represented by knowledge graph entries.

[0079] However, existing text-based search tools may not always capitalize on that data as effectively as tools that perform visual searches or combine visual searches and knowledge graph based searches. This description therefore provides a method for processing natural language input so a user can easily take advantage of both knowledge graph information (via a multi-turn dialog for example), and visual query information so that a user can easily take advantage of product image information.

[0080] The visual search service 800 may, in one embodiment, comprise software instructions called for execution on at least one hardware-based processor by the AI orchestrator 206 which coordinates information within the artificial intelligence framework 128 as previously described. In one embodiment, the visual search service 800 may be part of the computer

vision component 208, but in general the visual search service 800 may coordinate information between any number of components.

[0081] The visual search service 800 may receive an image query 802 originating from a user. The image query 802 may comprise one or more images the user believes will be helpful in finding a particular product. For simplicity, but not by limitation, this description may refer to image query 802 as comprising a single input query image. The input query image may comprise a photograph, a video frame, a sketch, or a diagram, for example. The input query image is typically a digital image file such as may be produced by a portable camera or smartphone, or such as may be copied from a web site or an electronic message.

[0082] The visual search service 800 may comprise or interact with a number of functional component blocks, which may each comprise a hardware processor implemented software program for example. A neural network (not shown) in the visual search service 800 for example may process the input query image. The neural network may comprise a fully convolutional neural network (FCN) as described in the previously cited article by Long *et al.* In another embodiment, the neural network may comprise a hybrid neural network (termed a CRF-RNN) including a fully convolutional neural network and a recurrent network (RNN) that includes conditional random fields (CRF) as described in the previously cited article by Zheng *et al.*

[0083] Images processed by the neural network may comprise an input query image from the image query 802 as well as any number of images associated with any number of candidate products in an electronic marketplace, for example. The neural network in the visual search service 800 may produce an image signature that concisely describes image content. In general, an image signature may numerically describe a number of image features and their relative dominance of overall image content. Each image signature may comprise a vector of binary numbers for example, also referred to as a binary hash. Any form of image signature may be considered to be within the scope of this description.

[0084] The visual search service 800 may generate metadata and image signatures from input query images. The visual search service 800 may also receive metadata and image signatures from product images, shown in block 814. Metadata may comprise for example a product identification (ID) number and a universal resource locator (URL) for a product listing in the electronic marketplace.

[0085] The visual search service 800 may then calculate a visual similarity measure between images, such as between a particular candidate product image and the input query image. The visual similarity measure may be estimated by calculating a distance value between two image signatures. The distance may comprise a Hamming distance, by way of example but not limitation. A Hamming distance generally describes the number of bits that are different in two binary vectors. Similar images being compared may therefore have a smaller Hamming distance between them, and thus a higher visual similarity measure. The visual similarity measure is therefore useful as a search result score, e.g., for the candidate product at hand.

[0086] In one embodiment, the external data in block 814 is computed offline for some or all of the products or items available in an electronic marketplace, and is stored in a database. For a marketplace with a large number of products available, substantially real time computation of image signatures may not be computationally feasible. The visual search service 800 may not only be used to help a shopper find a relevant product in the electronic marketplace, but may be used for other external visual search tasks assigned by external tools.

[0087] In one embodiment, each product image previously provided by sellers in an electronic marketplace may be processed to generate an image signature that may be stored in the index 814. The processing may be performed offline to build a catalog of image signatures without interfering with ongoing “live” operations of the electronic marketplace.

[0088] Any approach for calculating the visual similarity measure may provide the search result score described. Visual search result scores for any number of candidate products may for example be generated via visual

comparisons with an input query image as described above. The visual search result scores may determine the order in which ranked candidate products may be presented to a user in response to the image query 802. The end result of the visual search methodology described may comprise an output item list 818 that may correspond to available products in the electronic marketplace, for example.

[0089] The results of a visual search may be factored into an overall composite search scheme in any number of different formulations. In one example, a weighting coefficient may weight the visual search result score by a user-adjustable weighting factor, and the remaining weight may be applied to scores from a leaf category prediction from the knowledge graph. Figure 11 describes this type of composite weighting in more detail.

