



(22) Date de dépôt/Filing Date: 2004/05/13

(41) Mise à la disp. pub./Open to Public Insp.: 2004/11/29

(45) Date de délivrance/Issue Date: 2013/03/05

(30) Priorité/Priority: 2003/05/29 (US10/447,399)

(51) Cl.Int./Int.Cl. *G06F 3/00* (2006.01),
G06F 3/16 (2006.01), *G10L 15/22* (2006.01)


(72) Inventeur/Inventor:
WANG, KUANSAN, US

(73) Propriétaire/Owner:
MICROSOFT CORPORATION, US

(74) Agent: SMART & BIGGAR

(54) Titre : COMPREHENSION SYNCHRONE D'OBJETS SEMANTIQUES POUR INTERFACE HAUTEMENT
INTERACTIVE

(54) Title: SEMANTIC OBJECT SYNCHRONOUS UNDERSTANDING FOR HIGHLY INTERACTIVE INTERFACE



(57) Abrégé/Abstract:

A method and system provide a speech input mode which dynamically reports partial semantic parses, while audio captioning is still in progress. The semantic parses can be evaluated with an outcome immediately reported back to the user. The net effect is that task conventionally performed in the system turn are now carried out in the midst of the user turn thereby presenting a significant departure from the turn-taking nature of a spoken dialogue.



SEMANTIC OBJECT SYNCHRONOUS
UNDERSTANDING FOR HIGHLY INTERACTIVE
INTERFACE

ABSTRACT OF THE DISCLOSURE

5 A method and system provide a speech input
mode which dynamically reports partial semantic
parses, while audio captioning is still in progress.
The semantic parses can be evaluated with an outcome
immediately reported back to the user. The net effect
10 is that task conventionally performed in the system
turn are now carried out in the midst of the user
turn thereby presenting a significant departure from
the turn-taking nature of a spoken dialogue.

-1-

SEMANTIC OBJECT SYNCHRONOUS UNDERSTANDING FOR HIGHLY INTERACTIVE INTERFACE

BACKGROUND OF THE INVENTION

5 The present invention relates to access and rendering of information in a computer system. More particularly, the present invention relates to access of information using recognition and understanding.

10 Recently, technology has been advanced to allow the user to access information on a computer system by providing speech commands. Upon receipt of a user command, the computer system performs speech recognition on the user input and further processes the input in order to ascertain the intent of the
15 user in order that the computer system will perform a desired action.

 In some situations, the input provided by the user is incomplete or indefinite, which will require the computer system to solicit further information
20 from the user either in the form of visual or audible prompts. A dialog can thus be established between the user and the computer system, where each takes turns providing questions, answers and/or acknowledgments until the intent of the user is ascertained and an
25 action can be performed. In other situations, creating such a dialog is the preferred mode for interacting with the computer system.

 Speech Application Language Tags (SALT) has been introduced to facilitate speech as a viable
30 input/output modality for modern user interface

-2-

design. The design goal for SALT is to make common speech tasks simple to program, yet allow advanced capabilities with straightforward realization. SALT was designed for many applications. One being, for
5 example, a telephone-based, speech-only application that only interacts with users exclusively through spoken dialogue.

SALT includes speech input and output objects ("listen" and "prompt"), which have a mode design to
10 incorporate technologies to detect the start and the end of the user's turn. Accordingly, many speech applications employ user interfaces that require the user to signal the start of a user turn. Some computer systems include wearable computers, speech
15 enabled modal or multimodal (speech input provided for fields selected by an input device such as a mouse) devices and other eyes-free applications. Nevertheless, in each of these environments, a clean cut definition on the user versus computer system
20 turn in the dialog is still present.

Human conversation however does not generally follow a clean cut, turn-taking dialog between participants. Rather, conversations can include acknowledgements, confirmations, questions by one
25 participant, etc., while the other is providing information that may drastically affect, slightly affect or not even affect the manner in which the speaker is providing information. Human speakers enjoy this natural form of conversation. Likewise,

-3-

telephone systems employ full duplex technology in order to allow such conversations to take place.

In contrast, dialogue based interfaces employ a rigid turn-taking mode of operation between a user and a computer system, which causes the computer system to wait for the end of the user dialog before processing and taking subsequent action. Although simple feedback, such as visual indications like a series of dots progressing across a computer screen, may provide the user some assurance that the computer system is at least processing something, until the user finishes his/her turn and the computer system responds, the extent of understanding by the computer system is not known.

Accordingly, there is a need for improvements in a computer systems that is based on recognition and understanding. Such improvements would provide a system or method for accessing information that would be easier to use by being more natural for the user.

SUMMARY OF THE INVENTION

A method and system provide a speech input mode which dynamically reports partial semantic parses, while audio captioning is still in progress. The semantic parses can be evaluated with an outcome immediately reported back to the user. The net effect is that tasks conventionally performed in the system turn are now carried out in the midst of the user turn thereby presenting a significant departure from the turn-taking nature of a spoken dialogue.

-4-

Generally, one aspect of the present invention includes a computer implemented method for interacting with a computer system. The method includes receiving input from a user and capturing
5 the input for processing. Recognition is then performed on the input, sequentially or simultaneously, to ascertain semantic information pertaining to a first portion of the input and outputting the semantic object. The semantic object
10 includes data in a format to be processed by computer application that is in accordance with the input that has been recognized (e.g. text, encrypted text, compressed text, etc.) as well as semantic information for the first portion. As indicated
15 above, recognition is performed and the semantic object is outputted while capturing continues for subsequent portions of the input. This method can be implemented for audible input such as speech as well as non-audible input, for example, visual input or
20 handwriting recognition.

The above method can be implemented in a computer readable media that includes instructions readable by a computing device which, when implemented, cause the computing device to handle
25 information and perform the method. In a further embodiment, the computer readable medium can include instructions to establish a language model to perform recognition and understanding. The language model is adapted to provide data in a format in accordance

51039-8

5

with the recognized input and to provide semantic information for the input received.

According to one aspect of the present invention, there is provided a computer-implemented method for interacting with a computer system, the method comprising: receiving input comprising at least one command executable by an application from a user and capturing the input for processing; performing recognition on the input to ascertain semantic information pertaining to a first portion of the input and outputting a semantic object comprising data including data for executing the at least one command in a format to be processed by a computer application and being in accordance with the input that has been recognized and semantic information for the first portion, wherein performing recognition and outputting the semantic object are performed using a language model comprising a combination of an N-gram language model and a context free grammar and while capturing continues for subsequent portions of the input, the language model storing information related to words and semantic information to be recognized; and rendering information to the user while the user is providing the input, at least some of the information being responsive to the data regarding the at least one command within said semantic object.

