



(19) **United States**

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

(10) **Pub. No.: US 2012/0123777 A1**

(43) **Pub. Date: May 17, 2012**

(54) **ADJUSTING A SPEECH ENGINE FOR A MOBILE COMPUTING DEVICE BASED ON BACKGROUND NOISE**

Publication Classification

(51) **Int. Cl.**
G10L 15/20 (2006.01)

(52) **U.S. Cl.** **704/233; 704/E15.039**

(57) **ABSTRACT**

Methods, apparatus, and products are disclosed for adjusting a speech engine for a mobile computing device based on background noise, the mobile computing device operatively coupled to a microphone, that include: sampling, through the microphone, background noise for a plurality of operating environments in which the mobile computing device operates; generating, for each operating environment, a noise model in dependence upon the sampled background noise for that operating environment; and configuring the speech engine for the mobile computing device with the noise model for the operating environment in which the mobile computing device currently operates.

(75) Inventors: **Ciprian Agapi**, Hollywood, FL (US); **William K. Bodin**, Austin, TX (US); **Charles W. Cross, JR.**, Wellington, FL (US); **Paritosh D. Patel**, Parkland, FL (US)

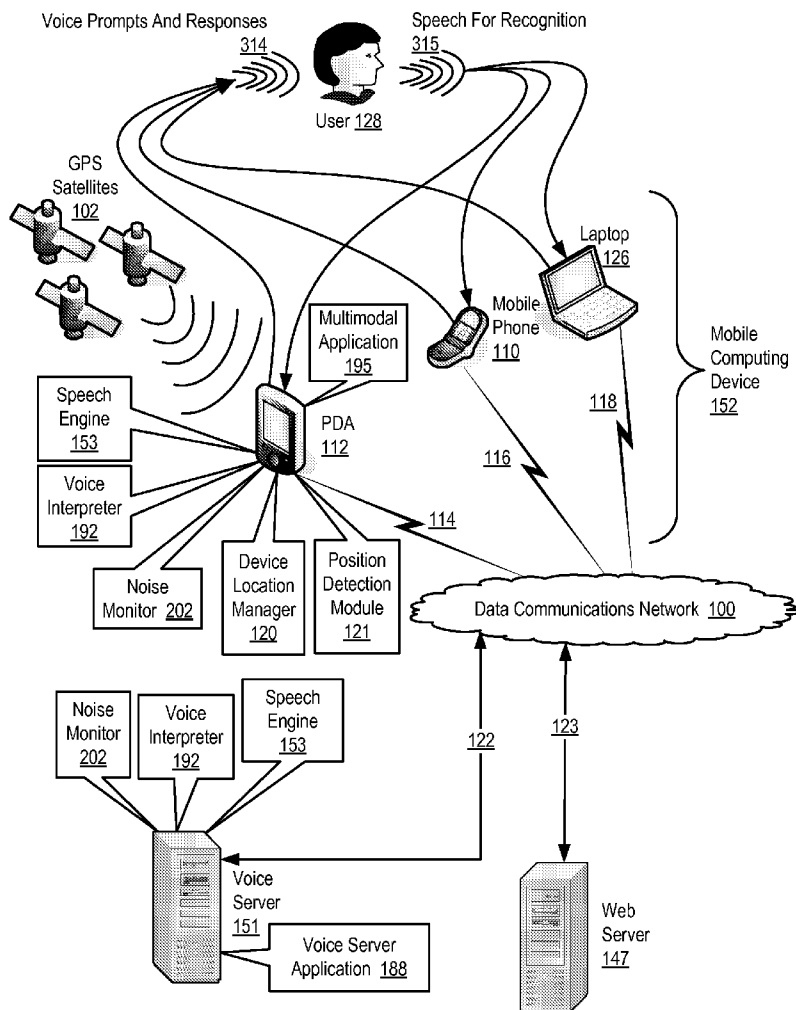
(73) Assignee: **Nuance Communications, Inc.**, Burlington, MA (US)

(21) Appl. No.: **13/358,185**

(22) Filed: **Jan. 25, 2012**

Related U.S. Application Data

(63) Continuation of application No. 12/109,151, filed on Apr. 24, 2008, now Pat. No. 8,121,837.