[0090] The knowledge graph may have categories, aspects, and aspect values provided by sellers to help buyers find the product in the electronic marketplace inventory. Similarly, the knowledge graph may include popular categories, aspects, and aspect values that buyers have frequently used when searching for particular items. Categories may describe predetermined product groupings or sub-groupings provided by the electronic marketplace (e.g., “wine”, “shoes”, “paint”), or may be open-ended for seller definition. Categories may be branched, so that a particularly narrow sub-category may be regarded as a leaf category (e.g., “men’s athletic shoes”) that may best narrow a given search to a small set of items best meeting a specified set of search constraints.

[0091] Aspects (also called attributes) may comprise descriptive parameters that may be specified by particular values, to provide further precise search keys for finding a particular product. Exemplary aspects or attributes may include but are not limited to “brand”, “color”, “style”, “material”, “size”. Corresponding exemplary values may include “Nike”, “red”, “running”, “canvas”, “ten”, for example. Knowledge graph construction and use is described further in the related applications previously incorporated by reference.

[0092] Some aspect values may be readily discernible via a visual search, while others may not. The combination of a visual search and a knowledge graph search may therefore enhance product searching significantly.

Broadly speaking, each type of search will find some results, but the combination finds the intersection of the results. The final output is often astonishingly accurate and unexpectedly fast.

[0093] Figure 9 shows a visual product search that generates product suggestions, according to some example embodiments. The user provides input query image 902 of a shoe. The image signature 814 generated from the input query image 902 may in itself denote that a particular category (“women’s heels”) and aspects and aspect value tuples (“color:red”, and “style:open toe”) are strongly indicated. The visual search may therefore be narrowed to those candidate products that share those category, aspect, and aspect values, rather than including all available candidate products. That is, a visual search alone may resolve categories, aspects, and aspect values just as a non-visual search may, depending on the inputs provided.

[0094] If the item list 818 includes several strong “hits”, based on the visual similarity measures obtained, the knowledge graph service 822 may merely confirm that the category, aspects, and aspect values are also strongly linked to the candidate products. The “bot” may thus elect to perform little or no further filtering of the candidate products in the item list. The “bot” may also elect to output several of the most strongly matching candidate products rather than generating a statement type or question type user prompt.

[0095] The number of candidate products displayed (three in this case, 908, 910, and 912) may thus generally be based on the amount of result filtering performed at block 824. That is, if the knowledge graph service 822 finds a strong indication that little re-ranking or filtering is required to reconcile the two types of searches, it may elect to output only a small number of the best matching candidate products. If the indication is sufficiently strong, as will be shown in Figure 11, only a single matching candidate product may need to be displayed.

[0096] In general, searching and prompting may iterate until a stopping condition occurs. A user may provide input that clearly indicates an intent change for example. A user may also submit an entirely new query to try a different approach to achieving the same or similar search intent; this is often the result of user frustration. And, of course, a user may be sufficiently pleased with the search results to make a final selection of a candidate product (or products) that have been systematically gathered during the multi-turn dialog.

[0097] Figure 10 shows a visual product search that generates product suggestions in response to user non-image input, according to some example embodiments. The visual search service 800 performs a visual search based on the image signature it generates for input query image 1002 of a dress, in comparison with image signatures of candidate product images in database 814. The result is an internal ranked list of potentially matching items 818, as previously described.

[0098] However, unlike the previous example, the “bot” acts to provide expertise via a statement prompt, a question prompt, and a candidate product image suggestion prompt via a dialog turn first. Why? In this example, a user also provides a natural language utterance (e.g., text or voice converted to text) input question.

[0099] The parsed input question “How about this for a formal dinner party?” contains the term “formal” that may be recognized as an aspect value 820 that also appears in the knowledge graph, perhaps for the aspect “style”. The term “dinner party” may also be recognized, perhaps as an aspect value for the aspect “occasion” in the knowledge graph. The “bot” therefore identifies the user’s intent from its processing of the natural language utterance. It further recognizes that the visual search results do not strongly correlate with candidate product images for “formal” and “dinner party” dresses. This mismatch would lead to a great deal of filtering, perhaps such that no acceptable candidate products may be found.