According to another aspect of the present invention, there is provided a computer readable storage media including instructions readable by a computing device which, when implemented, cause the computing device to handle information by performing steps comprising: establishing a language model to perform recognition and understanding, the language model comprising a combination of an N-gram language model and a context free grammar and storing information related to words and semantic information to be recognized, the language model providing data in a format to be processed by a computer application and being in accordance with an input received and to provide semantic information for the input received; receiving input comprising at least one command executable by an application from a user and capturing the input for processing; processing the input with the language model to perform recognition on the input to ascertain semantic information pertaining to a first portion of the input and outputting a semantic object comprising data

51039-8

5a

including data for executing the at least one command in a format to be processed by a computer application and being in accordance with the input that has been recognized and semantic information for the first portion, wherein performing recognition and outputting the semantic object are performed while capturing
5 continues for subsequent portions of the input; and rendering information to the user while the user is providing the input, the information being a function of said semantic object wherein at least some of the information includes information regarding the at least one command.

According to still another aspect of the present invention, there is
10 provided a computer readable storage media including instructions readable by a computing device comprising: a language model comprising a combination of an N-gram language model and a context-free grammar language model and storing information related to words and semantic information to be recognized; a module receiving input including commands executable by an application from a user and
15 capturing the input for processing, the module performing recognition on the input by accessing the language model and ascertaining semantic information pertaining to a first portion of the input and outputting a semantic object comprising data including data for executing commands in a format to be processed by a computer application and being in accordance with the input that has been recognized and semantic
20 information for the first portion, wherein performing recognition and outputting the semantic object are performed while capturing continues for subsequent portions of the input; and a second module performing a selected task, the second module receiving semantic objects synchronously in accordance with the user's input, each semantic object comprising data including commands in a format to be processed by
25 a computer application and in accordance with a portion of input from the user and semantic information pertaining to the portion of input from the user, the module taking action as a function of processing the semantic objects synchronously including providing information or utilizing other applications based on at least the data for executing the commands.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a block diagram of a data presentation system.

FIG. 2 is a plan view of a computing device operating environment.

FIG. 3 is a block diagram of the computing device of FIG. 2.

10 FIG. 4 is a plan view of a telephone.

FIG. 5 is a block diagram of a general purpose computer.

FIG. 6 is a block diagram of an architecture for a client/server system.

15 FIG. 7 is a block diagram of a speech recognition and understanding module.

DETAILED DESCRIPTION OF ILLUSTRATIVE EMBODIMENTS

FIG. 1 is a block diagram of a data presentation system 10 for rendering data based on voice input. 20 System 10 includes a speech interface module 12, a speech recognition and understanding module 14 and a data rendering module 16. A user provides input in the form of a voice query to speech interface module 12. Speech interface module 12 gathers speech information 25 from a user and provides a signal indicative thereof. After the input speech has been gathered by speech interface module 12, speech recognition and understanding module 14 recognizes the speech using a speech recognizer and also performs speech 30 understanding, wherein in one aspect of the present

-6-

invention, provides partial semantic parses of the input so far received, while audio speech input is still being captured.

The partial semantic parses, which typically
5 include the text for the input received (or other data indicative of the text of the input), but also ascertained semantic information, are provided to an application module 16 that can take many different forms. For instance, in one embodiment, the application
10 module 16 can be a personal information manager that is used to send, receive and reply to e-mails, arrange meetings and the like. In this manner, the user can provide audible commands to perform these tasks. More importantly though, the application module 16 can
15 provide interactive feedback and/or take actions on the partial semantic parses information as it is received, thereby providing the user with a highly interactive interface for the application module 16. For example, in a voice only mode of operation, an output 20 can
20 include audible statements back to the user, while, of course, possibly performing other tasks related to the application. The partial semantic parses or semantic objects can be used to execute dialog logic in the application. For instance, the dialog logic can present
25 an option, or a plurality or list of options, to the user based on one or more semantic objects.

This enables the system 10 to immediately report outcomes based on a partial utterance, namely, before the user turn has ended. In other words, by using back
30 channel communication to report and perform tasks

-7-

normally associated with a system turn, the definition of a user and a system turn is blurred. Most conventional dialog studies, especially those based on human to human dialogs, often view the back channel
5 communications as non-intrusive feedback that conveys only simple signals such as positive, negative, or neutral acknowledgement. However, the feedback provided by output 20 can potentially carry more information so as to be somewhat intrusive to the ongoing user
10 utterance, which may or may not cause the user to clarify the user's intention or direction. Nevertheless, this approach provides a more realistic human dialog between the user and the system 10, which in many cases will not be considered annoying, but
15 rather be more comfortable to the user and instill confidence that the user's desires will be met.

At this point it should be noted that the present invention is not limited to a voice-only operating environment, but rather can include other forms of
20 feedback to the user based on processing of partial semantic parses or objects. For example, in the application discussed above where the application module 16 performs e-mail tasks, the output 20 can include visual feedback such as activating an e-mail
25 module based only on receipt of a phrase such as "Send e-mail" in an uninterrupted command from the user comprising "Send e-mail to Bob", where processing of the phrase "to Bob", may cause the application module to access further information in data store 18 and
30 render a list of people having the first name "Bob".

-8-

Whereupon seeing the list, the user simply identifies the intended recipient as "Bob Green", which can then be selected because the system would have provided another semantic object for the partial utterance "Bob
5 Green", which upon receipt by the application and processed would cause "Bob Green" to be selected.

As indicated above, the application module 16 can take many forms where aspects of the present invention discussed further below may be beneficial. Without
10 limitation, the application module 16 can also be a dictation module to provide a textual output of the user's spoken input. However, by processing semantic information for partial input or phrases of input as well, a more accurate transcription can be obtained.

15 Although described above with respect to an input from the user comprising voice commands, aspects of the present invention can also be applied to other forms of input such as handwriting, DTMF, gesture or visual indications.

20 Given the broad applicability of processing partial semantic parses or objects, it may be useful to describe generally computing devices that can function in system 10 described above. As appreciated by those skilled in the art, the components of system 10 may be
25 located within a single computer or distributed across a distributed computing environment using network connections and protocols.

Referring now to FIG. 2, an exemplary form of a mobile device such as a data management device (PIM,
30 PDA or the like) is illustrated at 30. However, it is

-9-

contemplated that the present invention can also be practiced using other computing devices discussed below. For example, phones and/or data management devices will also benefit from the present invention.

5 Such devices will have an enhanced utility compared to existing portable personal information management devices and other portable electronic devices.