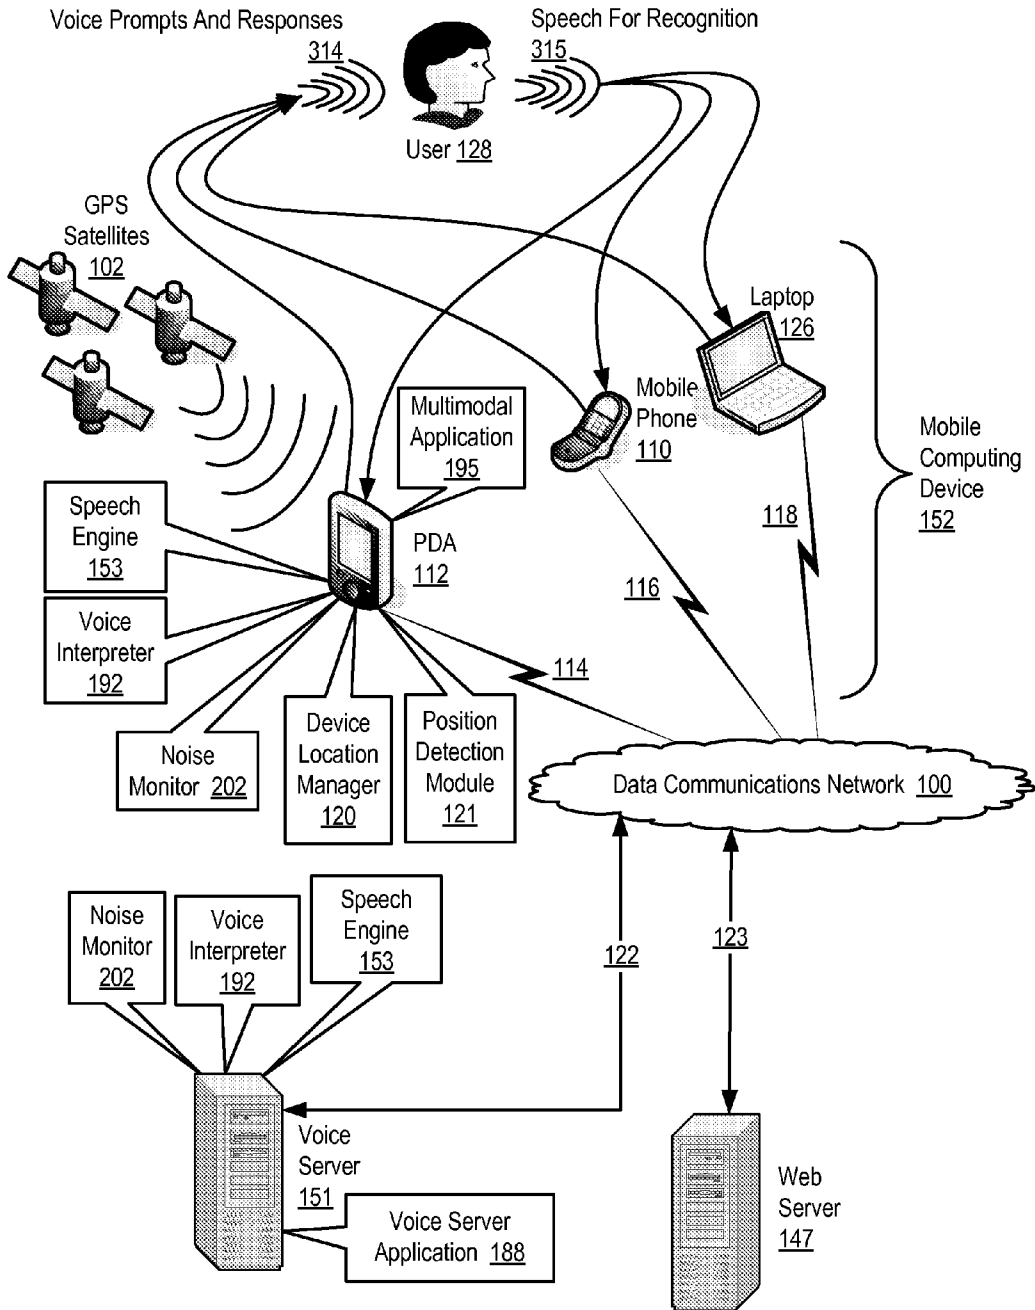


FIG. 1




FIG. 2

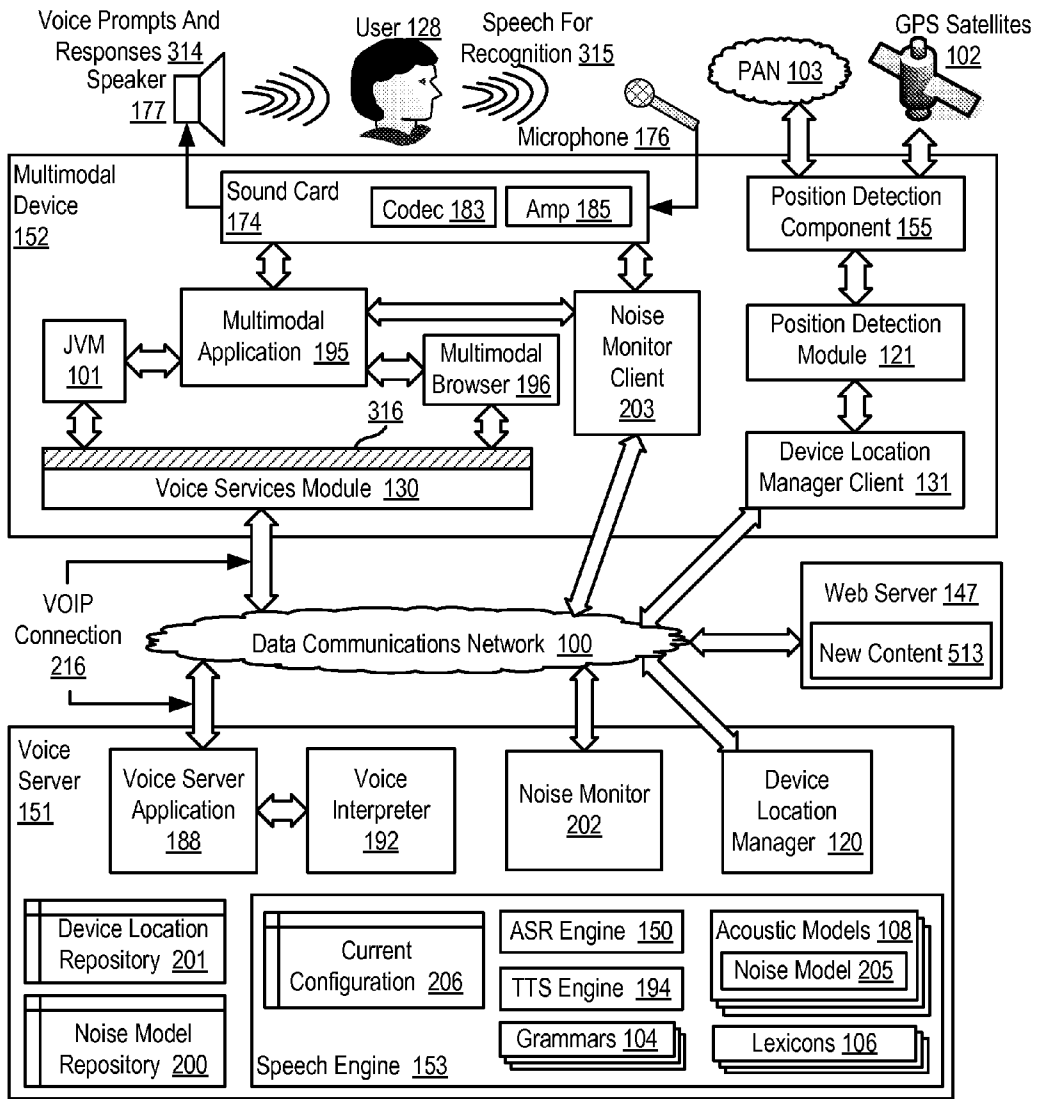


FIG. 3




FIG. 4

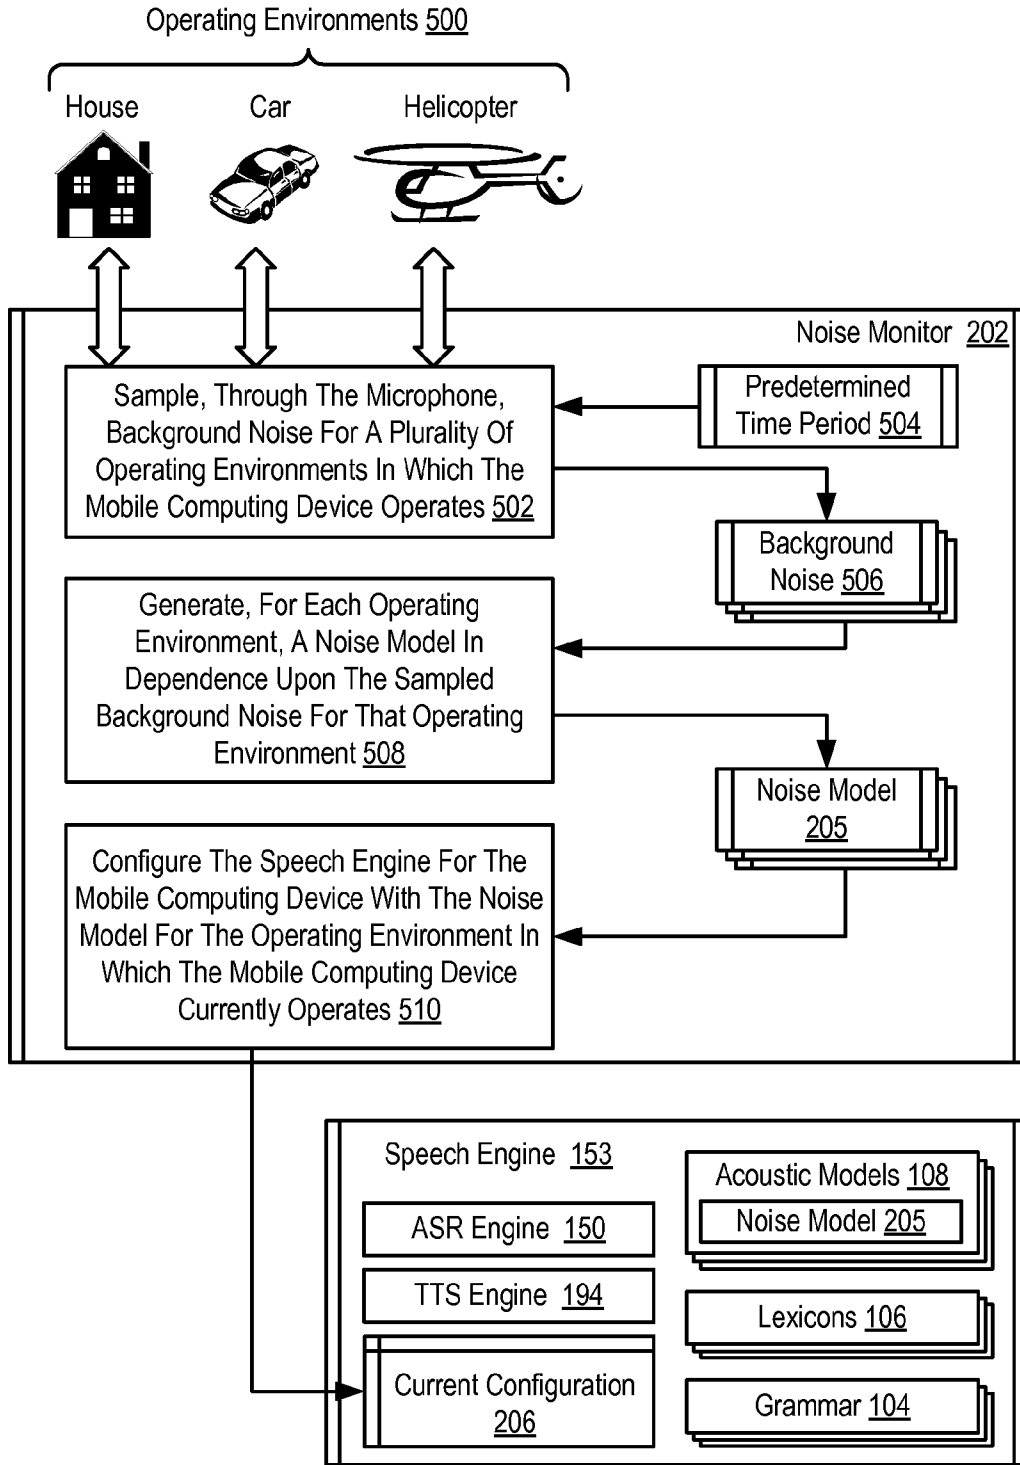


FIG. 5

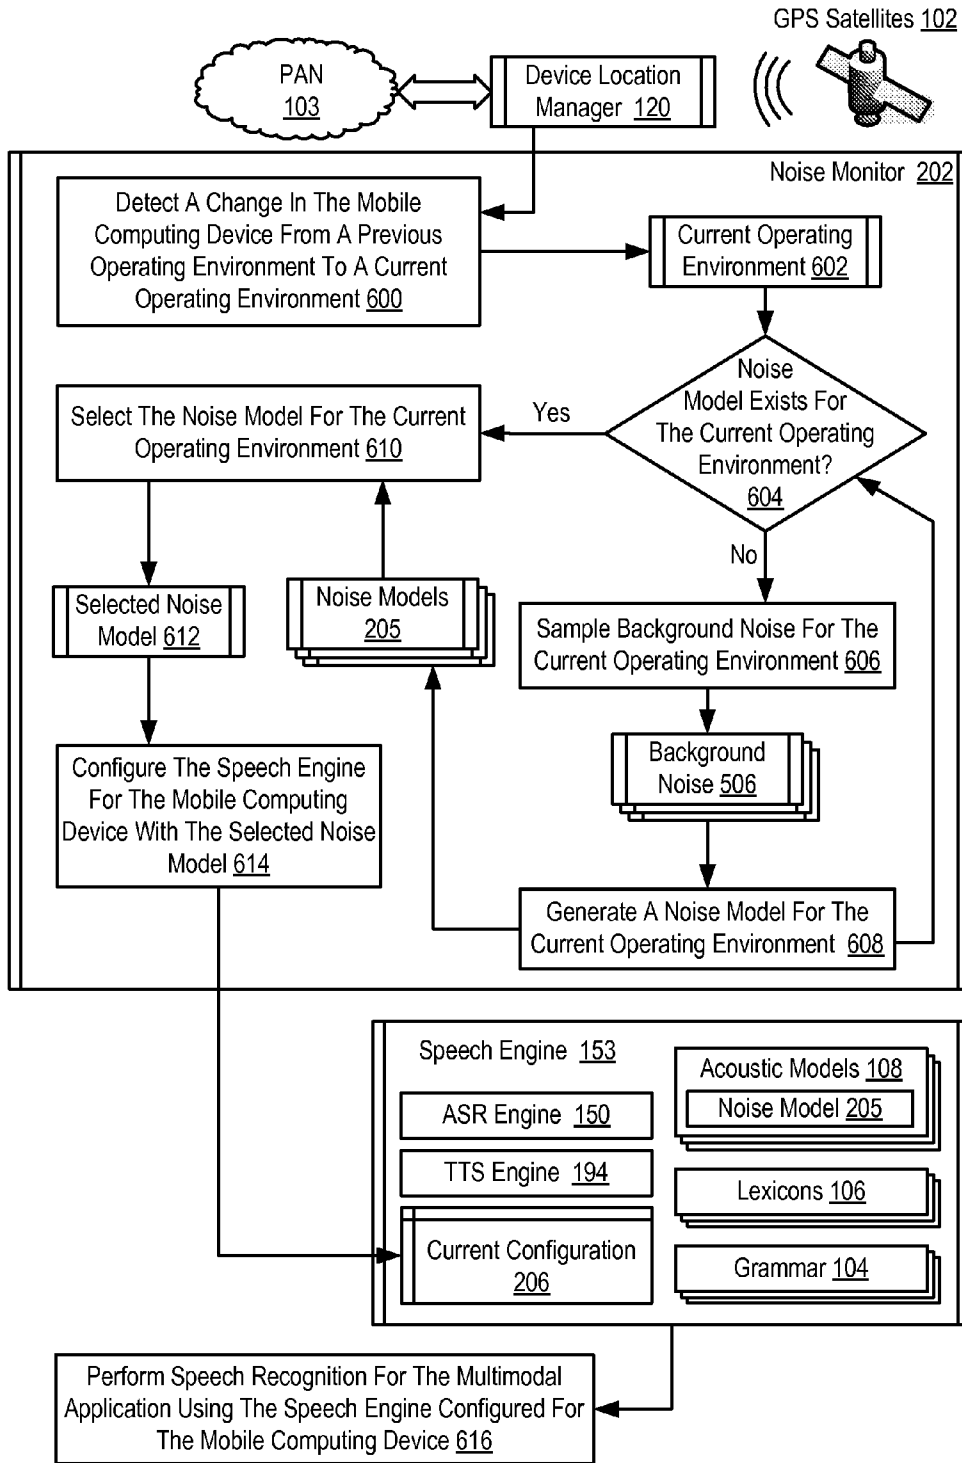


FIG. 6

**ADJUSTING A SPEECH ENGINE FOR A
MOBILE COMPUTING DEVICE BASED ON
BACKGROUND NOISE**

BACKGROUND OF THE INVENTION

[0001] 1. Field of the Invention

[0002] The field of the invention is data processing, or, more specifically, methods, apparatus, and products for adjusting a speech engine for a mobile computing device based on background noise.

[0003] 2. Description of Related Art

[0004] User interaction with applications running on small devices through a keyboard or stylus has become increasingly limited and cumbersome as those devices have become increasingly smaller. In particular, small handheld devices like mobile phones and PDAs serve many functions and contain sufficient processing power to support user interaction through multimodal access, that is, by interaction in non-voice modes as well as voice mode. Devices which support multimodal access combine multiple user input modes or channels in the same interaction allowing a user to interact with the applications on the device simultaneously through multiple input modes or channels. The methods of input include speech recognition, keyboard, touch screen, stylus, mouse, handwriting, and others. Multimodal input often makes using a small device easier.

[0005] Multimodal applications are often formed by sets of markup documents served up by web servers for display on multimodal browsers. A 'multimodal browser,' as the term is used in this specification, generally means a web browser capable of receiving multimodal input and interacting with users with multimodal output, where modes of the multimodal input and output include at least a speech mode. Multimodal browsers typically render web pages written in XHTML+Voice ('X+V'). X+V provides a markup language that enables users to interact with a multimodal application often running on a server through spoken dialog in addition to traditional means of input such as keyboard strokes and mouse pointer action. Visual markup tells a multimodal browser what the user interface is look like and how it is to behave when the user types, points, or clicks. Similarly, voice markup tells a multimodal browser what to do when the user speaks to it. For visual markup, the multimodal browser uses a graphics engine; for voice markup, the multimodal browser uses a speech engine. X+V adds spoken interaction to standard web content by integrating XHTML (eXtensible Hypertext Markup Language) and speech recognition vocabularies supported by VoiceXML. For visual markup, X+V includes the XHTML standard. For voice markup, X+V includes a subset of VoiceXML. For synchronizing the VoiceXML elements with corresponding visual interface elements, X+V uses events. XHTML includes voice modules that support speech synthesis, speech dialogs, command and control, and speech grammars. Voice handlers can be attached to XHTML elements and respond to specific events. Voice interaction features are integrated with XHTML and can consequently be used directly within XHTML content.

[0006] In addition to X+V, multimodal applications also may be implemented with Speech Application Tags ('SALT'). SALT is a markup language developed by the Salt Forum. Both X+V and SALT are markup languages for creating applications that use voice input/speech recognition and voice output/speech synthesis. Both SALT applications and X+V applications use underlying speech recognition and syn-

thesis technologies or 'speech engines' to do the work of recognizing and generating human speech. As markup languages, both X+V and SALT provide markup-based programming environments for using speech engines in an application's user interface. Both languages have language elements, markup tags, that specify what the speech-recognition engine should listen for and what the synthesis engine should 'say.' Whereas X+V combines XHTML, VoiceXML, and the XML Events standard to create multimodal applications, SALT does not provide a standard visual markup language or eventing model. Rather, it is a low-level set of tags for specifying voice interaction that can be embedded into other environments. In addition to X+V and SALT, multimodal applications may be implemented in Java with a Java speech framework, in C++, for example, and with other technologies and in other environments as well.

[0007] Current multimodal applications support a voice mode of user interaction using a speech engine. A speech engine provides recognition and generation or 'synthesis' of human speech though use of an acoustic model that associates speech waveform data representing recorded pronunciations of speech with textual representations of those pronunciations, also referred to as 'phonemes.' Speech recognition typically works best when the speech engine uses an acoustic model for speech recognition in the same environment in which the acoustic model was generated because the background noise embedded in the acoustic model and the speech for recognition are the same. Utilizing an acoustic model for speech recognition in an environment different from the environment in which the acoustic model was generated generally degrades overall speech recognition reliability and accuracy. Because multimodal devices are often tasked with operating in a variety of environments, each having a different background noise, the speech recognition performance of current multimodal devices varies greatly depending on the location and the surroundings of the device.