[00100] The knowledge graph service 822 therefore elects to generate both a statement type user prompt and a question type user prompt at 824 to solicit

further user input. The statement type user prompt “That looks casual.” therefore denotes the conflicting results between the visual search and the knowledge graph search. The statement explains the recognition of the contradiction in the results of different search types to the user.

[00101] The question type prompt asks the user to consider “more elegant” choices. This behavior is probably because the term “elegant” is associated in the knowledge graph historical data with prior buyer searches and/or seller listings for “formal” and “dinner party” dresses. The knowledge graph may contain sufficiently strong attribute values that link “formal” and “dinner party” to particular candidate products. That is, because many sellers historically described these particular candidate products using the terms “formal” or “dinner party”, they may be recognized as strong candidates. Similarly, if many buyers historically used “formal” or “dinner party” frequently to find these particular candidate products, that also may strengthen the linkage between the user utterance terms and the user input query image.

[00102] The “bot” therefore elects to provide product suggestion type user prompts to determine the user’s true intent, perhaps in spite of some user misunderstanding. Candidate products 1008, 1010, and 1012 are displayed for the user because there is sufficient confidence that the user will proceed to make a final product selection, rather than providing further utterances. This example therefore indicates that the intelligent online personal assistant can provide helpful expert guidance via a variety of output types, when necessary. That is, the user prompts may comprise selectable candidate product images for the highest ranked candidate products, or questions requesting further natural input data describing unresolved aspect values.

[00103] Figure 11 shows a visual product search that generates product suggestions in response to user category input, according to some example embodiments. In this example, a user provides an input query image 1102 of a tent, along with circled suggested categories (“Outdoor Play Tents” and “Camping Tents”). The knowledge graph provides a leaf category prediction in block 1104, based on historical data, e.g., denoting that past sellers and

buyers in aggregate have frequently used the sub-category of “Dome Camping Tent” to find items characterized by the suggested categories.

[00104] The visual search service 800 may perform a visual search based on the image signature it generates for input query image 1102, in comparison with image signatures of candidate product images in database 814. The result is an internal ranked list of potentially matching items 818, as previously described.

[00105] However, the knowledge graph service 822 and the visual search service 800 interaction may have led in this case to a suggested product other than either service may have selected separately. Thus, the filtered result in this case is a candidate product that bears both the leaf category predicted and has a candidate product image that very closely matches the provided input query image 1102. The “bot” therefore has such confidence (e.g., weighted support) that the filtered result satisfies the user’s search interest that the dialog manager 216 elects to output a single product suggestion rather than engage in further dialog turns with the user.

[00106] Figure 12 shows a flowchart of a methodology 1200 for visual search service and knowledge graph service interaction, according to some example embodiments. The methodology may process natural language input and an image query to generate an item recommendation in an electronic marketplace. This methodology may be implemented via the structural elements previously described, as well as via instructions executed by a processor in a computing machine.

[00107] At 1202, the methodology may generate an image signature for each of a plurality of candidate product images and an input query image, where the respective image signatures semantically represent respective visual content features. The image signature generation may be performed by a neural network. At 1204, the methodology may calculate a visual similarity measure between each candidate product image and the input query image based on the corresponding image signatures.

[00108] At 1206, the methodology may assemble a ranked list of candidate products based on the respective visual similarity measures. That is, the

visual similarity measure may serve as a search result score, and candidate products that are the most highly scored may be presented first. The candidate products may have various resolved categories, aspects, and aspect values associated with them, whether determined from a visual search process or a non-visual search process.

[00109] At 1208, for each candidate product in the ranked list, the methodology may provide its corresponding resolved category, aspects, and aspect values to a knowledge graph. The knowledge graph may contain aggregate historical electronic marketplace user interaction information, such as which categories, aspects, and/or aspect values are associated with a given item. These categories, aspects, and/or aspect values may help relate an input query image to a product of interest to the user even if they are not discernible via visual search.