In the exemplary form of a data management mobile device 30 is illustrated in FIG. 2, the mobile device

10 30 includes a housing 32 and has an user interface including a display 34, which uses a contact sensitive display screen in conjunction with a stylus 33. The stylus 33 is used to press or contact the display 34 at designated coordinates to select a field, to

15 selectively move a starting position of a cursor, or to otherwise provide command information such as through gestures or handwriting. Alternatively, or in addition, one or more buttons 35 can be included on the device 30 for navigation. In addition, other input mechanisms

20 such as rotatable wheels, rollers or the like can also be provided. However, it should be noted that the invention is not intended to be limited by these forms of input mechanisms. For instance, another form of input can include a visual input such as through

25 computer vision.

Referring now to FIG. 3, a block diagram illustrates the functional components comprising the mobile device 30. A central processing unit (CPU) 50 implements the software control functions. CPU 50 is

30 coupled to display 34 so that text and graphic icons

-10-

generated in accordance with the controlling software appear on the display 34. A speaker 43 can be coupled to CPU 50 typically with a digital-to-analog converter 59 to provide an audible output. Data that is
5 downloaded or entered by the user into the mobile device 30 is stored in a non-volatile read/write random access memory store 54 bi-directionally coupled to the CPU 50. Random access memory (RAM) 54 provides volatile storage for instructions that are executed by CPU 50,
10 and storage for temporary data, such as register values. Default values for configuration options and other variables are stored in a read only memory (ROM) 58. ROM 58 can also be used to store the operating system software for the device that controls the basic
15 functionality of the mobile device 30 and other operating system kernel functions (e.g., the loading of software components into RAM 54).

RAM 54 also serves as a storage for the code in the manner analogous to the function of a hard drive on
20 a PC that is used to store application programs. It should be noted that although non-volatile memory is used for storing the code, it alternatively can be stored in volatile memory that is not used for execution of the code.

25 Wireless signals can be transmitted/received by the mobile device through a wireless transceiver 52, which is coupled to CPU 50. An optional communication interface 60 can also be provided for downloading data directly from a computer (e.g., desktop computer), or
30 from a wired network, if desired. Accordingly,

-11-

interface 60 can comprise various forms of communication devices, for example, an infrared link, modem, a network card, or the like.

Mobile device 30 includes a microphone 29, and
5 analog-to-digital (A/D) converter 37, and an optional
recognition program (speech, DTMF, handwriting, gesture
or computer vision) stored in store 54. By way of
example, in response to audible information,
instructions or commands from a user of device 30,
10 microphone 29 provides speech signals, which are
digitized by A/D converter 37. The speech recognition
program can perform normalization and/or feature
extraction functions on the digitized speech signals to
obtain intermediate speech recognition results. Using
15 wireless transceiver 52 or communication interface 60,
speech data can be transmitted to a remote recognition
server 204 discussed below and illustrated in the
architecture of FIG. 6. Recognition results can then be
returned to mobile device 30 for rendering (e.g. visual
20 and/or audible) thereon, and eventual transmission to a
web server 202 (FIG. 6), wherein the web server 202 and
mobile device 30 operate in a client/server
relationship.

Similar processing can be used for other forms of
25 input. For example, handwriting input can be digitized
with or without pre-processing on device 30. Like the
speech data, this form of input can be transmitted to
the recognition server 204 for recognition wherein the
recognition results are returned to at least one of the
30 device 30 and/or web server 202. Likewise, DTMF data,

-12-

gesture data and visual data can be processed similarly. Depending on the form of input, device 30 (and the other forms of clients discussed below) would include necessary hardware such as a camera for visual
5 input.

FIG. 4 is a plan view of an exemplary embodiment of a portable phone 80. The phone 80 includes a display 82 and a keypad 84. Generally, the block diagram of FIG. 3 applies to the phone of FIG. 4, although
10 additional circuitry necessary to perform other functions may be required. For instance, a transceiver necessary to operate as a phone will be required for the embodiment of FIG. 3; however, such circuitry is not pertinent to the present invention.

15 In addition to the portable or mobile computing devices described above, it should also be understood that the present invention can be used with numerous other computing devices such as a general desktop computer. For instance, the present invention will
20 allow a user with limited physical abilities to input or enter text into a computer or other computing device when other conventional input devices, such as a full alpha-numeric keyboard, are too difficult to operate.

The invention is also operational with numerous
25 other general purpose or special purpose computing systems, environments or configurations. Examples of well known computing systems, environments, and/or configurations that may be suitable for use with the invention include, but are not limited to, regular
30 telephones (without any screen) personal computers,

-13-

server computers, hand-held or laptop devices, tablet computers, multiprocessor systems, microprocessor-based systems, set top boxes, programmable consumer electronics, network PCs, minicomputers, mainframe
5 computers, distributed computing environments that include any of the above systems or devices, and the like.

The following is a brief description of a general purpose computer 120 illustrated in FIG. 5. However,
10 the computer 120 is again only one example of a suitable computing environment and is not intended to suggest any limitation as to the scope of use or functionality of the invention. Neither should the computer 120 be interpreted as having any dependency or
15 requirement relating to any one or combination of components illustrated therein.

The invention may be described in the general context of computer-executable instructions, such as program modules, being executed by a computer.
20 Generally, program modules include routines, programs, objects, components, data structures, etc. that perform particular tasks or implement particular abstract data types. The invention may also be practiced in distributed computing environments where
25 tasks are performed by remote processing devices that are linked through a communications network. In a distributed computing environment, program modules may be located in both local and remote computer storage media including memory storage devices. Tasks
30 performed by the programs and modules are described

-14-

below and with the aid of figures. Those skilled in the art can implement the description and figures as processor executable instructions, which can be written on any form of a computer readable medium.

5 With reference to FIG. 5, components of computer 120 may include, but are not limited to, a processing unit 140, a system memory 150, and a system bus 141 that couples various system components including the system memory to the processing unit 140. The system
10 bus 141 may be any of several types of bus structures including a memory bus or memory controller, a peripheral bus, and a local bus using any of a variety of bus architectures. By way of example, and not limitation, such architectures include Industry
15 Standard Architecture (ISA) bus, Universal Serial Bus (USB), Micro Channel Architecture (MCA) bus, Enhanced ISA (EISA) bus, Video Electronics Standards Association (VESA) local bus, and Peripheral Component Interconnect (PCI) bus also known as
20 Mezzanine bus. Computer 120 typically includes a variety of computer readable mediums. Computer readable mediums can be any available media that can be accessed by computer 120 and includes both volatile and nonvolatile media, removable and non-
25 removable media. By way of example, and not limitation, computer readable mediums may comprise computer storage media and communication media. Computer storage media includes both volatile and nonvolatile, removable and non-removable media
30 implemented in any method or technology for storage

-15-

of information such as computer readable instructions, data structures, program modules or other data. Computer storage media includes, but is not limited to, RAM, ROM, EEPROM, flash memory or
5 other memory technology, CD-ROM, digital versatile disks (DVD) or other optical disk storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other medium which can be used to store the desired information
10 and which can be accessed by computer 120.