SUMMARY OF THE INVENTION

[0008] Methods, apparatus, and products are disclosed for adjusting a speech engine for a mobile computing device based on background noise, the mobile computing device operatively coupled to a microphone, that include: sampling, through the microphone, background noise for a plurality of operating environments in which the mobile computing device operates; generating, for each operating environment, a noise model in dependence upon the sampled background noise for that operating environment; and configuring the speech engine for the mobile computing device with the noise model for the operating environment in which the mobile computing device currently operates.

[0009] The foregoing and other objects, features and advantages of the invention will be apparent from the following more particular descriptions of exemplary embodiments of the invention as illustrated in the accompanying drawings wherein like reference numbers generally represent like parts of exemplary embodiments of the invention.

BRIEF DESCRIPTION OF THE DRAWINGS

[0010] FIG. 1 sets forth a network diagram illustrating an exemplary system for adjusting a speech engine for a mobile computing device based on background noise according to embodiments of the present invention.

[0011] FIG. 2 sets forth a block diagram of automated computing machinery comprising an example of a computer useful as a voice server in adjusting a speech engine for a mobile computing device based on background noise according to embodiments of the present invention.

[0012] FIG. 3 sets forth a functional block diagram of exemplary apparatus for adjusting a speech engine for a mobile computing device based on background noise according to embodiments of the present invention.

[0013] FIG. 4 sets forth a block diagram of automated computing machinery comprising an example of a computer useful as a multimodal device in adjusting a speech engine for a mobile computing device based on background noise according to embodiments of the present invention.

[0014] FIG. 5 sets forth a flow chart illustrating an exemplary method of adjusting a speech engine for a mobile computing device based on background noise according to embodiments of the present invention.

[0015] FIG. 6 sets forth a flow chart illustrating a further exemplary method of adjusting a speech engine for a mobile computing device based on background noise according to embodiments of the present invention.

DETAILED DESCRIPTION OF EXEMPLARY EMBODIMENTS

[0016] Exemplary methods, apparatus, and products for adjusting a speech engine for a mobile computing device based on background noise according to embodiments of the present invention are described with reference to the accompanying drawings, beginning with FIG. 1. FIG. 1 sets forth a network diagram illustrating an exemplary system for adjusting a speech engine (153) for a mobile computing device (152) based on background noise according to embodiments of the present invention. A mobile computing device is a computer capable of operating without being connected to traditional, immobile support infrastructures such as, for example, power infrastructures, networking infrastructures, and so on, while being transported from one location to another. Mobile computing devices are typically characterized by having their own built in batteries to supply power needed to operate the device's electronic components. Examples of mobile computing devices may include personal digital assistants, smart phones, laptops, computers embedded in vehicles, and so on.

[0017] The mobile computing device (152) throughout this specification is implemented as multimodal device having installed upon it a multimodal application (195). Readers will note however that implementing a mobile computing device as a multimodal device is for explanation only and not for limitation. The multimodal device (152) supports multiple modes of user interaction with the multimodal application (195) including a voice mode and one or more non-voice modes of user interaction with the multimodal application (195). The voice mode is represented here with audio output of voice prompts and responses (314) from the multimodal devices (152) and audio input of speech for recognition (315) from a user (128). Non-voice modes are represented by input/output devices such as keyboards and display screens on the multimodal devices (152). The voice mode of user interaction with the multimodal application is supported by a voice interpreter (192), through which the multimodal application (195) is operatively coupled to a speech engine (153). The operative coupling may be implemented with an application program-

ming interface ('API'), a voice service module, or a VOIP connection as explained in more detail below.

[0018] In the exemplary system of FIG. 1, the voice interpreter (192) supports the voice mode of user interaction with the multimodal application (195) by providing grammars, speech for recognition, and text prompts for speech synthesis to the speech engine (153), and by returning to the multimodal application (195) speech engine output in the form of recognized speech, semantic interpretation results, and digitized speech for voice prompts. The implementation of the voice interpreter (192) typically depends on the technology implementing the multimodal application (195). The voice interpreter (192) of FIG. 1 may be implemented as a VoiceXML interpreter when the multimodal application (195) implements X+V. A VoiceXML interpreter is a software module of computer program instructions that accepts voice dialog instructions from a multimodal application, typically in the form of a VoiceXML <form> element. The voice dialog instructions include one or more grammars, data input elements, event handlers, and so on, that advise the VoiceXML interpreter how to administer voice input from a user and voice prompts and responses to be presented to a user. The VoiceXML interpreter administers such dialogs by processing the dialog instructions sequentially in accordance with a VoiceXML Form Interpretation Algorithm ('FIA'). When the multimodal application (195) is implemented according to SALT, the voice interpreter (192) may be implemented as a SALT interpreter. When the multimodal application (195) is implemented using Java, the voice interpreter (192) may be implemented using a VoiceXML interpreter that exposes a Java interface.

[0019] In the example of FIG. 1, the mobile computing device (152) has installed upon a noise monitor (202), a software component that operates to adjust a speech engine (153) for a mobile computing device based on background noise according to embodiments of the present invention. The noise monitor (202) of FIG. 1 operates generally for adjusting the speech engine (153) for a mobile computing device based on background noise according to embodiments of the present invention by: sampling, through the microphone, background noise for a plurality of operating environments in which the mobile computing device (152) operates; generating, for each operating environment, a noise model in dependence upon the sampled background noise for that operating environment; and configuring the speech engine (153) for the mobile computing device (152) with the noise model for the operating environment in which the mobile computing device (152) currently operates. In such a manner, the noise monitor (202) may periodically tune the speech engine (153) enhance speech recognition performance based on the mobile computing device's current operating environment.

[0020] In some other embodiments, the noise monitor (202) may operate for adjusting a speech engine (153) for a mobile computing device based on background noise according to embodiments of the present invention by: detecting a change in the mobile computing device (152) from a previous operating environment to a current operating environment; determining whether a noise model exists for the current operating environment; sampling background noise for the current operating environment if a noise model does not exist for the current operating environment; generating a noise model for the current operating environment if a noise model does not exist for the current operating environment; and configuring the speech engine (153) for the mobile comput-

ing device (152) with the generated noise model. If a noise model does exist for the current operating environment, the noise monitor (202) may operate to adjust a speech engine (153) for a mobile computing device based on background noise according to embodiments of the present invention by: selecting the noise model for the current operating environment and configuring the speech engine (153) for the mobile computing device (152) with the selected noise model.

[0021] As illustrated in FIG. 1, the noise monitor (202) may be installed locally on the mobile computing device (152) or installed remotely with respect to the mobile computing device (152) on some other computing device such as, a voice server (151). When the noise monitor (202) is installed remotely on some other computing device, a noise monitor client may be installed on the mobile computing device (152) to provide an interface with the hardware and software components of the mobile device (152). Such a noise monitor client may communicate with the noise monitor through a data communications connection established through the data communications network (100).

[0022] In the example of FIG. 1, the noise monitor (202) may detect a change in the mobile computing device from a previous operating environment to a current operating environment by receiving a notification message from a device location manager (120). The device location manager (120) of FIG. 1 is a software module that includes computer program instructions for monitoring the location of the mobile computing device (152). The location of the mobile computing device (152) may be specified using the geographic coordinates such as, for example, (26°38'7.27"N, 80°13'37.22"W), using a network address or identifier from a proximate, adhoc personal area network such as, for example, a Bluetooth™ network, or in any other manner as will occur to those of skill in the art. The location of the mobile computing device (152), however, may also be specified as a semantic representation of the coordinates such as, for example, 'home,' 'work,' 'church,' or 'car.'

[0023] The device location manager (120) of FIG. 1 obtains location information from a position detection component (not shown) of the multimodal device (152) that is operatively coupled to the device location manager (120). A position detection component is a component, typically implemented using a combination of hardware and software, that communicates with other devices to determine the mobile computing device's position information. For example, the position detection component may be implemented as a global positioning system ('GPS') receiver that calculates the geographic coordinates of the device based on the device's position relative to a group of GPS satellites (102). The use of GPS, however, is for explanation and not for limitation. In other embodiments of the present invention, the position detection component may be implemented using a Long Range Navigation ('LORAN') receiver that calculates the geographic coordinates of the device based on the device's position relative to a group of LORAN radio transmission towers or any other component for calculating the geographic coordinate of a multimodal device as will occur to those of skill in the art. In still other embodiments, the position detection component may determine the geographic coordinates of the multimodal device using the signals from cell phone towers to triangulate the position of the multimodal device. In other examples, the position detection component may be implemented as a Bluetooth™ adapter or IEEE 802.11 adapter that determines when the mobile computing device

(152) establishes a connections with a personal area network ('PAN') such as, for example, a PAN in an automobile, train, airplane, portion of a building, or a Local Area Network ('LAN') such as for example, a home network, building network, office network, and so on.

[0024] In the example of FIG. 1, the device location manager (120) is 'operatively coupled' to a position detection component in the sense that the device location manager (120) communicates with a position detection module (121), a software driver that provides interaction with the position detection component through an exposed API. Such communication may occur locally on the multimodal device (152) when the device location manager (120) is installed on the multimodal device (152) as illustrated in FIG. 1 or may occur across a network when the device location manager (120) is installed on a server, which manages the locations of multiple multimodal devices (152).

[0025] As mentioned above, the mobile computing devices (152) of FIG. 1 are implemented as multimodal devices. A multimodal device is an automated device, that is, automated computing machinery or a computer program running on an automated device, that is capable of accepting from users more than one mode of input, keyboard, mouse, stylus, and so on, including speech input—and also rendering more than one mode of output, graphic, speech, and so on. A multimodal device is generally capable of accepting speech input from a user, digitizing the speech, and providing digitized speech to a speech engine for recognition. A multimodal device may be implemented, for example, as a voice-enabled browser on a laptop, a voice browser on a telephone handset, an online game implemented with Java on a personal computer, and with other combinations of hardware and software as may occur to those of skill in the art. Because multimodal applications may be implemented in markup languages (X+V, SALT), object-oriented languages (Java, C++), procedural languages (the C programming language), and in other kinds of computer languages as may occur to those of skill in the art, this specification uses the term 'multimodal application' to refer to any software application, server-oriented or client-oriented, thin client or thick client, that administers more than one mode of input and more than one mode of output, typically including visual and speech modes.

[0026] The system of FIG. 1 includes several example multimodal devices:

[0027] personal digital assistant ('PDA') (112) which is coupled for data communications to data communications network (100) through wireless connection (114),

[0028] mobile telephone (110) which is coupled for data communications to data communications network (100) through wireless connection (116), and

[0029] laptop computer (126) which is coupled for data communications to data communications network (100) through wireless connection (118).

[0030] Each of the example multimodal devices (152) in the system of FIG. 1 includes a microphone, an audio amplifier, a digital-to-analog converter, and a multimodal application capable of accepting from a user (128) speech for recognition (315), digitizing the speech, and providing the digitized speech to a speech engine for recognition. The speech may be digitized according to industry standard codecs, including but not limited to those used for Distributed Speech Recognition as such. Methods for 'CODing/DECoding' speech are referred to as 'codecs.' The European Telecommunications Standards Institute ('ETSI') provides sev-

eral codecs for encoding speech for use in DSR, including, for example, the ETSI ES 201 108 DSR Front-end Codec, the ETSI ES 202 050 Advanced DSR Front-end Codec, the ETSI ES 202 211 Extended DSR Front-end Codec, and the ETSI ES 202 212 Extended Advanced DSR Front-end Codec. In standards such as RFC3557 entitled

[0031] RTP Payload Format for European Telecommunications Standards Institute (ETSI) European Standard ES 201 108 Distributed Speech Recognition Encoding and the Internet Draft entitled

[0032] RTP Payload Formats for European Telecommunications Standards Institute (ETSI) European Standard ES 202 050, ES 202 211, and ES 202 212 Distributed Speech Recognition Encoding,

the IETF provides standard RTP payload formats for various codecs. It is useful to note, therefore, that there is no limitation in the present invention regarding codecs, payload formats, or packet structures. Speech for adjusting a speech engine for a mobile computing device based on background noise according to embodiments of the present invention may be encoded with any codec, including, for example:

[0033] AMR (Adaptive Multi-Rate Speech coder)

[0034] ARDOR (Adaptive Rate-Distortion Optimized sound codeR),

[0035] Dolby Digital (A/52, AC3),

[0036] DTS (DTS Coherent Acoustics),

[0037] MP1 (MPEG audio layer-1),

[0038] MP2 (MPEG audio layer-2) Layer 2 audio codec (MPEG-1, MPEG-2 and non-ISO MPEG-2.5),

[0039] MP3 (MPEG audio layer-3) Layer 3 audio codec (MPEG-1, MPEG-2 and non-ISO MPEG-2.5),

[0040] Perceptual Audio Coding,

[0041] FS-1015 (LPC-10),

[0042] FS-1016 (CELP),

[0043] G.726 (ADPCM),

[0044] G.728 (LD-CELP),

[0045] G.729 (CS-ACELP),

[0046] GSM,

[0047] HILN (MPEG-4 Parametric audio coding), and

[0048] others as may occur to those of skill in the art.