[00110] At 1210, the methodology may generate and output a user prompt that requests further user input in a multi-turn dialog. The user prompt may solicit information used for selecting the next most useful unresolved aspect values in the knowledge graph, for the most highly ranked candidate products. By combining visual search information and a knowledge graph based information, embodiments may iteratively filter the list of candidate products to refine the user query.

[00111] At 1212, the methodology may repeat the previously described visual search and user prompt operations to refine the ranked list of candidate products to better meet the user's interests. This methodology thus effectively transforms both the input query image and natural language user prompt information, in combination, into a list of the most relevant products.

[00112] Figure 13 shows an offline index generation service 1306 that interacts with an on-demand visual search service 800, according to some example embodiments. The present inventors have recognized, among other things, that offline processing of an immense product inventory may help produce data needed for a subsequent online search request. New products/items 1302 may be provided to an electronic marketplace, as items

to be sold for example. Millions of such items may be provided per week for some electronic marketplaces.

[00113] Therefore, the images provided for the new items 1302 may be assembled with item metadata at block 1308. Item metadata may include for example a product identification (ID) number and a universal resource locator (URL) assigned to an item listing page on the electronic marketplace's web site. Each product, which may be a candidate product of interest to a user at some point, may be processed to create descriptive image signatures for one, several, or all of the images provided for it by a seller, for example. Each image signature may comprise a binary hash of a number of floating point numbers that describe various aspects of image content. Thus each image signature is not merely a compressed version of a product image, but instead comprise a concise semantic summarization of its predominant content features.

[00114] A bulk image signature generation block is shown as item 1306. This block may comprise neural network software operating on a server farm to process a vast number of input images in parallel to rapidly produce an image signature for each. Block 1306 may process new item images on a periodic or more continuous ongoing basis. For a large electronic marketplace, a daily batch execution may process all of the images provided for all newly listed products, for example.

[00115] The product images, image signatures, and metadata may be combined and formatted by index generator 1308. The descriptive information for the new "live" products may then be put into an inventory database 814 that may for example contain the entire inventory for the electronic marketplace. The inventory database 814 may be centralized, or it may be replicated to multiple processing sites in a distributed computer network. Local availability of a copy of the inventory database 814 may help a local user view product listings on the electronic marketplace with less network latency, for example.

[00116] An online user who is interested in finding a particular product may then submit an image query 802, to be processed on demand. Image query

802 may comprise a number of input query images. Block 1310 may generate image signatures on demand for the input query images, for example following the same image signature format used by the electronic marketplace for the new product listing images.

[00117] The visual search service 800 may then calculate, on demand, a visual similarity measure between each candidate product image and the input query image based on the respective corresponding image signatures. The visual search service 800 may comprise neural network software executing on a number of hardware processors, for example. The visual search service 800 does not need to compute image signatures for the products available for sale because those have already been generated and stored in the database 814. The combined offline product image processing and the online query image processing may therefore perform a visual query for a user without significant user-discernible delay. This approach to product searching may therefore improve the efficiency of user product searching, and thus effectively improve the operation of the electronic marketplace computer system as a whole.

[00118] Figure 14 shows a flowchart of a methodology 1400 for offline index generation for visual search, according to some example embodiments. The methodology may identify a candidate product in an electronic marketplace based on an on-demand visual comparison between candidate product image visual content and visual content of an image query. This methodology may be implemented via the structural elements previously described, as well as via instructions executed by a processor in a computing machine.

[00119] At 1402, the methodology may generate an image signature for each candidate product image (or a selected subset), where the respective image signatures semantically represent the candidate product image visual content features. The image signature generation may be performed by a neural network. At 1404, the methodology may store, in a database, the image signature along with a product identification (ID) number and a universal resource locator (URL) associated with each candidate product image. The generating and storing may be performed for new products currently listed in

the electronic market place on a periodic basis or a more continuous ongoing basis.

[00120] Since an electronic marketplace may have millions of products/items listed at any given time, the generating and storing may be performed by a plurality of processors (e.g., in a server farm) operating in parallel. Each candidate product may have multiple candidate product images associated with it, and each such image may be similarly processed. The database may be replicated at multiple processing sites in a distributed computer network, to for example provide faster access to users in a given geographic region.