Communication media typically embodies computer readable instructions, data structures, program modules or other data in a modulated data signal such as a carrier wave or other transport mechanism and
15 includes any information delivery media. The term "modulated data signal" means a signal that has one or more of its characteristics set or changed in such a manner as to encode information in the signal. By way of example, and not limitation, communication
20 media includes wired media such as a wired network or direct-wired connection, and wireless media such as acoustic, FR, infrared and other wireless media. Combinations of any of the above should also be included within the scope of computer readable media.

25 The system memory 150 includes computer storage media in the form of volatile and/or nonvolatile memory such as read only memory (ROM) 151 and random access memory (RAM) 152. A basic input/output system
153 (BIOS), containing the basic routines that help
30 to transfer information between elements within

-16-

computer 120, such as during start-up, is typically stored in ROM 151. RAM 152 typically contains data and/or program modules that are immediately accessible to and/or presently being operated on by
5 processing unit 140. By way of example, and not limitation, FIG. 5 illustrates operating system 154, application programs 155, other program modules 156, and program data 157.

The computer 120 may also include other
10 removable/non-removable volatile/nonvolatile computer storage media. By way of example only, FIG. 5 illustrates a hard disk drive 161 that reads from or writes to non-removable, nonvolatile magnetic media, a magnetic disk drive 171 that reads from or writes
15 to a removable, nonvolatile magnetic disk 172, and an optical disk drive 175 that reads from or writes to a removable, nonvolatile optical disk 176 such as a CD ROM or other optical media. Other removable/non-removable, volatile/nonvolatile computer storage
20 media that can be used in the exemplary operating environment include, but are not limited to, magnetic tape cassettes, flash memory cards, digital versatile disks, digital video tape, solid state RAM, solid state ROM, and the like. The hard disk drive 161 is
25 typically connected to the system bus 141 through a non-removable memory interface such as interface 160, and magnetic disk drive 171 and optical disk drive 175 are typically connected to the system bus 141 by a removable memory interface, such as interface 170.

-17-

The drives and their associated computer storage media discussed above and illustrated in FIG. 5, provide storage of computer readable instructions, data structures, program modules and other data for the computer 120. In FIG. 5, for example, hard disk drive 161 is illustrated as storing operating system 164, application programs 165, other program modules 166, and program data 167. Note that these components can either be the same as or different from operating system 154, application programs 155, other program modules 156, and program data 157. Operating system 164, application programs 165, other program modules 166, and program data 167 are given different numbers here to illustrate that, at a minimum, they are different copies.

A user may enter commands and information into the computer 120 through input devices such as a keyboard 182, a microphone 183, and a pointing device 181, such as a mouse, trackball or touch pad. Other input devices (not shown) may include a joystick, game pad, satellite dish, scanner, or the like. These and other input devices are often connected to the processing unit 140 through a user input interface 180 that is coupled to the system bus, but may be connected by other interface and bus structures, such as a parallel port, game port or a universal serial bus (USB). A monitor 184 or other type of display device is also connected to the system bus 141 via an interface, such as a video interface 185. In addition to the monitor, computers may also include other

-18-

peripheral output devices such as speakers 187 and printer 186, which may be connected through an output peripheral interface 188.

The computer 120 may operate in a networked
5 environment using logical connections to one or more remote computers, such as a remote computer 194. The remote computer 194 may be a personal computer, a hand-held device, a server, a router, a network PC, a peer device or other common network node, and
10 typically includes many or all of the elements described above relative to the computer 120. The logical connections depicted in FIG. 5 include a local area network (LAN) 191 and a wide area network (WAN) 193, but may also include other networks. Such
15 networking environments are commonplace in offices, enterprise-wide computer networks, intranets and the Internet.

When used in a LAN networking environment, the computer 120 is connected to the LAN 191 through a
20 network interface or adapter 190. When used in a WAN networking environment, the computer 120 typically includes a modem 192 or other means for establishing communications over the WAN 193, such as the Internet. The modem 192, which may be internal or
25 external, may be connected to the system bus 141 via the user input interface 180, or other appropriate mechanism. In a networked environment, program modules depicted relative to the computer 120, or portions thereof, may be stored in the remote memory
30 storage device. By way of example, and not

-19-

limitation, FIG. 5 illustrates remote application programs 195 as residing on remote computer 194. It will be appreciated that the network connections shown are exemplary and other means of establishing a communications link between the computers may be used.

FIG. 6 illustrates architecture 200 for web based recognition and data rendering, which is one exemplary environment for the present invention. Generally, information stored in a web server 202 can be accessed through a client 100 such as mobile device 30 or computer 120 (which herein represent other forms of computing devices having a display screen, a microphone, a camera, a touch sensitive panel, etc., as required based on the form of input), or through phone 80 wherein information is requested audibly or through tones generated by phone 80 in response to keys depressed and wherein information from web server 202 is provided only audibly back to the user.

In this embodiment, architecture 200 is unified in that whether information is obtained through client 100 or phone 80 using speech recognition, a single recognition server 204 can support either mode of operation. In addition, architecture 200 operates using an extension of well-known mark-up languages (e.g. HTML, XHTML, cHTML, XML, WML, and the like). Thus, information stored on web server 202 can also be accessed using well-known GUI methods found in these mark-up languages. By using an extension of

-20-

well-known mark-up languages, authoring on the web server 202 is easier, and legacy applications currently existing can be also easily modified to include voice recognition.

5 Generally, client 100 executes HTML pages, scripts, or the like, generally indicated at 206, provided by web server 202 using a browser. When voice recognition is required, by way of example, speech data, which can be digitized audio signals or
10 speech features wherein the audio signals have been preprocessed by client 100 as discussed above, are provided to recognition server 204 with an indication of a grammar or language model 220 to use during speech recognition, which may be provided by client
15 100. Alternatively, speech server 204 may include the language model 220. The implementation of the recognition server 204 can take many forms, one of which is illustrated, but generally includes a recognizer 211. The results of recognition are
20 provided back to client 100 for local rendering if desired or appropriate. If desired, text-to-speech module 222 can be used to provide spoken text to client 100. Upon compilation of information through recognition and any graphical user interface if used,
25 client 100 sends the information to web server 202 for further processing and receipt of further HTML pages/scripts, if necessary.

As illustrated in FIG. 6, client 100, web server 202 and recognition server 204 are commonly
30 connected, and separately addressable, through a

-21-

network 205, herein a wide area network such as the Internet. It therefore is not necessary that any of these devices be physically located adjacent each other. In particular, it is not necessary that web server 202 includes recognition server 204. In this manner, authoring at web server 202 can be focused on the application to which it is intended without the authors needing to know the intricacies of recognition server 204. Rather, recognition server 204 can be independently designed and connected to the network 205, and thereby, be updated and improved without further changes required at web server 202. Web server 202 can also include an authoring mechanism that can dynamically generate client-side markups and scripts. In a further embodiment, the web server 202, recognition server 204 and client 100 may be combined depending on the capabilities of the implementing machines. For instance, if the client 100 comprises a general purpose computer, e.g. a personal computer, the client may include the recognition server 204. Likewise, if desired, the web server 202 and recognition server 204 can be incorporated into a single machine.