[0049] As mentioned, a multimodal device according to embodiments of the present invention is capable of providing speech for recognition to a speech engine (153) adjusted for the mobile computing device (152) based on background noise according to embodiments of the present invention. A speech engine is a functional module, typically a software module, although it may include specialized hardware also, that does the work of recognizing and generating or 'synthesizing' human speech. The speech engine implements speech recognition by use of a further module referred to in this specification as a ASR engine, and the speech engine carries out speech synthesis by use of a further module referred to in this specification as a text-to-speech ('TTS') engine. As shown in FIG. 1, a speech engine (153) may be installed locally in the mobile computing device (152) itself, or a speech engine (153) may be installed remotely with respect to the mobile computing device, across a data communications network (100) in a voice server (151). A mobile computing device that itself contains its own speech engine is said to implement a 'thick client' or 'thick client,' because the thick client itself contains all the functionality needed to carry out speech recognition and speech synthesis—through API calls to speech recognition and speech synthesis modules in the mobile computing device itself with no need to send requests

for speech recognition across a network and no need to receive synthesized speech across a network from a remote voice server. A mobile computing device that does not contain its own speech engine is said to implement a 'thin client' because the thin client itself contains only a relatively thin layer of application software that obtains speech recognition and speech synthesis services from a voice server located remotely across a network from the thin client. For ease of explanation, only one (112) of the mobile computing devices (152) in the system of FIG. 1 is shown with a speech engine (153), but readers will recognize that any multimodal device may have a speech engine according to embodiments of the present invention.

[0050] As shown in FIG. 1, a speech engine (153) and a voice interpreter (192) may be installed locally in the multimodal device (112) itself, or a speech engine (153) and a voice interpreter (192) may be installed remotely with respect to the multimodal device, across a data communications network (100) in a voice server (151). In a thick client architecture, a multimodal device (152) includes both its own speech engine (153) and its own voice interpreter (192). The voice interpreter (192) exposes an API to the multimodal application (195) for use in providing speech recognition and speech synthesis for the multimodal application. The multimodal application provides, for example, dialog instructions, VoiceXML <form> elements, grammars, input elements, event handlers, and so on, through the API to the voice interpreter (192), and the voice interpreter (192) administers the speech engine on behalf of the multimodal application. In the thick client architecture, a multimodal application, including for example, VoiceXML dialogs, is interpreted by a voice interpreter on the multimodal device. In the thin client architecture, a multimodal application, including for example VoiceXML dialogs, is interpreted by a voice interpreter on a voice server (151) located remotely across a data communications network (100) from the multimodal device running the multimodal application (195).

[0051] In a thin client architecture, the speech engine (153) and the voice interpreter (192) are located remotely from the multimodal client device in a voice server (151), the API for the voice interpreter is still implemented in the multimodal device, with the API modified to communicate voice dialog instructions, speech for recognition, and text and voice prompts to and from the voice interpreter on the voice server. For ease of explanation, only one (112) of the multimodal devices (152) in the system of FIG. 1 is shown with a voice interpreter (192), but readers will recognize that any multimodal device may have a voice interpreter according to embodiments of the present invention.

[0052] The use of these three example multimodal devices (152) is for explanation only, not for limitation of the invention. Any automated computing machinery capable of accepting speech from a user, providing the speech digitized to a speech engine through a voice interpreter, and receiving and playing speech prompts and responses from the voice interpreter may be improved to function as a multimodal device for adjusting a speech engine based on background noise according to embodiments of the present invention.

[0053] The system of FIG. 1 also includes a voice server (151) which is connected to data communications network (100) through wireline connection (122). The voice server (151) is a computer that runs a speech engine (153) adjusted based on background noise according to embodiments of the present invention. The voice server (151) provides voice rec-

ognition services for multimodal devices by accepting requests for speech recognition and returning text representing recognized speech. Voice server (151) also provides speech synthesis, text to speech ('TTS') conversion, for voice prompts and voice responses (314) to user input in multimodal applications such as, for example, X+V applications, SALT applications, or Java voice applications.

[0054] The system of FIG. 1 includes a data communications network (100) that connects the multimodal devices (152) and the voice server (151) for data communications. A data communications network for adjusting a speech engine for a mobile computing device based on background noise according to embodiments of the present invention is a data communications network composed of a plurality of computers that function as data communications routers connected for data communications with packet switching protocols. Such a data communications network may be implemented with optical connections, wireline connections, or with wireless connections. Such a data communications network may include intranets, internets, local area data communications networks ('LANs'), and wide area data communications networks ('WANs'). Such a data communications network may implement, for example:

[0055] a link layer with the Ethernet™ Protocol or the Wireless Ethernet™ Protocol,

[0056] a data communications network layer with the Internet Protocol ('IP'),

[0057] a transport layer with the Transmission Control Protocol ('TCP') or the User Datagram Protocol ('UDP'),

[0058] an application layer with the HyperText Transfer Protocol ('HTTP'), the Session Initiation Protocol ('SIP'), the Real Time Protocol ('RTP'), the Distributed Multimodal Synchronization Protocol ('DMSP'), the Wireless Access Protocol ('WAP'), the Handheld Device Transfer Protocol ('HDTP'), the ITU protocol known as H.323, and

[0059] other protocols as will occur to those of skill in the art.

[0060] The system of FIG. 1 includes a web server (147) connected for data communications through wireline connection (123) to network (100) and therefore to the multimodal devices (152). The web server (147) may be any server that provides to client devices markup documents that compose multimodal applications. The web server (147) typically provides such markup documents via a data communications protocol, HTTP, HDTP, WAP, or the like. That is, although the term 'web' is used to describe the web server generally in this specification, there is no limitation of data communications between multimodal devices and the web server to HTTP alone. The markup documents also may be implemented in any markup language that supports non-speech display elements, data entry elements, and speech elements for identifying which speech to recognize and which words to speak, grammars, form elements, and the like, including, for example, X+V and SALT. A multimodal application in a multimodal device then, upon receiving from the web sever (147) a markup document as part of a multimodal application, may execute speech elements by use of a voice interpreter (192) and speech engine (153) in the multimodal device itself or by use of a voice interpreter (192) and speech engine (153) located remotely from the multimodal device in a voice server (151).

[0061] The arrangement of the multimodal devices (152), the web server (147), the voice server (151), and the data communications network (100) making up the exemplary system illustrated in FIG. 1 are for explanation, not for limitation. Data processing systems useful for adjusting a speech engine for a mobile computing device based on background noise according to various embodiments of the present invention may include additional servers, routers, other devices, and peer-to-peer architectures, not shown in FIG. 1, as will occur to those of skill in the art. Data communications networks in such data processing systems may support many data communications protocols in addition to those noted above. Various embodiments of the present invention may be implemented on a variety of hardware platforms in addition to those illustrated in FIG. 1.

[0062] Adjusting a speech engine for a mobile computing device based on background noise according to embodiments of the present invention in a thin client architecture may be implemented with one or more voice servers, computers, that is, automated computing machinery, that provide speech recognition and speech synthesis. For further explanation, therefore, FIG. 2 sets forth a block diagram of automated computing machinery comprising an example of a computer useful as a voice server (151) in adjusting a speech engine for a mobile computing device based on background noise according to embodiments of the present invention. The voice server (151) of FIG. 2 includes at least one computer processor (156) or 'CPU' as well as random access memory (168) ('RAM') which is connected through a high speed memory bus (166) and bus adapter (158) to processor (156) and to other components of the voice server.

[0063] Stored in RAM (168) is a voice server application (188), a module of computer program instructions capable of operating a voice server in a system that is configured to carry out adjusting a speech engine for a mobile computing device based on background noise according to embodiments of the present invention. Voice server application (188) provides voice recognition services for multimodal devices by accepting requests for speech recognition and returning speech recognition results, including text representing recognized speech, text for use as variable values in dialogs, and text as string representations of scripts for semantic interpretation. Voice server application (188) also includes computer program instructions that provide text-to-speech ('TTS') conversion for voice prompts and voice responses to user input in multimodal applications such as, for example, X+V applications, SALT applications, or Java Speech applications.

[0064] Voice server application (188) may be implemented as a web server, implemented in Java, C++, or another language, that supports X+V, SALT, VoiceXML, or other multimodal languages, by providing responses to HTTP requests from X+V clients, SALT clients, Java Speech clients, or other multimodal clients. Voice server application (188) may, for a further example, be implemented as a Java server that runs on a Java Virtual Machine and supports a Java voice framework by providing responses to HTTP requests from Java client applications running on multimodal devices. And voice server applications that support adjusting a speech engine for a mobile computing device based on background noise may be implemented in other ways as may occur to those of skill in the art, and all such ways are well within the scope of the present invention.

[0065] Also stored in RAM is a voice interpreter (192), a module of computer program instructions that supports the

voice mode of user interaction with a multimodal application operating on a multimodal device. The voice interpreter (192) provides speech engine input such as grammars, speech for recognition, and text prompts for speech synthesis to the speech engine (153) and returns to the multimodal application speech engine output in the form of recognized speech, semantic interpretation results, and digitized speech for voice prompts. Input to voice interpreter (192) may originate, for example, from VoiceXML clients running remotely on multimodal devices, from X+V clients running remotely on multimodal devices, from SALT clients running on multimodal devices, or from Java client applications running remotely on multimodal devices. In this example, voice interpreter (192) interprets and executes VoiceXML segments representing voice dialog instructions received from remote multimodal devices and provided to voice interpreter (192) through voice server application (188).

[0066] When implemented in X+V, a multimodal application in a thin client architecture may provide voice dialog instructions, VoiceXML segments, VoiceXML <form> elements, and the like, to voice interpreter (149) through data communications across a network with the multimodal application. The voice dialog instructions include one or more grammars, data input elements, event handlers, and so on, that advise the voice interpreter how to administer voice input from a user and voice prompts and responses to be presented to a user. The voice interpreter (192) administers such dialogs by processing the dialog instructions sequentially in accordance with a VoiceXML Form Interpretation Algorithm ('FIA'). The voice interpreter (192) interprets VoiceXML dialogs provided to the voice interpreter (192) by a multimodal application.

[0067] To provide voice services to a multimodal application, the voice server (151) in this example includes a speech engine (153). The speech engine is a functional module, typically a software module, although it may include specialized hardware also, that does the work of recognizing and generating human speech. The speech engine (153) includes an automated speech recognition ('ASR') engine for speech recognition and a text-to-speech ('TTS') engine for generating speech. The speech engine also includes grammars (104), lexicons (106), and language-specific acoustic models (108). Each lexicon (106) and language-specific acoustic model (108) may be associated with a separate language in a configuration parameter repository (200).

[0068] The acoustic models (108) associate speech waveform data representing recorded pronunciations of speech with textual representations of those pronunciations, which are referred to as 'phonemes.' The speech waveform data may be implemented as a Speech Feature Vector ('SFV') that may be represented, for example, by the first twelve or thirteen Fourier or frequency domain components of a sample of digitized speech waveform. Accordingly, the acoustic models (108) may be implemented as data structures or tables in a database, for example, that associates these SFVs with phonemes representing, to the extent that it is practically feasible to do so, all pronunciations of all the words in various human languages, each language having a separate acoustic model (108). The lexicons (106) are associations of words in text form with phonemes representing pronunciations of each word; the lexicon effectively identifies words that are capable of recognition by an ASR engine. Each language has a separate lexicon (106). Also stored in RAM (168) is a Text To Speech ('TTS') Engine (194), a module of computer program

instructions that accepts text as input and returns the same text in the form of digitally encoded speech, for use in providing speech as prompts for and responses to users of multimodal systems.

[0069] The grammars (104) communicate to the ASR engine (150) the words and sequences of words that currently may be recognized. For precise understanding, distinguish the purpose of the grammar and the purpose of the lexicon. The lexicon associates with phonemes all the words that the ASR engine can recognize. The grammar communicates the words currently eligible for recognition. The set of words currently eligible for recognition and the set of words capable of recognition may or may not be the same.

