



(51) International Patent Classification:

G10L 15/20 (2006.01) G10L 21/02 (2013.01)
G10L 15/26 (2006.01)

(21) International Application Number:

PCT/US2021/031512

(22) International Filing Date:

10 May 2021 (10.05.2021)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

63/022,269 08 May 2020 (08.05.2020) US

(71) Applicant: NUANCE COMMUNICATIONS, INC.

[US/US]; One Wayside Road, Burlington, Massachusetts 01803 (US).

(72) Inventors: SHARMA, Dushyant; 36 Cambridge Road,

Unit 27, Woburn, Massachusetts 01801 (US). NAYLOR,

Patrick A.; 63 Culver Lane, Reading, Berkshire RG6 1DX (GB).

(74) Agent: SARGEANT, Heath M. et al.; Holland & Knight LLP, 10 St. James Avenue, Boston, Massachusetts 02116 (US).

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

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ,

(54) Title: SYSTEM AND METHOD FOR MULTI-MICROPHONE AUTOMATED CLINICAL DOCUMENTATION

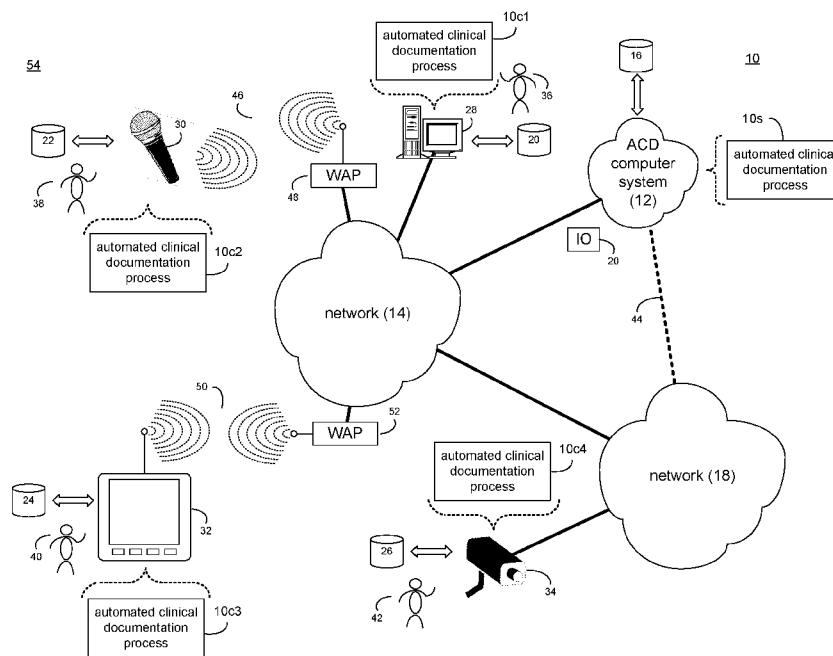


FIG. 1

(57) Abstract: A method, computer program product, and computing system for receiving a plurality of predefined beams associated with a microphone array. A plurality of predefined nulls associated with the microphone array may be received. One or more predefined beams from the plurality of predefined beams or one or more predefined nulls from the plurality of predefined nulls may be selected. A microphone array may obtain audio encounter information, via the microphone array, using at least one of the one or more selected beams and the one or more selected nulls.



UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

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

Published:

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

System and Method for Multi-Microphone Automated Clinical Documentation

Related Application(s)

[001] This application claims the benefit of the following U.S. Provisional Application No.: 63/022,269 filed on 08 May 2020, the contents of which are all incorporated herein by reference.

Background

[002] Automated Clinical Documentation (ACD) may be used, e.g., to turn transcribed conversational (e.g., physician, patient, and/or other participants such as patient's family members, nurses, physician assistants, etc.) speech into formatted (e.g., medical) reports. Such reports may be reviewed, e.g., to assure accuracy of the reports by the physician, scribe, etc.

[003] The process of capturing speech may include distant conversation automated speech recognition (DCASR). DCASR includes multiple microphone systems configured to record and recognize speech of one or more speakers (e.g., a microphone array and a single mobile microphone carried by a speaker). Conventional approaches to DCASR are subject to various challenges. For example, signal acquisition via a number of microphone systems (e.g., one or more microphone arrays and one or more mobile microphones) may be problematic. For example, when beamforming with single microphones, there are single channel noise reduction and de-reverberation methods that cannot be applied to multiple microphone systems. Additionally, conventional approaches to DCASR are unable to manage the combination of enhanced signals from various devices in a room. Conventional approaches to DCASR may encounter issues when locating where the speakers are at any given time, tracking their movements, and then identifying their identify. This may result in less accurate text transcriptions of recorded speech. In

addition, an automated speech recognition (ASR) system, that converts the audio to text, may be unable to process the various metadata and audio signals from multiple sources (e.g., multiple microphone systems).