Access to web server 202 through phone 80 includes connection of phone 80 to a wired or wireless telephone network 208, that in turn, connects phone 80 to a third party gateway 210. Gateway 210 connects phone 80 to a telephony voice browser 212. Telephony voice browser 212 includes a media server 214 that provides a telephony interface

-22-

and a voice browser 216. Like client 100, telephony voice browser 212 receives HTML pages/scripts or the like from web server 202. In one embodiment, the HTML pages/scripts are of the form similar to HTML pages/scripts provided to client 100. In this manner, web server 202 need not support client 100 and phone 80 separately, or even support standard GUI clients separately. Rather, a common mark-up language can be used. In addition, like client 100, voice recognition from audible signals transmitted by phone 80 are provided from voice browser 216 to recognition server 204, either through the network 205, or through a dedicated line 207, for example, using TCP/IP. Web server 202, recognition server 204 and telephone voice browser 212 can be embodied in any suitable computing environment such as the general purpose desktop computer illustrated in FIG. 5.

Having described various environments and architectures functioning in system 10, a more detailed description of various components and the function of system 10 is provided. FIG. 7 illustrates a block diagram of speech recognition and understanding module 14. Input speech received from speech interface module 12 is sent to speech recognition and understanding module 14. Speech recognition and understanding module 14 includes a recognition engine 306, which has an associated language model 310. Recognition engine 306 uses language model 310 to identify possible surface semantic structures to represent each of the phrases

-23-

forming the input, providing partial semantic parses or objects as the input is received. Unlike systems that wait for the user to complete the utterance and then process the complete input received, the module
5 14 continuously provides semantic objects based only what has so far been received.

Recognition engine 306 provides at least one surface semantic output object based on the partial utterance. In some embodiments, the recognition
10 engine 306 is capable of providing more than one alternative surface semantic object for each alternative structure.

Although illustrated in FIG. 7 wherein speech input is provided, the present invention can be used
15 with handwriting recognition, gesture recognition or graphical user interfaces (where the user interacts with a keyboard or other input device). In these other embodiments, the speech recognizer 306 is replaced with a suitable recognition engine as is
20 known in the art. For graphical user interfaces, a grammar (having the language model) is associated with the user input such as through an input box. Accordingly, a user's input is processed in a consistent way without significant modification based
25 on the manner of input.

The interactive dialog discussed above, which also includes other forms of informative feedback by the system 10 based on partial semantic parses or objects can be implemented using SALT (Speech
30 Application Language Tags) or other speech,

-24-

handwriting, and pattern recognition APIs (Application Program Interfaces) that support a language model construct that can provide semantic information based on a selected schema for a given application and semantic object synchronous decoding. SALT is a developing standard for enabling access to information, applications and web services from personal computers, telephones, tablet PCs and wireless mobile devices, for example, but can also be applied to application interfaces as well without interconnection across a network. SALT extends existing markup languages such as HTML, XHTML and XML. The SALT 1.0 specification may be found online at <http://www.SALTforum.org>. It should be noted that SALT can provide semantic information based upon the user's input, for example from speech server 204, which such information forms the objects provided to data rendering module 16 after completion of an utterance; however as discussed further below SALT can be used in manner not previously contemplated to provide partial semantic parses or objects. Use of SALT extensions, or similar extensions in other APIs, provides support for highly interactive event driven user interaction.

Using SALT by way of example, the SALT *listen* object can be used to perform both speech recognition and understanding tasks. This is because the design follows the viewpoint and formulation that treats speech understanding as a pattern recognition problem, just like speech recognition. Both strive to

-25-

find a pattern from a collection of possible outcomes that best match a given speech signal. For speech recognition, the pattern to be found is a string of words whereas, for understanding, a tree of semantic
5 objects. A traditional speech recognition task instructs the search process with a language model in composing the likely word strings. In a similar fashion, a speech understanding task can guide the same search engine to compose the suitable semantic
10 object trees with a semantic model. Like a language model that often implies a lexicon and the rules of composing phrase segments from the lexicon entries, a semantic model implies a dictionary of all semantic objects and the rules of composing them. While the
15 recognition outcome is a text string, the understanding result is a tree of semantic objects.

Although it is possible to extend N-gram to return a structured search outcome, most typical speech understanding applications are based on
20 probabilistic context free grammar (PCFG) where the designers can specify rules of composing semantic objects without massive tree-bank annotated training data. One method of specifying such rules is to associate each PCFG rule with production directives
25 for the search engine on how to transform the partial PCFG parse tree into a semantic object tree. An example written in Microsoft Speech Application Interface (SAPI) (which is also an example of a speech API that can be used in the present invention)
30 format is shown below:

-26-

```

<rule name="nyc">
  <list>
    <phrase>new york ?city</phrase>
5    <phrase>?the big apple</phrase>
  </list>
  <output>
    <city_location>
      <city>New York</city>
10    <state>New York</state>
      <country>USA</country>
    </city_location>
  </output>
</rule>
15 ...
<rule name="NewMeeting">
  <ruleref min="0" name="CarrierPhrase"/>
  <ruleref max="inf" name="ApptProperty"/>
  <output>
20    <NewMeeting>
      <DateTime>
        <xsl:apply-templates select="//Date"/>
        <xsl:apply-templates select="//Time"/>
        <xsl:apply-templates select="//Duration"/>
25    </DateTime>
      <Invitees>
        <xsl:apply-templates select="//Person"/>
      </Invitees>
      ....
30    </NewMeeting>

```

-27-

```

    </output>
  </rule>

  <rule name="ApptProperty"/>
5    <list>
      <ruleref name="Date"/>
      <ruleref name="Duration"/>
      <ruleref name="Time"/>
      <ruleref name="Person" max="inf"/>
10    <ruleref name="ApptSubject"/>
      .. ..
    </list>
  </rule>
  .. ..
15
  The grammar segment contains three rules. The first
  one, a pre-terminal named "nyc" lists the expressions
  for New York City. The <output> tags in this example
  enclose the rules for constructing semantic objects.
20 They are invoked when the search path exits the
  grammar node denoted by the token immediately
  preceding it. In the case, a semantic object,
  represented in XML with a <city_location> element, is
  created when a search path exits the "nyc" rule. This
25 semantic object is in turn composed of three semantic
  objects: the city name, state and country name
  abbreviations, respectively.

  The composition of semantic objects can also be
  a dynamic process, for example, in scheduling a new
30 meeting. For instance, a NewMeeting semantic object