[0070] Grammars for use in adjusting a speech engine for a mobile computing device based on background noise according to embodiments of the present invention may be expressed in any format supported by any ASR engine, including, for example, the Java Speech Grammar Format ('JSGF'), the format of the W3C Speech Recognition Grammar Specification ('SRGS'), the Augmented Backus-Naur Format ('ABNF') from the IETF's RFC2234, in the form of a stochastic grammar as described in the W3C's Stochastic Language Models (N-Gram) Specification, and in other grammar formats as may occur to those of skill in the art. Grammars typically operate as elements of dialogs, such as, for example, a VoiceXML <menu> or an X+V<form>. A grammar's definition may be expressed in-line in a dialog. Or the grammar may be implemented externally in a separate grammar document and referenced from with a dialog with a URI. Here is an example of a grammar expressed in JSGF:

```
<grammar scope="dialog" >![CDATA[
#JSGF V1.0;
grammar command;
<command> = [remind me to] call | phone | telephone <name>
<when>;
<name> = bob | martha | joe | pete | chris | john | artoush | tom;
<when> = today | this afternoon | tomorrow | next week;
]]>
</grammar>
```

[0071] In this example, the elements named <command>, <name>, and <when> are rules of the grammar. Rules are a combination of a rulename and an expansion of a rule that advises an ASR engine or a voice interpreter which words presently can be recognized. In this example, expansion includes conjunction and disjunction, and the vertical bars '|' mean 'or.' An ASR engine or a voice interpreter processes the rules in sequence, first <command>, then <name>, then <when>. The <command> rule accepts for recognition 'call' or 'phone' or 'telephone' plus, that is, in conjunction with, whatever is returned from the <name> rule and the <when> rule. The <name> rule accepts 'bob' or 'martha' or 'joe' or 'pete' or 'chris' or 'john' or 'artoush' or 'tom', and the <when> rule accepts 'today' or 'this afternoon' or 'tomorrow' or 'next week.' The command grammar as a whole matches utterances like these, for example:

- [0072] "phone bob next week,"
- [0073] "telephone martha this afternoon,"
- [0074] "remind me to call chris tomorrow," and
- [0075] "remind me to phone pete today."

[0076] The voice server application (188) in this example is configured to receive, from a multimodal client located remotely across a network from the voice server, digitized

speech for recognition from a user and pass the speech along to the ASR engine (150) for recognition. ASR engine (150) is a module of computer program instructions, also stored in RAM in this example. In carrying out automated speech recognition, the ASR engine receives speech for recognition in the form of at least one digitized word and uses frequency components of the digitized word to derive a SFV for the digitized word. The ASR engine can use the SFV to infer phonemes for the word from the language-specific acoustic model (108). The ASR engine then uses the phonemes to find the word in the lexicon (106).

[0077] In the example of FIG. 2, the speech engine (153) operates according to the current configuration (206) for the speech engine (153). The current configuration (206) may specify the active lexicon and the active acoustic model and noise model used by the speech engine (153) to recognize and synthesize human speech. The current configuration (206) may also specify a current configuration for the automatic speech recognition ('ASR') engine such as, for example, speech transition times, silence detection times, speech timeouts, and gain maps. Furthermore, the current configuration (206) may also specify a configuration for use by a text-to-speech ('TTS') engine such as, for example, the voice used in synthesizing speech from text.

[0078] In the example of FIG. 2, each acoustic model (108) includes a noise model (205) that represents the waveform of the recorded background noise in an environment, often the environment in which the speech waveform for the acoustic model (108) was recorded. The noise model (205) may also be implemented as the first twelve or thirteen Fourier or frequency domain components of a sample of the background noise waveform. When the speech engine (153) compares Fourier or frequency domain representations of a voice utterance with the SFVs in the acoustic models (108) to identify phonemes in the acoustic model (108), the speech engine (153) may subtract out the background noise specified by the noise model (205) from the Fourier or frequency domain representations of a voice utterance. In such a manner, speech recognition performance is generally increased when the background noise represented by the noise model (205) matches the background noise in the environment in which the voice utterance for recognition is recorded.

[0079] Stored in RAM (168) of FIG. 2 is a noise monitor (202), a software component that operates to adjust a speech engine (153) for a mobile computing device based on background noise according to embodiments of the present invention. The noise monitor (202) of FIG. 2 operates generally for adjusting the speech engine (153) for a mobile computing device based on background noise according to embodiments of the present invention by: sampling, through a microphone of the mobile computing device, background noise for a plurality of operating environments in which the mobile computing device operates; generating, for each operating environment, a noise model (205) in dependence upon the sampled background noise for that operating environment; and configuring the speech engine (153) for the mobile computing device with the noise model (205) for the operating environment in which the mobile computing device currently operates. In such a manner, the noise monitor (202) may periodically tune the speech engine (153) enhance speech recognition performance based on the mobile computing device's current operating environment. The noise monitor (202) may configure the speech engine (153) for the mobile computing device with the noise model (205) for the operat-

ing environment in which the mobile computing device currently operates by: updating a current acoustic model (108) for the speech engine (153) with the noise model (205) for the operating environment in which the mobile computing device currently operates or replacing the current acoustic model (108) for the speech engine (153) with another acoustic model (108) already incorporating the noise model (205) for the operating environment in which the mobile computing device currently operates.

[0080] In some other embodiments, the noise monitor (202) of FIG. 2 may operate for adjusting a speech engine (153) for a mobile computing device based on background noise according to embodiments of the present invention by: detecting a change in the mobile computing device from a previous operating environment to a current operating environment; determining whether a noise model (205) exists for the current operating environment; sampling background noise for the current operating environment if a noise model (205) does not exist for the current operating environment; generating a noise model (205) for the current operating environment if a noise model (205) does not exist for the current operating environment; and configuring the speech engine (153) for the mobile computing device with the generated noise model (205). If a noise model (205) does exist for the current operating environment, the noise monitor (202) may operate to adjust a speech engine (153) for a mobile computing device based on background noise according to embodiments of the present invention by: selecting the noise model (205) for the current operating environment and configuring the speech engine (153) for the mobile computing device with the selected noise model (205).

[0081] Also stored in RAM (168) of FIG. 2 is a device location manager (120), a software module that includes computer program instructions for monitoring the locations of the mobile computing devices based on location information provided by position detection components installed in the mobile computing devices. The location of a mobile computing device may be specified using the geographic coordinates such as, for example, (26°38'7.27"N, 80°13'37.22"W), using a network address or network identifier from a proximate, adhoc personal area network such as, for example, a Bluetooth™ network, or in any other manner as will occur to those of skill in the art. The locations of the mobile computing devices, however, may also be specified as a semantic representation of the coordinates such as, for example, 'home,' 'work,' 'church,' or 'car.' To implement the location of a multimodal device using a semantic representation, the device location manager may access a device location repository (201) that associates location semantics with a set of geographic coordinates. For further explanation, consider the exemplary device location repository, network address, network identifier, or the like, illustrated below using XML schema,

```
<device location repository>
  <deviceid = "1">
    <location semantic = "church"
      centercoordinate = "(26.6351245°N, 80.2261248°W)"
      radius = "200 feet"/>
    <location semantic = "house"
      coordinate1 = "(26.0202777°N, 80.0400210°W)"
      coordinate2 = "(26.0204166°N, 80.0400210°W)"
```

```

-continued
-----
coordinate3 = "(26.0204166°N, 80.0398822°W)"
coordinate4 = "(26.0202777°N, 80.0398822°W)"/>
<location semantic = "car"
networkDeviceConnection =
  CAR_ADAPTER_MAC/>
...
</device>
<device id = "2">
...
</device>
...
</device location repository>,
-----

```

which maintains location semantics for multiple multimodal devices that request voice services from voice server (151). Using the exemplary device location repository above, the device location manager (120) may determine that the current location of the multimodal is 'church' if the coordinates received from the position detection component of the device indicate that the device is within a circular region centered at the coordinate (26.6351245° N, 80.2261248° W) and having a radius of 200 feet. Using the exemplary device location repository above, the device location manager (120) may determine that the current location of the multimodal device is 'house' if the coordinates received from the position detection component of the device indicate that the device is within an enclosed region defined by the coordinates (26.0202777° N, 80.0400210° W), (26.0204166° N, 80.0400210° W), (26.0204166° N, 80.0398822° W), and (26.0202777° N, 80.0398822° W). Using the exemplary device location repository above, the device location manager (120) may determine that the current location of the multimodal device is 'car' if the mobile computing device connects to a PAN in the user's car through the car's network adapter having a particular MAC address. Readers will note that the values and the format for the exemplary device location repository above are for explanation and not for limitation.

[0082] Readers will also note that the device location repository (201) may contain location semantics for more than one multimodal device, as illustrated in the exemplary device location repository above. Although such an implementation is useful when the device location manager (120) manages the location of multiple multimodal devices, the device location repository (201) may contain only the location semantics relevant to a single multimodal device when the device location manager (120) is installed on the multimodal device itself.

[0083] In the example of FIG. 2, the device location manager (120) maintains the device location repository (201). The device location manager (120) may receive location definitions that associate a location semantic with a geographic coordinate from a user through a device location manager client operating on a multimodal device. Through a user interface provided by a device location manager client, a user may associate particular location semantics with geographic coordinates. For example, a user may manually enter coordinates into a user interface and associate those coordinates with a location semantic chosen by the user. The user may also instruct the device location manager client to capture the coordinates of the device's current position and associate those captured coordinates with a location semantic. The device location manager client then provides the data received from the user to the device location manager (120),

which in turn stores the location definition into the device location repository (201). In other embodiments, the device location manager (120) may retrieve coordinates and their associated location semantics from data storage (170) or some other repository connected to the network (100).

[0084] As mentioned above, the noise monitor (202) adjusts a speech engine for a mobile computing device based on background noise by configuring the speech engine with the noise model for the operating environment in which the mobile computing device currently operates. The noise monitor (202) of FIG. 2 may identify the noise model for the operating environment in which the mobile computing device currently operates using a noise model repository (200). The noise model repository (200) of FIG. 2 may be implemented as a data structure or table that associates various noise models (205) with different operating environments. For example, consider the following exemplary noise model repository:

TABLE 1

EXEMPLARY NOISE MODEL REPOSITORY	
NOISE MODEL ID	OPERATING ENVIRONMENT ID
Noise_Model_0	Home
Noise_Model_1	Office
Noise_Model_2	Car
...	...

[0085] Each record of the exemplary noise model repository above includes a noise model identifier and an operating environment identifier. The noise model identifier specifies the particular noise model for use in the associated operating environment. The operating environment identifier specifies various operating environments. By adding an additional 'Device ID' field to table 1 above, readers will note that the exemplary table may be augmented to accommodate multiple mobile computing devices. Readers will further note that the exemplary table above is for explanation only and not for limitation.

[0086] Using the exemplary noise model repository above, the noise monitor (202) of FIG. 2 may configure the speech engine (153) with the noise model identified by 'Noise_Model_0' when the mobile computing device is in the 'Home' operating environment. The 'Home' operating environment may be further specified by specific geographic coordinates or network identifiers in the device location repository (201) as described above. Similarly, using the exemplary noise model repository above, the noise monitor (202) of FIG. 2 may configure the speech engine (153) with the noise model identified by 'Noise_Model_1' when the mobile computing device is in the 'Office' operating environment. Furthermore, the noise monitor (202) of FIG. 2 may configure the speech engine (153) with the noise model identified by 'Noise_Model_2' when the mobile computing device is in the 'Car' operating environment.

[0087] In the example of FIG. 2, the noise monitor (202) maintains the noise model repository (200) that associates noise models (205) with various operating environments of a multimodal device. The noise monitor (202) may create entries for the noise model repository (200) automatically as the mobile computing device changes location from one operating environment to another or by receiving user-specified associations through a user interface provided by the noise monitor (202). In a thin client architecture, such as the one illustrated in FIG. 2, where the noise monitor (202) may not

be installed on the multimodal device, a noise monitor client may be installed on the multimodal device that provides a user interface to a user to receive entries for the noise model repository (200). The noise monitor client then in turn provides the user input to the noise monitor (202) installed on the voice server (151). To ensure that the environments for which a user specifies noise models match the locations managed by the device location manager (120), the noise monitor (202) may register with the device location manager (120) to receive notification of changes to the location definitions specified in the device location repository (201).

[0088] Also stored in RAM (168) is an operating system (154). Operating systems useful in voice servers according to embodiments of the present invention include UNIX™, Linux™, Microsoft Vista™, IBM's AIX™, IBM's i5/OS™, and others as will occur to those of skill in the art. Operating system (154), voice server application (188), voice interpreter (192), speech engine (153), device location manager (120), device location repository (201), and configuration parameter repository (200) in the example of FIG. 2 are shown in RAM (168), but many components of such software typically are stored in non-volatile memory also, for example, on a disk drive (170).

[0089] Voice server (151) of FIG. 2 includes bus adapter (158), a computer hardware component that contains drive electronics for high speed buses, the front side bus (162), the video bus (164), and the memory bus (166), as well as drive electronics for the slower expansion bus (160). Examples of bus adapters useful in voice servers according to embodiments of the present invention include the Intel Northbridge, the Intel Memory Controller Hub, the Intel Southbridge, and the Intel I/O Controller Hub. Examples of expansion buses useful in voice servers according to embodiments of the present invention include Peripheral Component Interconnect ('PCI') and PCI-Extended ('PCI-X') bus, as well as PCI Express ('PCIe') point to point expansion architectures and others.