[00121] At 1406, the methodology may generate a further image signature for an input query image. This input query image signature may be generated by a neural network, and may be generated “online” or on demand, with little or no user-discernible delay. At 1408, the methodology may calculate a visual similarity measure between each candidate product image and the input query image. Although described as processing each candidate product image here, the methodology may incorporate other operations that very effectively reduce the product search space so that a more relevant subset of the candidate product images may be compared to the input query image via the respective image signatures.

[00122] At 1410, the methodology may output a ranked list of candidate products based on the visual similarity measure. That is, the visual similarity measure may serve as a search result score, and candidate products that are the most highly scored may be presented first. At 1412, the methodology may update the database with the image signature for the input query image, even if the input query image is not of an item being placed for sale on the electronic marketplace. Thus, although the database may encompass information regarding some or all of the items available for sale on the electronic marketplace, this feature is exemplary and not limiting. This methodology thus effectively transforms the input query image into a list of the most relevant products.

[00123] The specific ordered combinations described herein may improve the overall operation of a computer system used to help a user find products of interest in an electronic marketplace. The input images (e.g., candidate product images and query images) are effectively transformed into a ranked list of products that may most satisfy a user's shopping interests. The overall speed of searching is improved to reduce the time that a user has to spend while still searching a given (typically large) number of candidate products, compared with purely text-based search approaches of the past. The number of search iterations may also be reduced via the specialized approaches described, so that overall efficiency of user interaction with an electronic marketplace is measurably increased.

[00124] Although the subject matter has been described with reference to specific example embodiments, it will be evident that various modifications and changes may be made to these embodiments without departing from the broader scope of the disclosed subject matter. Accordingly, the specification and drawings are to be regarded in an illustrative rather than a restrictive sense. The accompanying drawings that form a part hereof, show by way of illustration, and not of limitation, specific embodiments in which the subject matter may be practiced. The embodiments illustrated are described in sufficient detail to enable those skilled in the art to practice the teachings disclosed herein. Other embodiments may be utilized and derived therefrom, such that structural and logical substitutions and changes may be made without departing from the scope of this disclosure. This Description, therefore, is not to be taken in a limiting sense, and the scope of various embodiments is defined only by any appended claims, along with the full range of equivalents to which such claims are entitled.

[00125] Such embodiments of the inventive subject matter may be referred to herein, individually and/or collectively, by the term "invention" merely for convenience and without intending to voluntarily limit the scope of this application to any single invention or inventive concept if more than one is in fact disclosed. Thus, although specific embodiments have been illustrated and described herein, it should be appreciated that any arrangement calculated to achieve the same purpose may be substituted for the specific

embodiments shown. This disclosure is intended to cover any and all adaptations or variations of various embodiments. Combinations of the above embodiments, and other embodiments not specifically described herein, will be apparent to those of skill in the art upon reviewing the above description.

[00126] The following numbered examples are embodiments.

1. A method for processing natural language input and an image query to generate an item recommendation in an electronic marketplace, the method comprising:

- generating an image signature for each of a plurality of candidate product images and an input query image, the respective image signatures semantically representing respective visual content features, using a neural network;

- calculating a visual similarity measure between each candidate product image and the input query image based on the respective corresponding image signatures, using one or more processors;

- assembling a ranked list of candidate products based on the respective visual similarity measures, the candidate products having various resolved categories, aspects, and aspect values;

- for each candidate product in the ranked list, providing its corresponding resolved category, aspects, and aspect values to a knowledge graph containing aggregate historical electronic marketplace user interaction information;

- outputting a user prompt requesting further user input in a multi-turn dialog for selecting next most useful unresolved aspect values in the knowledge graph for the most highly ranked candidate products, for filtering the list of candidate products;

- repeating the previous operations until a stopping condition occurs; and

- outputting the candidate products selected by the user via the multi-turn dialog.

[00127] 2. The method of example 1, wherein the image query comprises a plurality of input query images, each similarly processed.

[00128] 3. The method of example 1 or example 2, wherein the input query image comprises one of: a photograph, a video frame, a sketch, and a diagram.