```

-28-

will be produced when the user finishes specifying the meeting properties such as date, time, duration and attendees. Templates can be used to paste other semantic objects as constituents into the NewMeeting semantic object. The same principle can also be applied to other rules not shown here. As an example, an utterance "schedule a meeting with Li Deng and Alex Acero on January first for one hour" will result in the following semantic object:

10

```
<NewMeeting>
```

```
  <DateTime>
```

```
    <Date>01/01/2003</Date>
```

```
    <Duration>3600</Duration>
```

15

```
  </DateTime>
```

```
  <Invitees>
```

```
    <Person>Li Deng</Person>
```

```
    <Person>Alex Acero</Person>
```

```
  </Invitees>
```

20

```
</NewMeeting>
```

In real applications, improving PCFG coverage is a daunting task. It is therefore desirable to be able to use an N-gram to model, among other things, *functional phrases* that do not carry critical semantic information but usually have sizeable variations in the syntactic structure (e.g., "May I...", "Could you show me...", "Please show me..."). In one embodiment, the language model 310 comprises a semantic language model that combines PCFG with N-

30

-29-

gram. The technique is slightly different from a unified language model, which can also be used. The unified language model is a natural extension to the conventional class N-gram except it allows CFG
5 fragments, not just a list of words, to be modeled as an individual token in N-gram. The recognizer 306 using this model still produces text string that has to be subsequently parsed. The unified language model thus is designed to incorporate certain linguistic
10 structure to assist text transcription.

The semantic language model, on the other hand, aims at using the decoder or recognizer to search for the semantic structure, which is usually better captured by PCFG. Therefore, instead of embedding CFG
15 fragments into N-gram, the PCFG is used to contain N-gram by creating a special PCFG pre-terminal that corresponds to an N-gram of interests. In Microsoft SAPI grammar format, this can be denoted using a pre-terminal with an XML <dictation> tag, as in

20

```
LCFG <dictation max="inf"/> RCFG
```

where LCFG and RCFG denotes the left and right context of the embedded N-gram, respectively. The
25 search process treats the <dictation> tag as a token and expands into the N-gram as if entering a regular non-terminal. The max attribute on the tag specifies the maximum number of words that can be consumed by the N-gram. Inside this N-gram, the word string
30 probability is computed by interpolating the backoff

51039-8

-30-

N-gram with the PCFG, more specifically,

$$P(w_n | w_{n-1}, w_{n-2}, \dots) = \lambda P(w_n | Ngram, w_{n-1}, w_{n-2}, \dots) + (1 - \lambda) P(w_n | RCFG) P(RCFG | w_{n-1}, w_{n-2}, \dots) \quad (1)$$

where λ is the N-gram interpolation weight and $P(RCFG | w_{n-1}, \dots)$ uses the back-off probability of the N-gram, i.e., w_n is treated as if it is an out of vocabulary word. In one embodiment, the term $P(w_n | RCFG)$ assumes only binary value depending on whether the maximum N-gram word string size is reached and the word is in the coverage of the CFG fragment or not. Since words drawn from PCFG often have a higher probability, paths that really belong to be covered by CFG have tendency to win out their N-gram counterparts even when the maximum N-gram word count is set to infinite. In addition to functional phrases, the embedded N-gram can also be used to model semantic object with a dictation like property. For example, the meeting subject is model in our task as

```
20 <rule name="ApptSubject">
    <p> <dictation max="inf"/> </p>
```

Further details regarding the semantic language model are described in K. Wang, "Semantic modeling for dialog systems in a pattern recognition framework," in *Proc. ASRU-2001*, Trento Italy, 2001.

51039-8

-31-

A further aspect of the present invention includes a new use of the SALT listen object . SALT provides a set of XML elements, with associated attributes and DOM object properties, events and methods, which may be used in conjunction with a source markup document to apply a speech interface to a source page. In general, the main elements include:

10

<prompt ...> for speech synthesis configuration and prompt playing

<listen ...> for speech recognizer configuration, recognition execution and post-processing, and recording

15

<dtmf ...> for configuration and control of DTMF

<smex ...> for general purpose communication with platform components

20 The listen and dtmf objects also contain grammar and binding controls:

<grammar ...> for specifying input grammar resources

<bind ...> for processing of recognition results

25

The listen element can include a "mode" attribute to distinguish three modes of recognition, which instruct the recognition server (e.g. 204) how and when to return results. The return of results implies providing the "onReco" event or activating the "bind" elements as appropriate.

30

In a first mode, "automatic", the speech platform, rather than the application, is in control of when to stop the recognition process. This mode

-32-

has been advanced for telephony or hands-free scenarios. As soon as a recognition result is available, and/or where a time period has elapsed indicating silence, the speech platform
5 automatically stops the recognizer and returns its result, which can then be associated with the appropriate field through the bind element.

In a second mode of operation, "single", the return of a recognition result is under the control
10 of an explicit "stop" call. The stop call would correspond to an event such as "pen-up" by the user, and this mode is advanced for being used in a multimodal environment where the device allows speech input, but the user is generally under control of
15 which and when to select fields for example through the use of the stylus 33 (FIG.1).

A third mode of operation of the speech recognizer is a "multiple mode". This mode of operation is used for an "open-microphone" or in a
20 dictation scenario. Generally, in this mode of operation, recognition results are returned at intervals until an explicit stop call is received or the or other time periods associated with unrecognized input or a maximum time to listen are
25 exceeded. Generally, in this mode of operation, for each phrase that is recognized, a "onReco" event is issued and the result is returned until the stop () call is received.

However this mode of operation, as another
30 aspect of the present invention, can provide a means

-33-

for the search engines to expose more interactive capabilities to the users by allowing them to report immediately whenever a salient linguistic landmark is reached. Search algorithms based on time synchronous
5 decoding are well known and can be employed for this mode in a straightforward manner. One such algorithm is described in H. Ney, S. Ortmanns, "Dynamic programming search for continuous speech recognition," *IEEE Signal Processing Magazine*, pp.
10 64-83, 1999. For speech recognition, the linguistic landmark usually corresponds to a word or a phrase boundary. A SALT *multiple* mode recognition can therefore be used to display dynamically the word string hypotheses as soon as they are available, a UI
15 effect commonly seen in many commercial dictation software. However, in the present invention, the *multiple* mode can treat the instantiations of semantic objects as linguistic landmarks and report, i.e. provide some response as a function of what was
20 understood, back to the application on them in a dynamic manner as well. This appears to the application designer as if SALT is performing a semantic object synchronous understanding.