[0090] Voice server (151) of FIG. 2 includes disk drive adapter (172) coupled through expansion bus (160) and bus adapter (158) to processor (156) and other components of the voice server (151). Disk drive adapter (172) connects non-volatile data storage to the voice server (151) in the form of disk drive (170). Disk drive adapters useful in voice servers include Integrated Drive Electronics ('IDE') adapters, Small Computer System Interface ('SCSI') adapters, and others as will occur to those of skill in the art. In addition, non-volatile computer memory may be implemented for a voice server as an optical disk drive, electrically erasable programmable read-only memory (so-called 'EEPROM' or 'Flash' memory), RAM drives, and so on, as will occur to those of skill in the art.

[0091] The example voice server of FIG. 2 includes one or more input/output ('I/O') adapters (178). I/O adapters in voice servers implement user-oriented input/output through, for example, software drivers and computer hardware for controlling output to display devices such as computer display screens, as well as user input from user input devices (181) such as keyboards and mice. The example voice server of FIG. 2 includes a video adapter (209), which is an example of an I/O adapter specially designed for graphic output to a display device (180) such as a display screen or computer monitor. Video adapter (209) is connected to processor (156) through a high speed video bus (164), bus adapter (158), and the front side bus (162), which is also a high speed bus.

[0092] The exemplary voice server (151) of FIG. 2 includes a communications adapter (167) for data communications with other computers (182) and for data communications with a data communications network (100). Such data communications may be carried out serially through RS-232 connections, through external buses such as a Universal Serial Bus ('USB'), through data communications data communications networks such as IP data communications networks, and in other ways as will occur to those of skill in the art. Communications adapters implement the hardware level of data communications through which one computer sends data communications to another computer, directly or through a data communications network. Examples of communications adapters useful for adjusting a speech engine for a mobile computing device based on background noise according to embodiments of the present invention include modems for wired dial-up communications, Ethernet (IEEE 802.3) adapters for wired data communications network communications, and 802.11 adapters for wireless data communications network communications.

[0093] For further explanation, FIG. 3 sets forth a functional block diagram of exemplary apparatus for adjusting a speech engine for a mobile computing device based on background noise in a thin client architecture according to embodiments of the present invention. The example of FIG. 3 includes a mobile computing device implemented as a multimodal device (152) and a voice server (151) connected for data communication by a VOIP connection (216) through a data communications network (100). A multimodal application (195) runs on the multimodal device (152), and a voice server application (188) runs on the voice server (151). The multimodal client application (195) may be a set or sequence of X+V or SALT documents that execute on multimodal browser (196), a Java voice application that executes on the Java Virtual Machine (101), or a multimodal application implemented in other technologies as may occur to those of skill in the art. The example multimodal device of FIG. 3 also includes a sound card (174), which is an example of an I/O adapter specially designed for accepting analog audio signals from a microphone (176) and converting the audio analog signals to digital form for further processing by a codec (183).

[0094] In addition to the multimodal sever application (188), the voice server (151) also has installed upon it a speech engine (153) with an ASR engine (150), grammars (104), lexicons (106), acoustic models (108) with noise models (205), and a TTS engine (194), as well as a voice interpreter (192), configuration parameter repository (200), a device location manager (120), and a device location repository (201). The voice interpreter (192) of FIG. 3 is implemented as a VoiceXML interpreter that interprets and executes VoiceXML dialog instructions received from the multimodal application and provided to voice interpreter (192) through voice server application (188). VoiceXML input to the voice interpreter (192) may originate from the multimodal application (195) implemented as an X+V client running remotely on the multimodal device (152). As noted above, the multimodal application (195) also may be implemented as a Java client application running remotely on the multimedia device (152), a SALT application running remotely on the multimedia device (152), and in other ways as may occur to those of skill in the art.

[0095] The voice server (151) of FIG. 3 also has installed upon it a noise monitor (202). The noise monitor (202) of FIG. 3 operates generally for adjusting the speech engine

(153) for a mobile computing device based on background noise according to embodiments of the present invention. The noise monitor (202) of FIG. 3 operates generally for adjusting the speech engine (153) for a mobile computing device based on background noise according to embodiments of the present invention by: sampling, through a microphone of the mobile computing device, background noise for a plurality of operating environments in which the mobile computing device operates; generating, for each operating environment, a noise model in dependence upon the sampled background noise for that operating environment; and configuring the speech engine (153) for the mobile computing device with the noise model for the operating environment in which the mobile computing device currently operates.

[0096] The noise monitor (202) of FIG. 3 is operatively coupled to the microphone (176) of the mobile computing device (152) through the data communication network (100) and the noise monitor client (203) installed on the mobile computing device (152). The noise monitor client (203) interfaces with the hardware and software components of the multimodal device (152) on behalf of the noise monitor (202) installed on the voice server (151). Through the noise monitor client (203), the noise monitor (202) may periodically sample background noise in the environment in which the mobile computing device (152) operates. In addition, the noise monitor client (203) may expose an API that allows the JVM (101), the multimodal application (195), or the multimodal browser (196) to communicate with the noise monitor (202).

[0097] The noise monitor (202) of FIG. 3 may also operate for adjusting a speech engine (153) for a mobile computing device based on background noise according to embodiments of the present invention by: detecting a change in the mobile computing device (152) from a previous operating environment to a current operating environment; determining whether a noise model exists for the current operating environment; sampling background noise for the current operating environment if a noise model does not exist for the current operating environment; generating a noise model for the current operating environment if a noise model does not exist for the current operating environment; and configuring the speech engine (153) for the mobile computing device (152) with the generated noise model. If a noise model does exist for the current operating environment, the noise monitor (202) may operate to adjust a speech engine (153) for a mobile computing device based on background noise according to embodiments of the present invention by: selecting the noise model for the current operating environment and configuring the speech engine (153) for the mobile computing device (152) with the selected noise model. The noise monitor (202) of FIG. 3 may determine whether a noise model exists for the current operating environment using the noise model repository (200) that associates noise models with various operating environments.

[0098] In the example of FIG. 3, the noise monitor (202) detects a change in the mobile computing device (152) from a previous operating environment to a current operating environment by receiving a notification message from a device location manager (120). The device location manager (120) of FIG. 3 monitors the location of the mobile computing device (152) and provides the noise monitor (202) with notification messages regarding the location of the mobile computing device (152). The device location manager (120) of FIG. 3 monitors the location of the mobile computing device (152) based on position information provided by a position detection component (155) of the device (152). The position detection component (155) communicates with other devices to determine the mobile computing device's position infor-

mation. In the example of FIG. 3, the position detection component (155) communicates with a group of GPS satellites (102) or the devices in a PAN (103) to determine the mobile computing device's position information. Using the position information provided by the position detection component (155), the device location manager (120) may identify the location of the device (152) using the device location repository (201) that associates position information such as geographic coordinates or network identifiers with location or environment semantics such as house, car, work, school, and so on.

[0099] In the example of FIG. 3, the device location manager (120) is operatively coupled to the position detection component (155) of the multimodal device (152) through network (100), a device location manager client (131), and position detection module (121). The position detection module (121) operates as a software driver for the position detection component and provides an interface for other software components, such as the device location manager (120) and the device location manager client (131), to interact with the position detection component. The device location manager client (131) of FIG. 3 is a software module that provides data communications between the position detection module (121) on the multimodal device (152) and the device location manager (120) on voice server (151). In addition, the device location manager client (131) may also provide a user interface to user (128) for receiving a location definition that associates location semantics with geographic coordinates provided by the position detection component (155). The device location manager client (131) of FIG. 3 may, in turn, provide the location definition to the device location manager (120), which stores the location definition in the device location repository (201).

[0100] As mentioned above, the VOIP connection (216) connects for data communication the multimodal device (152) and the voice server (151). VOIP stands for 'Voice Over Internet Protocol,' a generic term for routing speech over an IP-based data communications network. The speech data flows over a general-purpose packet-switched data communications network, instead of traditional dedicated, circuit-switched voice transmission lines. Protocols used to carry voice signals over the IP data communications network are commonly referred to as 'Voice over IP' or 'VOIP' protocols. VOIP traffic may be deployed on any IP data communications network, including data communications networks lacking a connection to the rest of the Internet, for instance on a private building-wide local area data communications network or 'LAN.'

[0101] Many protocols are used to effect VOIP. The two most popular types of VOIP are effected with the IETF's Session Initiation Protocol ('SIP') and the ITU's protocol known as 'H.323.' SIP clients use TCP and UDP port 5060 to connect to SIP servers. SIP itself is used to set up and tear down calls for speech transmission. VOIP with SIP then uses RTP for transmitting the actual encoded speech. Similarly, H.323 is an umbrella recommendation from the standards branch of the International Telecommunications Union that defines protocols to provide audio-visual communication sessions on any packet data communications network.

[0102] The apparatus of FIG. 3 operates in a manner that is similar to the operation of the system of FIG. 2 described above. Multimodal application (195) is a user-level, multimodal, client-side computer program that presents a voice interface to user (128), provides audio prompts and responses (314) and accepts input speech for recognition (315). Multimodal application (195) provides a speech interface through which a user may provide oral speech for recognition through

microphone (176) and have the speech digitized through an audio amplifier (185) and a coder/decoder ('codec') (183) of a sound card (174) and provide the digitized speech for recognition to ASR engine (150). Multimodal application (195) then packages the digitized speech in a recognition request message according to a VOIP protocol, and transmits the speech to voice server (151) through the VOIP connection (216) on the network (100).

[0103] Voice server application (188) provides voice recognition services for multimodal devices by accepting dialog instructions, VoiceXML segments, and returning speech recognition results, including text representing recognized speech, text for use as variable values in dialogs, and output from execution of semantic interpretation scripts as well as voice prompts. Voice server application (188) supports text-to-speech ('TTS') conversion for voice prompts and voice responses to user input in multimodal applications such as, for example, X+V applications, SALT applications, or Java Speech applications.

[0104] The voice server application (188) receives speech for recognition from a user and passes the speech through API calls to voice interpreter (192) which in turn uses an ASR engine (150) for speech recognition. The ASR engine receives digitized speech for recognition, uses frequency components of the digitized speech to derive an SFV, uses the SFV to infer phonemes for the word from the language-specific acoustic model (108), and uses the phonemes to find the speech in the lexicon (106). The ASR engine then compares speech found as words in the lexicon to words in a grammar (104) to determine whether words or phrases in speech are recognized by the ASR engine.

[0105] The multimodal application (195) is operatively coupled to the ASR engine (150). In this example, the operative coupling between the multimodal application and the ASR engine (150) is implemented with a VOIP connection (216) through a voice services module (130), then through the voice server application (188) and the voice interpreter (192). Depending on whether the multimodal application is implemented in X+V, Java, or SALT, the voice interpreter (192) may be implemented using a VoiceXML interpreter, a VoiceXML interpreter exposing a Java interface, or a SALT interpreter. The voice services module (130) is a thin layer of functionality, a module of computer program instructions, that presents an API (316) for use by an application level program in providing dialog instructions and speech for recognition to a voice server application (188) and receiving in response voice prompts and other responses. In this example, application level programs are represented by multimodal application (195), JVM (101), and multimodal browser (196).

[0106] The voice services module (130) provides data communications services through the VOIP connection and the voice server application (188) between the multimodal device (152) and the voice interpreter (192). The API (316) is the same API presented to applications by a voice interpreter when the voice interpreter is installed on the multimodal device in a thick client architecture. So from the point of view of an application calling the API (316), the application is calling the VoiceXML interpreter directly. The data communications functions of the voice services module (130) are transparent to applications that call the API (316). At the application level, calls to the API (316) may be issued from the multimodal browser (196), which provides an execution environment for the multimodal application (195) when the multimodal application is implemented with X+V. And calls to the API (316) may be issued from the JVM (101), which

provides an execution environment for the multimodal application (195) when the multimodal application is implemented with Java.

[0107] Adjusting a speech engine for a mobile computing device based on background noise according to embodiments of the present invention in thick client architectures is generally implemented with multimodal devices, that is, automated computing machinery or computers. In the system of FIG. 1, for example, all the multimodal devices (152) are implemented to some extent at least as computers. For further explanation, therefore, FIG. 4 sets forth a block diagram of automated computing machinery comprising an example of a computer useful as a multimodal device (152) in adjusting a speech engine for a mobile computing device based on background noise according to embodiments of the present invention. In a multimodal device implementing a thick client architecture as illustrated in FIG. 4, the multimodal device (152) has no connection to a remote voice server containing a voice interpreter and a speech engine. All the components needed for speech synthesis and voice recognition in adjusting a speech engine for a mobile computing device based on background noise according to embodiments of the present invention are installed or embedded in the multimodal device itself.