[00129] 4. The method of any one of examples 1 to 3, wherein the image signature comprises a binary hash of a number of floating point numbers semantically representing the visual content features.

[00130] 5. The method of any one of examples 1 to 4, wherein the user prompt comprises at least one of: a number of selectable candidate product images for the highest ranked candidate products, and a number of questions requesting further natural language input data describing the unresolved aspect values.

[00131] 6. The method of any one of examples 1 to 5, wherein the resolved categories, aspects, and aspect values are determined from both visual searches and non-visual searches.

[00132] 7. The method of any one of examples 1 to 6, wherein the stopping condition comprises at least one of: a detection of a user intent change, a submission of an entirely new query, and a final user selection of a candidate product.

[00133] 8. A non-transitory computer-readable storage medium having embedded therein a set of instructions which, when executed by one or more processors of a computer, causes the computer to execute operations for processing natural language input and an image query to generate an item recommendation in an electronic marketplace, the operations comprising:

generating an image signature for each of a plurality of candidate product images and an input query image, the respective image signatures semantically representing respective visual content features, using a neural network;

calculating a visual similarity measure between each candidate product image and the input query image based on the respective corresponding image signatures;

assembling a ranked list of candidate products based on the respective visual similarity measures, the candidate products having various resolved categories, aspects, and aspect values;

for each candidate product in the ranked list, providing its corresponding resolved category, aspects, and aspect values to a knowledge graph containing aggregate historical electronic marketplace user interaction information;

outputting a user prompt requesting further user input in a multi-turn dialog for selecting next most useful unresolved aspect values in the knowledge graph for the most highly ranked candidate products, for filtering the list of candidate products;

repeating the previous operations until a stopping condition occurs; and

outputting the candidate products selected by the user via the multi-turn dialog.

[00134] 9. The medium of example 8, wherein the image query comprises a plurality of input query images, each similarly processed.

[00135] 10. The medium of examples 8 or 9, wherein the input query image comprises one of: a photograph, a video frame, a sketch, and a diagram.

[00136] 11. The medium of any one of examples 8 to 10, wherein the image signature comprises a binary hash of a number of floating point numbers semantically representing the visual content features.

[00137] 12. The medium of any one of examples 8 to 11, wherein the user prompt comprises at least one of: a number of selectable candidate product images for the highest ranked candidate products, and a number of questions requesting further natural language input data describing the unresolved aspect values.

[00138] 13. The medium of any one of examples 8 to 12, wherein the resolved categories, aspects, and aspect values are determined from both visual searches and non-visual searches.

[00139] 14. The medium of any one of examples 8 to 13, wherein the stopping condition comprises at least one of: a detection of a user intent change, a submission of an entirely new query, and a final user selection of a candidate product.

[00140] 15. A system for identifying a candidate product in an electronic marketplace based on a visual comparison between candidate product image visual text content and input query image visual text content, the system comprising:

- a memory comprising instructions; and

- one or more hardware-based computer processors, wherein the instructions, when executed by the one or more computer processors, cause the one or more computer processors to perform operations comprising:

 - generating an image signature for each of a plurality of candidate product images and an input query image, the respective image signatures semantically representing respective visual content features, using a neural network;

 - calculating a visual similarity measure between each candidate product image and the input query image based on the respective corresponding image signatures;

 - assembling a ranked list of candidate products based on the respective visual similarity measures, the candidate products having various resolved categories, aspects, and aspect values;

 - for each candidate product in the ranked list, providing its corresponding resolved category, aspects, and aspect values to a knowledge graph containing aggregate historical electronic marketplace user interaction information;

 - outputting a user prompt requesting further user input in a multi-turn dialog for selecting next most useful unresolved aspect values in the knowledge graph for the most highly ranked candidate products, for filtering the list of candidate products;

 - repeating the previous operations until a stopping condition occurs; and

 - outputting the candidate products selected by the user via the multi-turn dialog.

[00141] 16. The system of example 15, wherein the image query comprises a plurality of input query images, each similarly processed.

[00142] 17. The system of examples 15 or 16, wherein the image signature comprises a binary hash of a number of floating point numbers semantically representing the visual content features.