This mode of operation may be better understood
25 by comparing it with a multimodal scenario. In the multimodal scenario, a user indicates a field, for example, by pointing and holding the stylus in an input field while speaking. While the user can tap onto a generic field and utter an elaborated sentence
30 to fill many fields in a single sentence, the tap-

-34-

and-talk interface nevertheless engages the eyes and hands of the user, a scheme not appropriate for many situations. In addition, although tap-and-talk features a rich back channel communication that displays the volume and a progress bar of the underlying spoken language process, those feedbacks provide only very primitive clues to the quality of the spoken language processing in terms of speed and accuracy. This can be potentially more problematic for longer sentences in which errors can propagate to a wider scope that eventually requires more efforts in merely verifying and correcting the recognition and understanding outcomes. Since the usability studies seem to indicate that long sentences are a key differentiating factor that demonstrates the utility of speech as more than a keyboard enhancement or alternative, a satisfying UI experience is absolutely necessary to the success of using speech as a viable modality.

20 In order to promote the perception of human-computer as closely collaborative partners in achieving a common goal, semantic object synchronous understanding, by reporting the partial semantic parses or objects as soon as they are available, is effective. In one embodiment, this is achieved by using the *multiple* mode of the listen element in SALT. In particular, for the listen element a multiple mode is designated and all of the recognition grammars are then specified for the input speech to be recognized. The assignment of results is

-35-

also specified in the listen object. For example, HTML input code to acquire necessary information, such as date, time, location, subject, and meeting attendees, etc., for creating a new meeting might
 5 take the form:

```

    <listen mode="multiple" ...>
      <grammar src="subject.grm"/>
      <grammar src="date.grm"/>
10    <grammar src="time_duration.grm"/>
      <grammar src="attendees.grm"/>
      <bind targetElement="subject"
        value="//ApptSubject"/>
      <bind targetElement="date"
15    value="//DateTime"/>
      <bind targetElement="start_time"
        value="//start_time"
        targetElement="end_time"
        value="//end_time"
20    targetElement="duration"
        value="//DateTime/duration"/>
      ...
    </listen>

```

25 The multiple grammars compose a parallel search space for the recognition with a null transition looping back to entry point. In this mode, SALT allows the listen object to raise an event as soon as a grammar is exited. The event forks a parallel
 30 process to invoke the bind directives in sequence

-36-

while the underlying audio collection and recognition are ongoing, thus creating the effect to the user that relevant fields on the form are being filled while a spoken command is still being uttered for an application that has a visual rendering of fields.

For user interface for eyes free applications, accompanied speech outputs might be desired. In that case, SALT *prompt* objects can be used to give immediate feedbacks. For example, the following SALT *prompt* object can be used to synthesize response based on the dynamic contents in the date field, and the speech synthesis can be triggered with additional SALT *bind* directives as follows:

```

15  <prompt id="say_date">
      on <value targetElement="date"/>
    </prompt>
    ...
    <listen ...>
20    ...
      <bind targetElement="date"
        value="//date"
        targetElement="say_date"
        targetMethod="Start"/>
25    ...
    </listen>

```

The net effect is the user feels like he/she is talking to another party that not only jots down but repeats what is heard, as in "Schedule a meeting (new

51039-8

-37-

meeting) at two (starting at two o'clock PM) next Tuesday (on 10/29/02) for two hours (duration: two hours)", where the phrases provided in parentheses represent the audible and/or visual prompts (which
5 also can be synchronized) back to the user.

It should be noted that SALT allows designers to attach customized recognition event handlers that perform sophisticated computations beyond the simple assignments as with the SALT *bind* directives. In the
10 above example, the date normalization can be accomplished in the semantic grammar which, however, cannot facilitate advanced reference resolution (e.g., "Schedule a meeting with Li Deng and his manager"). For such cases, algorithms can be
15 implemented as script objects accessible to proper event handlers to access stored data to ascertain indefinite references. Such algorithms are described in K. Wang, "A plan based dialog system with probabilistic inferences", in *Proc. ICSLP-2000*,
20 Beijing China, 2000 and European Patent Application EP 1199630A2, published April 24, 2002.

It should be noted that although the multiple mode operation existed for the listen object, in
25 current implementations, this mode provides only text for the received input such as in the dictation scenario. However, in this aspect of the present invention, the partial results as input is being received is not only the text, but also includes the
30 corresponding semantic information pertaining to the

51039-8

-38-

text, and thus the output comprises partial semantic
parses or objects, which can be used as described
above to provide the user more quality feedback that
the computer has properly understood what has been
5 received. Depending on the sophistication of the
application receiving the partial semantic parses or
objects, the system can provide back to the user
confirmations, alternatives, corrections and
clarifications based on the partial semantic parses
10 received.

Although in multimodal applications it is known to
include multiple grammars to allow the user to speak
ahead, and thereby, provide the ability to provide
information that has not been indicated, using the
15 listen element in the multiple mode operation can be
better because it provides to the user a higher
indication of understanding. In a voice-only
application, a natural form of dialog is generated,
whereas if using visual renderings are employed, an
20 application can begin processing (taking actions,
displaying interim results or options e.g. through
pop-up windows) based only on partial semantic parses
of what the user has so far provided and while the
user continues speaking.

25 Although the present invention has been
described with reference to particular embodiments,
workers skilled in the art will recognize that
changes may be made in form and detail without
departing from the scope of the invention.

51039-8

39

CLAIMS:

1. A computer-implemented method for interacting with a computer system, the method comprising:

5 receiving input comprising at least one command executable by an application from a user and capturing the input for processing;

10 performing recognition on the input to ascertain semantic information pertaining to a first portion of the input and outputting a semantic object comprising data including data for executing the at least one command in a format to be processed by a computer application and being in accordance with the input that has been recognized and semantic information for the first portion, wherein performing recognition and outputting the semantic object are performed using a language model comprising a combination of an N-gram language model and a context free grammar and while capturing continues for subsequent portions of the input, the language model storing information related to words and semantic information to be
15 recognized; and

rendering information to the user while the user is providing the input, at least some of the information being responsive to the data regarding the at least one command within said semantic object.

2. The computer-implemented method of claim 1 wherein outputting the semantic object comprises outputting the data in a textual form.
20

3. The computer-implemented method of claim 2 wherein receiving and capturing input comprises receiving and capturing audible input from the user.

4. The computer-implemented method of claim 3 wherein receiving and capturing input comprises receiving and capturing speech input from the user.

51039-8

40

5. The computer-implemented method of claim 1 wherein receiving and capturing input comprises receiving and capturing handwriting input from the user.

6. The computer-implemented method of claim 1 wherein receiving and capturing input comprises receiving and capturing visual input from the user.

5 7. The computer-implemented method of claim 1 and further comprising rendering information to the user while the user is providing the input, the information being a function of the semantic object received.

8. The computer-implemented method of claim 7 wherein rendering information comprises acknowledging what has been received by rendering
10 information as a function of the semantic information for the semantic object.