[0108] The example multimodal device (152) of FIG. 4 includes several components that are structured and operate similarly as do parallel components of the voice server, having the same drawing reference numbers, as described above with reference to FIG. 2: at least one computer processor (156), frontside bus (162), RAM (168), high speed memory bus (166), bus adapter (158), video adapter (209), video bus (164), expansion bus (160), communications adapter (167), I/O adapter (178), disk drive adapter (172), an operating system (154), a voice interpreter (192), a speech engine (153), and so on. As in the system of FIG. 2, the speech engine in the multimodal device of FIG. 4 includes an ASR engine (150), grammars (104), lexicons (106), acoustic models (108) that each have a noise model (205), and a TTS engine (194). The voice interpreter (192) may be implemented as a VoiceXML interpreter that administers dialogs by processing the dialog instructions sequentially in accordance with a VoiceXML Form Interpretation Algorithm ('FIA').

[0109] The speech engine (153) in this kind of embodiment, a thick client architecture, often is implemented as an embedded module in a small form factor device such as a handheld device, a mobile phone, PDA, and the like. An example of an embedded speech engine that may be improved for adjusting a speech engine for a mobile computing device based on background noise according to embodiments of the present invention is IBM's Embedded ViaVoice Enterprise. The example multimodal device of FIG. 4 also includes a sound card (174), which is an example of an I/O adapter specially designed for accepting analog audio signals from a microphone (176) and converting the audio analog signals to digital form for further processing by a codec (183). The sound card (174) is connected to processor (156) through expansion bus (160), bus adapter (158), and front side bus (162).

[0110] Also stored in RAM (168) in this example is a multimodal application (195), a module of computer program instructions capable of operating a multimodal device as an apparatus that supports multiple modes of user interaction, including a voice mode and one or more non-voice modes. The multimodal application (195) implements speech recognition by accepting speech for recognition from a user and sending the speech for recognition through API calls to the ASR engine (150). The multimodal application (195) imple-

ments speech synthesis generally by sending words to be used as prompts for a user to the TTS engine (194). As an example of thick client architecture, the multimodal application (195) in this example does not send speech for recognition across a network to a voice server for recognition, and the multimodal application (195) in this example does not receive synthesized speech, TTS prompts and responses, across a network from a voice server. All grammar processing, voice recognition, and text to speech conversion in this example is performed in an embedded fashion in the multimodal device (152) itself.

[0111] More particularly, multimodal application (195) in this example is a user-level, multimodal, client-side computer program that provides a speech interface through which a user may provide oral speech for recognition through microphone (176), have the speech digitized through an audio amplifier (185) and a coder/decoder ('codec') (183) of a sound card (174) and provide the digitized speech for recognition to ASR engine (150). The multimodal application (195) may be implemented as a set or sequence of X+V documents executing in a multimodal browser (196) or microbrowser that passes VoiceXML grammars and digitized speech by calls through an API (316) directly to an embedded voice interpreter (192) for processing. The embedded voice interpreter (192) may in turn issue requests for speech recognition through API calls directly to the embedded ASR engine (150). Multimodal application (195) also can provide speech synthesis, TTS conversion, by API calls to the embedded TTS engine (194) for voice prompts and voice responses to user input.

[0112] In a further class of exemplary embodiments, the multimodal application (195) may be implemented as a Java voice application that executes on Java Virtual Machine (102) and issues calls through an API of the voice interpreter (192) for speech recognition and speech synthesis services. In further exemplary embodiments, the multimodal application (195) may be implemented as a set or sequence of SALT documents executed on a multimodal browser (196) or microbrowser that issues calls through an API of the voice interpreter (192) for speech recognition and speech synthesis services. In addition to X+V, SALT, and Java implementations, multimodal application (195) may be implemented in other technologies as will occur to those of skill in the art, and all such implementations are well within the scope of the present invention.

[0113] Stored in RAM (168) of FIG. 4 is a noise monitor (202) that operates generally for adjusting the speech engine (153) for a mobile computing device based on background noise according to embodiments of the present invention. The noise monitor (202) of FIG. 4 operates generally for adjusting the speech engine (153) for a mobile computing device (152) based on background noise according to embodiments of the present invention by: sampling, through a microphone (176) of the mobile computing device (152), background noise for a plurality of operating environments in which the mobile computing device (152) operates; generating, for each operating environment, a noise model in dependence upon the sampled background noise for that operating environment; and configuring the speech engine (153) for the mobile computing device (152) with the noise model for the operating environment in which the mobile computing device currently operates.

[0114] The noise monitor (202) of FIG. 4 may also operate for adjusting a speech engine (153) for a mobile computing device (152) based on background noise according to embodiments of the present invention by: detecting a change in the mobile computing device (152) from a previous oper-

ating environment to a current operating environment; determining whether a noise model exists for the current operating environment; sampling background noise for the current operating environment if a noise model does not exist for the current operating environment; generating a noise model for the current operating environment if a noise model does not exist for the current operating environment; and configuring the speech engine (153) for the mobile computing device (152) with the generated noise model. If a noise model does exist for the current operating environment, the noise monitor (202) may operate to adjust a speech engine (153) for a mobile computing device based on background noise according to embodiments of the present invention by: selecting the noise model for the current operating environment and configuring the speech engine (153) for the mobile computing device (152) with the selected noise model. The noise monitor (202) of FIG. 4 may determine whether a noise model exists for the current operating environment using the noise model repository (200) that associates noise models with various operating environments.

[0115] In the example of FIG. 4, the noise monitor (202) detects a change in the mobile computing device (152) from a previous operating environment to a current operating environment by receiving a notification message from a device location manager (120). The device location manager (120) of FIG. 4 monitors the location of the mobile computing device (152) and provides the noise monitor (202) with notification messages regarding the location of the mobile computing device (152). The device location manager (120) of FIG. 4 monitors the location of the mobile computing device (152) based on position information provided by a position detection component (155) through a position detection module (121). The position detection component (155) is a component that communicates with other devices to determine the mobile computing device's position information. In the example of FIG. 4, the position detection component (155) communicates with a group of GPS satellites (102) or the devices in a PAN (103) to determine the mobile computing device's position information. Using the position information provided by the position detection component (155), the device location manager (120) may identify the location of the device (152) using the device location repository (201) that associates position information such as geographic coordinates or network identifiers with location or environment semantics such as house, car, work, school, and so on. The position detection module (121) operates in RAM (168) as a software driver for the position detection component (155) and provides an interface for other software components, such as the device location manager (120), to interact with the position detection component.

[0116] The multimodal application (195) of FIG. 4 is operatively coupled to the ASR engine (150). In this example, the operative coupling between the multimodal application and the ASR engine (150) is implemented through the voice interpreter (192). Depending on whether the multimodal application is implemented in X+V, Java, or SALT, the voice interpreter (192) may be implemented as a VoiceXML interpreter, a VoiceXML interpreter exposing a Java interface, or SALT interpreter. When the multimodal application (195) is implemented in X+V, the operative coupling is effected through the multimodal browser (196), which provides an operating environment and an interpreter for the X+V application, and then through a VoiceXML interpreter, which passes grammars and voice utterances for recognition to the ASR engine. When the multimodal application (195) is implemented in Java Speech, the operative coupling is effected through the JVM (101), which provides an operating

environment for the Java application and the voice interpreter (192), which passes grammars and voice utterances for recognition to the ASR engine (150). When the multimodal application (195) is implemented in SALT, the operative coupling is effected through the multimodal browser (196), which provides an operating environment and an interpreter for the X+V application, and then through a SALT interpreter, which passes grammars and voice utterances for recognition to the ASR engine.

[0117] The multimodal application (195) in this example, running on a multimodal device (152) that contains its own voice interpreter (192) and its own speech engine (153) with no network or VOIP connection to a remote voice server containing a remote VoiceXML interpreter or a remote speech engine, is an example of a so-called ‘thick client architecture,’ so-called because all of the functionality for processing voice mode interactions between a user and the multimodal application is implemented on the multimodal device itself.

[0118] For further explanation, FIG. 5 sets forth a flow chart illustrating an exemplary method of adjusting a speech engine (153) for a mobile computing device based on background noise according to embodiments of the present invention. The mobile computing device described with reference to FIG. 5 is operatively coupled to a microphone used to record the background noise for the environments in which the mobile computing device operates. The example of FIG. 5 illustrates three operating environments (500) in which the mobile computing device operates: a house, a car, and a helicopter.

[0119] In the example of FIG. 5, the speech engine (153) adjusted based on background noise is similar to the speech engines described above. The speech engine (153) includes acoustic models (108), lexicons (106), and grammars (104) used by the speech engine (153) to provide speech recognition and synthesization services. The speech engine (153) provides speech recognition using an ASR engine (150) and speech synthesization using a TTS engine (194).

[0120] The method of FIG. 5 includes sampling (502), through the microphone, background noise (506) for a plurality of operating environments (500) in which the mobile computing device operates. The noise monitor (202) may sample (502) background noise (506) for a plurality of operating environments (500) according to the method of FIG. 5 by periodically at a predetermined time period (504): determining the lack of any foreground noise and digitizing the background noise (506) for the environment (500) in which the mobile computing device is operating. The predetermined time period (504) at which the noise monitor (202) samples the background noise (506) may be a static time period or a time period that is dynamically generated. For example, the predetermined time period (504) may specify sampling the background noise every five minutes, three minutes after the most recent user interaction with the mobile computing device, and so on. Because foreground noise generally results in a sudden change in the input to a microphone, the noise monitor (202) may determine the lack of any foreground noise by identifying over a period of time when the input to the microphone does not vary beyond a predefined threshold. The period of time used to identify the lack of any foreground noise may be a several milliseconds, several seconds, or any other time period as will occur to those of skill in the art.

[0121] In some embodiments, the noise monitor (202) may sample the background noise (506) while a user is not interacting with the mobile computing device to avoid having additional noise from the user interaction included the background noise for the operating environment. In other embodi-

ments, the noise monitor (202) may sample the background noise (506) while the user is interacting with the device. For example, the noise monitor may sample the background noise immediately before or after the user provides a voice utterance for speech recognition.

[0122] The method of FIG. 5 includes generating (508), for each operating environment (500), a noise model (205) in dependence upon the sampled background noise (506) for that operating environment (500). Each noise model (205) of FIG. 5 represents the waveform of the recorded background noise (506) used to adjust the speech engine (153). The noise model (205) of FIG. 5 may be implemented as the first twelve or thirteen Fourier or frequency domain components of the waveform of the background noise (506). The noise monitor (202) may generate (508) a noise model (205) for each operating environment (500) according to the method of FIG. 5 by transforming the background noise (506) for each operating environment (500) from the time domain to the frequency domain using, for example, a Fourier transformation algorithm and selecting a predetermined number of frequency domain components to represent the background noise waveform in the noise model (205).

[0123] The method of FIG. 5 includes configuring (510) the speech engine (153) for the mobile computing device with the noise model (205) for the operating environment (500) in which the mobile computing device currently operates. The noise monitor (202) may configure (510) the speech engine (153) with the noise model (205) for the operating environment (500) in which the mobile computing device currently operates according to the method of FIG. 5 by altering the current configuration (206) of the speech engine (153). The current configuration (206) of FIG. 5 specifies the operating settings, parameters, and other variable used by the speech engine (153) to provide speech recognition and synthesization services. The current configuration (206) may specify the active lexicon and the active acoustic model and noise model used by the speech engine (153) to recognize and generate human speech. The current configuration (206) may also specify a current configuration for the automatic speech recognition (‘ASR’) engine such as, for example, speech transition times, silence detection times, speech timeouts, and gain maps. Furthermore, the current configuration (206) may also specify a configuration for use by a text-to-speech (‘TTS’) engine such as, for example, the voice used to synthesize the text.

[0124] The noise monitor (202) may alter the current configuration (206) of the speech engine (153) to configure the speech engine (153) with a noise model (205) for the current operating environment according to the method of FIG. 5 by updating the current acoustic model for the speech engine (153) with the noise model for the operating environment (500) in which the mobile computing device currently operates. Rather than updating the current acoustic model, the noise monitor (202) may also alter the current configuration (206) of the speech engine (153) to configure the speech engine (153) with a noise model (205) for the current operating environment according to the method of FIG. 5 by replacing the current acoustic model for the speech engine (153) with another acoustic model already incorporating the noise model for the operating environment (500) in which the mobile computing device currently operates.