[00143] 18. The system of any one of examples 15 to 17, wherein the user prompt comprises at least one of: a number of selectable candidate product images for the highest ranked candidate products, and a number of questions requesting further natural language input data describing the unresolved aspect values.

[00144] 19. The system of any one of examples 15 to 18, wherein the resolved categories, aspects, and aspect values are determined from both visual searches and non-visual searches.

[00145] 20. The system of any one of examples 15 to 19, wherein the stopping condition comprises at least one of: a detection of a user intent change, a submission of an entirely new query, and a final user selection of a candidate product.

[00146] 21. A machine-readable medium carrying a set of instructions which, when executed by one or more processors of a computer, causes the computer to carry out the method of any one of examples 1 to 7.

CLAIMS

1. A method for processing natural language input and an image query to generate an item recommendation in an electronic marketplace, the method comprising:
 - generating an image signature for each of a plurality of candidate product images and an input query image, the respective image signatures semantically representing respective visual content features, using a neural network;
 - calculating a visual similarity measure between each candidate product image and the input query image based on the respective corresponding image signatures, using one or more processors;
 - assembling a ranked list of candidate products based on the respective visual similarity measures, the candidate products having various resolved categories, aspects, and aspect values;
 - for each candidate product in the ranked list, providing its corresponding resolved category, aspects, and aspect values to a knowledge graph containing aggregate historical electronic marketplace user interaction information;
 - outputting a user prompt requesting further user input in a multi-turn dialog for selecting next most useful unresolved aspect values in the knowledge graph for the most highly ranked candidate products, for filtering the list of candidate products;
 - repeating the previous operations until a stopping condition occurs; and
 - outputting the candidate products selected by the user via the multi-turn dialog.
2. The method of claim 1, wherein the image query comprises a plurality of input query images, each similarly processed.
3. The method of claim 1, wherein the input query image comprises one of: a photograph, a video frame, a sketch, and a diagram.

4. The method of claim 1, wherein the image signature comprises a binary hash of a number of floating point numbers semantically representing the visual content features.
5. The method of claim 1, wherein the user prompt comprises at least one of: a number of selectable candidate product images for the highest ranked candidate products, and a number of questions requesting further natural language input data describing the unresolved aspect values.
6. The method of claim 1, wherein the resolved categories, aspects, and aspect values are determined from both visual searches and non-visual searches.
7. The method of claim 1, wherein the stopping condition comprises at least one of: a detection of a user intent change, a submission of an entirely new query, and a final user selection of a candidate product.
8. A non-transitory computer-readable storage medium having embedded therein a set of instructions which, when executed by one or more processors of a computer, causes the computer to execute operations for processing natural language input and an image query to generate an item recommendation in an electronic marketplace, the operations comprising:
 - generating an image signature for each of a plurality of candidate product images and an input query image, the respective image signatures semantically representing respective visual content features, using a neural network;
 - calculating a visual similarity measure between each candidate product image and the input query image based on the respective corresponding image signatures;
 - assembling a ranked list of candidate products based on the respective visual similarity measures, the candidate products having various resolved categories, aspects, and aspect values;
 - for each candidate product in the ranked list, providing its corresponding resolved category, aspects, and aspect values to a knowledge graph containing aggregate historical electronic marketplace user interaction information;

outputting a user prompt requesting further user input in a multi-turn dialog for selecting next most useful unresolved aspect values in the knowledge graph for the most highly ranked candidate products, for filtering the list of candidate products;

repeating the previous operations until a stopping condition occurs; and

outputting the candidate products selected by the user via the multi-turn dialog.

9. The medium of claim 8, wherein the image query comprises a plurality of input query images, each similarly processed.

10. The medium of claim 8, wherein the input query image comprises one of: a photograph, a video frame, a sketch, and a diagram.

11. The medium of claim 8, wherein the image signature comprises a binary hash of a number of floating point numbers semantically representing the visual content features.

12. The medium of claim 8, wherein the user prompt comprises at least one of: a number of selectable candidate product images for the highest ranked candidate products, and a number of questions requesting further natural language input data describing the unresolved aspect values.

13. The medium of claim 8, wherein the resolved categories, aspects, and aspect values are determined from both visual searches and non-visual searches.