9. The computer-implemented method of claim 8 wherein rendering information comprises rendering information as a function of the data in accordance with the input that has been recognized.

10. The computer-implemented method of claim 9 wherein rendering
15 information comprising rendering an option to the user as a function of the data in accordance with the input that has been received.

11. The computer-implemented method of claim 10 wherein rendering information comprises rendering a plurality of options to the user as a function of the data in accordance with the input that has been received.

20 12. The computer-implemented method of claim 10 wherein rendering information comprises rendering an audible prompt to the user.

13. The computer-implemented method of claim 11 wherein rendering information comprises rendering visual indications to the user.

51039-8

41

14. The computer-implemented method of claim 13 wherein rendering information comprises rendering synchronized audible and visual indications to the user.

15. A computer readable storage media including instructions readable by a computing device which, when implemented, cause the computing device to handle information by performing steps comprising:

establishing a language model to perform recognition and understanding, the language model comprising a combination of an N-gram language model and a context free grammar and storing information related to words and semantic information to be recognized, the language model providing data in a format to be processed by a computer application and being in accordance with an input received and to provide semantic information for the input received;

receiving input comprising at least one command executable by an application from a user and capturing the input for processing;

15 processing the input with the language model to perform recognition on the input to ascertain semantic information pertaining to a first portion of the input and outputting a semantic object comprising data including data for executing the at least one command in a format to be processed by a computer application and being in accordance with the input that has been recognized and semantic information for the first portion, wherein performing recognition and outputting the semantic object are performed while capturing continues for subsequent portions of the input; and

rendering information to the user while the user is providing the input, the information being a function of said semantic object wherein at least some of the information includes information regarding the at least one command.

25 16. The computer readable storage media of claim 15 wherein establishing a language model comprises defining multiple grammars to be used for recognition.

51039-8

42

17. The computer readable storage media of claim 16 wherein establishing a language model comprises using application program interfaces to define the multiple grammars to be used for recognition.

18. The computer readable storage media of claim 17 wherein establishing
5 a language model comprises using SALT.

19. The computer readable storage media of claim 18 wherein processing the input comprise implementing a listen object of SALT in multiple mode.

20. The computer readable storage media of claim 15 wherein outputting the semantic object comprises outputting the data in a textual form.

10 21. The computer readable storage media of claim 20 wherein receiving and capturing input comprises receiving and capturing audible input from the user.

22. The computer readable storage media of claim 21 wherein receiving and capturing input comprises receiving and capturing speech input from the user.

15 23. The computer readable storage media of claim 15 wherein receiving and capturing input comprises receiving and capturing handwriting input from the user.

24. The computer readable storage media of claim 15 wherein receiving and capturing input comprises receiving and capturing visual input from the user.

20 25. The computer readable storage media of claim 15 and further comprising rendering information to the user while the user is providing the input, the information being a function of the semantic object received.

26. The computer readable storage media of claim 25 wherein rendering information comprises acknowledging what has been received by rendering information as a function of the semantic information for the semantic object.

51039-8

43

27. The computer readable storage media of claim 26 wherein rendering information comprises rendering information as a function of the data in accordance with the input that has been recognized.

28. The computer readable storage media of claim 27 wherein rendering
5 information comprising rendering an option to the user as a function of the data in accordance with the input that has been received.

29. The computer readable storage media of claim 28 wherein rendering information comprises rendering a plurality of options to the user as a function of the data in accordance with the input that has been received.

10 30. The computer readable storage media of claim 25 wherein rendering information comprises rendering an audible prompt to the user.

31. The computer readable storage media of claim 30 wherein rendering information comprises rendering visual indications to the user.

32. The computer readable storage media of claim 31 wherein rendering
15 information comprises rendering synchronized audible and visual indications to the user.

33. A computer readable storage media including instructions readable by a computing device comprising:

a language model comprising a combination of an N-gram language
20 model and a context-free grammar language model and storing information related to words and semantic information to be recognized;

a module receiving input including commands executable by an application from a user and capturing the input for processing, the module performing recognition on the input by accessing the language model and ascertaining semantic
25 information pertaining to a first portion of the input and outputting a semantic object

51039-8

44

comprising data including data for executing commands in a format to be processed by a computer application and being in accordance with the input that has been recognized and semantic information for the first portion, wherein performing recognition and outputting the semantic object are performed while capturing
5 continues for subsequent portions of the input; and

a second module performing a selected task, the second module receiving semantic objects synchronously in accordance with the user's input, each semantic object comprising data including commands in a format to be processed by a computer application and in accordance with a portion of input from the user and
10 semantic information pertaining to the portion of input from the user, the module taking action as a function of processing the semantic objects synchronously including providing information or utilizing other applications based on at least the data for executing the commands.

34. The computer readable storage media of claim 33 wherein the module
15 provides a semantic object that includes the data in a textual form.

35. The computer readable storage media of claim 34 wherein the module captures audible input from the user.

36. The computer readable storage media of claim 35 wherein the module captures speech input from the user.

20 37. The computer readable storage media of claim 33 wherein the module captures handwriting input from the user.

38. The computer readable storage media of claim 33 wherein the module captures visual input from the user.

39. The computer readable storage media of claim 33 wherein the second
25 module processes the data in a textual form.

51039-8

45

40. The computer readable storage media of claim 33 wherein the second module takes action comprising rendering information to the user while semantic objects related to subsequent portions of the user's input are further received.

41. The computer readable storage media of claim 33 wherein the
5 information rendered is a function of the semantic information for the semantic object.

42. The computer readable storage media of claim 39 wherein the information rendered is a function of the data in accordance with the input.

43. The computer readable storage media of claim 40 wherein the information rendered comprises an option to the user as a function of the semantic
10 objects received.

44. The computer readable storage media of claim 43 wherein rendering information comprises rendering a plurality of options to the user as a function of the semantic objects received.

45. The computer readable storage media of claim 40 wherein information
15 rendered comprises an audible prompt to the user.

46. The computer readable storage media of claim 45 wherein the information comprises rendering visual indications to the user.

47. The computer readable storage media of claim 46 wherein rendering information comprises rendering synchronized audible and visual indications to the
20 user.

48. The computer readable storage media of claim 33 wherein utilizing other applications based upon at least the data regarding commands comprises accessing an e-mail server and scheduling a meeting with a third party.

Application number numéro de demande: 2467134

Figures: 2.

Pages: _____

Drawing

Unscannable items
received with this application
(Request original documents in File Prep. Section on the 10th floor)

Documents reçu avec cette demande ne pouvant être balayés
(Commander les documents originaux dans la section de préparation des dossiers au
10^{ème} étage)




FIG. 1




FIG. 3




FIG. 4




FIG. 5




FIG. 6




FIG. 7