[0125] The explanation above with reference to FIG. 5 describes a noise monitor that periodically adjusts a speech engine for a mobile computing device based on the background noise for the device’s current operating environment without taking into account the device’s location. In some other embodiments, however, the noise monitor may incor-

porate the device's location in adjusting a speech engine for the device based on the background noise. For further explanation, consider FIG. 6 that sets forth a flow chart illustrating a further exemplary method of adjusting a speech engine for a mobile computing device based on background noise according to embodiments of the present invention. The mobile computing device described with reference to FIG. 6 is operatively coupled to a microphone used to record the background noise for the environments in which the mobile computing device operates.

[0126] In the example of FIG. 6, the speech engine (153) adjusted based on background noise is similar to the speech engines described above. The speech engine (153) includes acoustic models (108), lexicons (106), and grammars (104) used by the speech engine (153) to provide speech recognition and synthesization services. The speech engine (153) provides speech recognition using an ASR engine (150) and speech synthesization using a TTS engine (194). The speech engine (153) also includes a current configuration (206) that specifies the operating settings, parameters, and other variable used by the speech engine (153) to provide speech recognition and synthesization services.

[0127] The method of FIG. 6 includes detecting (600) a change in the mobile computing device from a previous operating environment to a current operating environment (602). The noise monitor (202) may detect (600) a change in the mobile computing device from a previous operating environment to a current operating environment (602) according to the method of FIG. 6 by receiving a location change notification message from a device location manager for the mobile computing device. The location change notification message may provide the noise monitor (202) with location information in the form of position information such as geographic coordinates, network identifiers, and the like, as well as in the form of semantics representing such position information such as, for example, 'home,' 'work,' 'church,' or 'car.' The location device manager (120) may have obtained the position information from a position detection component installed in the mobile computing device that communicates with GPS satellites (102) or other devices in a PAN (103).

[0128] The method of FIG. 6 also includes determining (604) whether a noise model exists for the current operating environment (602). The noise monitor (202) may determine (604) whether a noise model (205) exists for the current operating environment (602) according to the method of FIG. 6 by identifying whether a noise model is associated with an identifier for the current operating environment (602) in a noise model repository. As described above, a noise model repository may be implemented as a data structure or table that associates various noise models (205) with different operating environments. If a noise model is associated with an identifier for the current operating environment (602) in the noise model repository, then a noise model (205) exists for the current operating environment (602). If a noise model is not associated with an identifier for the current operating environment (602) in the noise model repository, then a noise model (205) does not exist for the current operating environment (602).

[0129] The method of FIG. 6 includes sampling (606) background noise (506) for the current operating environment (602) if a noise model (205) does not exist for the current operating environment (602). The noise monitor (202) may sample (606) background noise (506) for the current operating environment (602) according to the method of FIG. 6 by determining the lack of any foreground noise and digitizing

the background noise (506) for the current operating environment (602) in which the mobile computing device is operating.

[0130] The method of FIG. 6 includes generating (608) a noise model (205) for the current operating environment (602) if a noise model (205) does not exist for the current operating environment (602). The noise monitor (202) may generate (608) a noise model (205) for the current operating environment (602) according to the method of FIG. 6 by transforming the background noise (506) for the current operating environment (602) from the time domain to the frequency domain using, for example, a Fourier transformation algorithm and selecting a predetermined number of frequency domain components to represent the background noise waveform in the noise model (205). Generating (608) a noise model (205) for the current operating environment (602) according to the method of FIG. 6 may also include associating the generated noise model (205) with the current operating environment (602) in a noise model repository such as, for example, the noise model repository described above with reference to FIG. 2.

[0131] The method of FIG. 6 includes selecting (610) the noise model (612) for the current operating environment (602) if a noise model (205) exists for the current operating environment (602). The noise monitor (202) may select (610) the noise model (612) for the current operating environment (602) according to the method of FIG. 6 by retrieving the noise model (612) associated with the current operating environment (602) in a noise model repository.

[0132] The method of FIG. 6 also includes configuring (614) the speech engine (153) for the mobile computing device with the selected noise model (612). The noise monitor (202) may configure (614) the speech engine (153) for the mobile computing device with the selected noise model (612) according to the method of FIG. 6 by altering the current configuration (206) of the speech engine (153). The noise monitor (202) may alter the current configuration (206) of the speech engine (153) to configure (614) the speech engine (153) for the mobile computing device with the selected noise model (612) according to the method of FIG. 5 by updating the current acoustic model for the speech engine (153) with the noise model for the operating environment (500) in which the mobile computing device currently operates. Rather than updating the current acoustic model, the noise monitor (202) may also alter the current configuration (206) of the speech engine (153) to configure (614) the speech engine (153) for the mobile computing device with the selected noise model (612) according to the method of FIG. 5 by replacing the current acoustic model for the speech engine (153) with another acoustic model already incorporating the noise model for the operating environment (500) in which the mobile computing device currently operates.

[0133] In the example of FIG. 6, the mobile computing device is implemented as a multimodal device having installed upon it a multimodal application operatively coupled to the speech engine (153). The multimodal device described with reference to FIG. 6 supports multiple modes of user interaction with the multimodal application. The modes of user interaction include a voice mode and one or more non-voice modes.

[0134] The method of FIG. 6 includes performing (616) speech recognition for a multimodal application using the speech engine (153) configured for the mobile computing device with the noise model (612) for the operating environment (602) in which the mobile computing device currently operates. Performing (616) speech recognition for a multimodal application using the speech engine (153) according to

the method of FIG. 6 may be carried out by receiving a user's voice utterance from the multimodal application, passing the voice utterance to the ASR engine (150) for speech recognition using the acoustic model (108) with the selected noise model (612), the lexicon (106), and the grammar (104), and receiving from the ASR engine (150) a textual representation of the voice utterance provided by the user. When the multimodal application is implemented in part using VoiceXML, the multimodal application may access the representation of the voice utterance in the 'application.lastresult\$' array. Readers will note that adjusting a speech engine for a mobile computing device based on background noise according to embodiments of the present invention advantageously increases the speech recognition accuracy of the recognition results returned to the multimodal application.

[0135] Exemplary embodiments of the present invention are described largely in the context of a fully functional computer system for adjusting a speech engine for a mobile computing device based on background noise. Readers of skill in the art will recognize, however, that the present invention also may be embodied in a computer program product disposed on computer readable media for use with any suitable data processing system. Such computer readable media may be transmission media or recordable media for machine-readable information, including magnetic media, optical media, or other suitable media. Examples of recordable media include magnetic disks in hard drives or diskettes, compact disks for optical drives, magnetic tape, and others as will occur to those of skill in the art. Examples of transmission media include telephone networks for voice communications and digital data communications networks such as, for example, EthernetTM and networks that communicate with the Internet Protocol and the World Wide Web. Persons skilled in the art will immediately recognize that any computer system having suitable programming means will be capable of executing the steps of the method of the invention as embodied in a program product. Persons skilled in the art will recognize immediately that, although some of the exemplary embodiments described in this specification are oriented to software installed and executing on computer hardware, nevertheless, alternative embodiments implemented as firmware or as hardware are well within the scope of the present invention.

[0136] It will be understood from the foregoing description that modifications and changes may be made in various embodiments of the present invention without departing from its true spirit. The descriptions in this specification are for purposes of illustration only and are not to be construed in a limiting sense. The scope of the present invention is limited only by the language of the following claims.

1. A method of adjusting a speech engine for a mobile computing device based on background noise, the mobile computing device operatively coupled to a microphone, the method comprising:

determining whether a previously-generated background noise model exists for a current operating environment of the mobile computing device; and

if it is determined that the previously-generated background noise model exists for the current operating environment, configuring the speech engine for the mobile computing device with the previously-generated background noise model,

wherein the previously-generated noise model was generated at least in part by sampling background noise for the current operating environment.

2. The method of claim 1, wherein if it is determined that no previously-generated background noise model exists for the current operating environment, the method further comprises:

sampling, through the microphone, background noise for the current operating environment;
generating a background noise model based at least in part on the sampled background noise; and
configuring the speech engine for the mobile computing device with the background noise model.

3. The method of claim 1, further comprising:

detecting a change in an operating environment of the mobile computing device.

4. The method of claim 1, wherein configuring the speech engine for the mobile computing device with the previously-generated background noise model for the current operating environment comprises updating a current acoustic model for the speech engine with the previously-generated background noise model.

5. The method of claim 1, wherein configuring the speech engine for the mobile computing device with the previously-generated background noise model for the current operating environment comprises replacing a current acoustic model for the speech engine with another acoustic model already incorporating the previously-generated background noise model.

6. The method of claim 1, wherein the mobile computing device is implemented as a multimodal device having installed upon it a multimodal application, the multimodal device supporting multiple modes of user interaction with the multimodal application, the modes of user interaction including a voice mode and one or more non-voice modes, the multimodal application operatively coupled to the speech engine, further comprising performing speech recognition for the multimodal application using the speech engine configured for the mobile computing device with the previously-generated background noise model for the current operating environment.

7. The method of claim 1, wherein the current operating environment includes a location, wherein the location is one of a home of a user of the mobile computing device, a car of the user, and/or a workplace of the user.

8. An apparatus for adjusting a speech engine for a mobile computing device based on background noise, the mobile computing device operatively coupled to a microphone, the apparatus comprising:

at least one processor configured to:

determine whether a previously-generated background noise model exists for a current operating environment of the mobile computing device; and

if it is determined that the previously-generated background noise model exists for the current operating environment, configuring the speech engine for the mobile computing device with the previously-generated background noise model,

wherein the previously-generated noise model was generated at least in part by sampling background noise for the current operating environment.

9. The apparatus of claim 8, wherein if it is determined that no previously-generated background noise model exists for the current operating environment, the at least one processor is further configured to:

sample, through the microphone, background noise for the current operating environment;

generate a background noise model based at least in part on the sampled background noise; and configure the speech engine for the mobile computing device with the background noise model.

10. The apparatus of claim 8, wherein the at least one processor is further configured to:

detect a change in an operating environment of the mobile computing device.

11. The apparatus of claim 8, wherein the at least one processor is configured to configure the speech engine for the mobile computing device with the previously-generated background noise model for the current operating environment by updating a current acoustic model for the speech engine with the previously-generated background noise model.

12. The apparatus of claim 8, wherein the at least one processor is configured to configure the speech engine for the mobile computing device with the previously-generated background noise model for the current operating environment by replacing a current acoustic model for the speech engine with another acoustic model already incorporating the previously-generated background noise model.

13. The apparatus of claim 8, wherein:

the mobile computing device is implemented as a multimodal device having installed upon it a multimodal application, the multimodal device supporting multiple modes of user interaction with the multimodal application, the modes of user interaction including a voice mode and one or more non-voice modes, the multimodal application operatively coupled to the speech engine; and

the at least one processor is further configured to perform speech recognition for the multimodal application using the speech engine configured for the mobile computing device with the previously-generated background noise model for the current operating environment.

14. The apparatus of claim 8, wherein the current operating environment includes a location, wherein the location is one of a home of a user of the mobile computing device, a car of the user, and/or a workplace of the user.

15. At least one computer-readable recordable storage medium storing processor-executable instructions that, when executed by at least one processor, perform a method comprising:

determining whether a previously-generated background noise model exists for a current operating environment of the mobile computing device; and

if it is determined that the previously-generated background noise model exists for the current operating environment, configuring the speech engine for the mobile computing device with the previously-generated background noise model,

wherein the previously-generated noise model was generated at least in part by sampling background noise for the current operating environment.

16. The at least one computer-readable recordable storage medium of claim 15, wherein if it is determined that no previously-generated background noise model exists for the current operating environment, the method further comprises:

sampling, through the microphone, background noise for the current operating environment;

generating a background noise model based at least in part on the sampled background noise; and

configuring the speech engine for the mobile computing device with the background noise model.

17. The at least one computer readable recordable storage medium of claim 15, wherein the method further comprises:

detecting a change in an operating environment of the mobile computing device.

18. The at least one computer-readable recordable storage medium of claim 15, wherein configuring the speech engine for the mobile computing device with the previously-generated background noise model for the current operating environment comprises updating a current acoustic model for the speech engine with the previously-generated background noise model.

19. The at least one computer-readable recordable storage medium of claim 15, wherein configuring the speech engine for the mobile computing device with the previously-generated background noise model for the current operating environment comprises replacing a current acoustic model for the speech engine with another acoustic model already incorporating the previously-generated background noise model.

20. The at least one computer-readable recordable storage medium of claim 15, wherein the mobile computing device is implemented as a multimodal device having installed upon it a multimodal application, the multimodal device supporting multiple modes of user interaction with the multimodal application, the modes of user interaction including a voice mode and one or more non-voice modes, the multimodal application operatively coupled to the speech engine, further comprising performing speech recognition for the multimodal application using the speech engine configured for the mobile computing device with the previously-generated background noise model for the current operating environment.

21. The at least one computer-readable recordable storage medium of claim 15, wherein the current operating environment includes a location, wherein the location is one of a home of a user of the mobile computing device, a car of the user, and/or a workplace of the user.

* * * * *