14. The medium of claim 8, wherein the stopping condition comprises at least one of: a detection of a user intent change, a submission of an entirely new query, and a final user selection of a candidate product.

15. A system for identifying a candidate product in an electronic marketplace based on a visual comparison between candidate product image visual text content and input query image visual text content, the system comprising:

a memory comprising instructions; and
one or more hardware-based computer processors, wherein the instructions, when executed by the one or more computer processors, cause the one or more computer processors to perform operations comprising:
generating an image signature for each of a plurality of candidate product images and an input query image, the respective image signatures semantically representing respective visual content features, using a neural network;
calculating a visual similarity measure between each candidate product image and the input query image based on the respective corresponding image signatures;
assembling a ranked list of candidate products based on the respective visual similarity measures, the candidate products having various resolved categories, aspects, and aspect values;
for each candidate product in the ranked list, providing its corresponding resolved category, aspects, and aspect values to a knowledge graph containing aggregate historical electronic marketplace user interaction information;
outputting a user prompt requesting further user input in a multi-turn dialog for selecting next most useful unresolved aspect values in the knowledge graph for the most highly ranked candidate products, for filtering the list of candidate products;
repeating the previous operations until a stopping condition occurs; and
outputting the candidate products selected by the user via the multi-turn dialog.

16. The system of claim 15, wherein the image query comprises a plurality of input query images, each similarly processed.

17. The system of claim 15, wherein the image signature comprises a binary hash of a number of floating point numbers semantically representing the visual content features.

18. The system of claim 15, wherein the user prompt comprises at least one of: a number of selectable candidate product images for the highest ranked candidate


products, and a number of questions requesting further natural language input data describing the unresolved aspect values.

19. The system of claim 15, wherein the resolved categories, aspects, and aspect values are determined from both visual searches and non-visual searches.

20. The system of claim 15, wherein the stopping condition comprises at least one of: a detection of a user intent change, a submission of an entirely new query, and a final user selection of a candidate product.

21. A machine-readable medium carrying a set of instructions which, when executed by one or more processors of a computer, causes the computer to carry out the method of any one of claims 1 to 7.

1/14

**FIG. 1**

2/14




FIG. 2




FIG. 3A




FIG. 3B

4/14




FIG. 4

5/14




FIG. 5

6/14




FIG. 6





FIG. 7

FIG. 8






FIG. 9

10/14




FIG. 10

11/14





FIG. 11

12/14

FIG. 12







FIG. 13

14/14

FIG. 14



INTERNATIONAL SEARCH REPORT

International application No.

PCT/US 17/56116

A. CLASSIFICATION OF SUBJECT MATTER

IPC(8) - G06Q 30/00 (2017.01)

CPC - G06Q 30/0631, G06Q 30/02, G06Q 30/0601, G06Q 30/0641, G06Q 30/0633, G06Q 30/0605, G06Q 30/06

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)

See Search History Document

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

See Search History Document

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

See Search History Document

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2003/0063779 A1 (Wrigley) 03 April 2003 (03.04.2003), entire document especially paras [0073], [0079], [0080], [0177], [0248], [0255]	1-21
A	US 2013/0166276 A1 (Raichelgauz et al.) 27 July 2013 (27.07.2013), entire document	1-21
A	US 2008/0199075 A1 (Gokturk et al.) 21 August 2008 (21.08.2008), entire document	1-21
A	US 2008/0097859 A1 (Schrenk) 24 April 2008 (24.04.2008), entire document	1-21
A	US 6,901,378 B1 (Linker et al.) 31 May 2005 (31.05.2005), entire document	1-21
A	US 5,579,471 A (Barber et al.) 26 November 1996 (26.11.1996), entire document	1-21

☐ 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

11 December 2017 (11.12.2017)

Date of mailing of the international search report

09 JAN 2018

Name and mailing address of the ISA/US

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents
P.O. Box 1450, Alexandria, Virginia 22313-1450

Facsimile No. 571-273-8300

Authorized officer:

Lee W. Young

PCT Helpdesk: 571-272-4300
PCT OSP: 571-272-7774