Summary of Disclosure

[004] In one implementation, a computer-implemented method executed by a computer may include but is not limited to receiving a plurality of predefined beams associated with a microphone array. A plurality of predefined nulls associated with the microphone array may be received. One or more predefined beams from the plurality of predefined beams or one or more predefined nulls from the plurality of predefined nulls may be selected. A microphone array may obtain audio encounter information, via the microphone array, using at least one of the one or more selected beams and the one or more selected nulls.

[005] One or more of the following features may be included. The plurality of predefined beams may include one or more predefined beams configured to receive audio encounter information from one or more target speaker locations within an acoustic environment. The plurality of predefined nulls may include one or more predefined nulls configured to limit receiving audio encounter information from the one or more target speaker locations within the acoustic environment. Speaker metadata associated with one or more portions of audio encounter information may be received. Selecting at least one of the one or more beams and the one or more nulls may include selecting at least one of the one or more beams and the one or more nulls based upon, at least in part, the speaker metadata associated with the one or more portions of the audio encounter information. Selecting at least one of the one or more beams and the one or more nulls based upon, at least in part, the speaker metadata associated with the one or more portions of the audio encounter information may include selecting the one or more predefined beams configured to receive audio encounter information from the one or more target speaker locations within

the acoustic environment based upon, at least in part, the speaker metadata associated with the one or more portions of the audio encounter information. Selecting at least one of the one or more beams and the one or more nulls based upon, at least in part, the speaker metadata associated with the one or more portions of the audio encounter information may include selecting the one or more predefined nulls configured to limit receiving audio encounter information from the one or more target speaker locations within the acoustic environment based upon, at least in part, the speaker metadata associated with the one or more portions of the audio encounter information.

[006] In another implementation, a computer program product resides on a computer readable medium and has a plurality of instructions stored on it. When executed by a processor, the instructions cause the processor to perform operations including but not limited to receiving a plurality of predefined beams associated with a microphone array. A plurality of predefined nulls associated with the microphone array may be received. One or more predefined beams from the plurality of predefined beams or one or more predefined nulls from the plurality of predefined nulls may be selected. A microphone array may obtain audio encounter information, via the microphone array, using at least one of the one or more selected beams and the one or more selected nulls.

[007] One or more of the following features may be included. The plurality of predefined beams may include one or more predefined beams configured to receive audio encounter information from one or more target speaker locations within an acoustic environment. The plurality of predefined nulls may include one or more predefined nulls configured to limit receiving audio encounter information from the one or more target speaker locations within the acoustic environment. Speaker metadata associated with one or more portions of audio encounter information may be received. Selecting at least one of the one or more beams and the one or more nulls may include selecting at least one of the one or more beams and the one or more nulls based upon, at least in part, the speaker

metadata associated with the one or more portions of the audio encounter information. Selecting at least one of the one or more beams and the one or more nulls based upon, at least in part, the speaker metadata associated with the one or more portions of the audio encounter information may include selecting the one or more predefined beams configured to receive audio encounter information from the one or more target speaker locations within the acoustic environment based upon, at least in part, the speaker metadata associated with the one or more portions of the audio encounter information. Selecting at least one of the one or more beams and the one or more nulls based upon, at least in part, the speaker metadata associated with the one or more portions of the audio encounter information may include selecting the one or more predefined nulls configured to limit receiving audio encounter information from the one or more target speaker locations within the acoustic environment based upon, at least in part, the speaker metadata associated with the one or more portions of the audio encounter information.

[008] In another implementation, a computing system includes a processor and memory is configured to perform operations including but not limited to, receiving a plurality of predefined beams associated with a microphone array. The processor may be further configured to receive a plurality of predefined nulls associated with the microphone array. The processor may be further configured to select at least one of: one or more predefined beams from the plurality of predefined beams, thus defining one or more selected beams and one or more predefined nulls from the plurality of predefined nulls, thus defining one or more selected nulls. The processor may be further configured to obtain audio encounter information, via the microphone array, using at least one of the one or more selected beams and the one or more selected nulls.

[009] One or more of the following features may be included. The plurality of predefined beams may include one or more predefined beams configured to receive audio encounter information from one or more target speaker locations within an acoustic

environment. The plurality of predefined nulls may include one or more predefined nulls configured to limit receiving audio encounter information from the one or more target speaker locations within the acoustic environment. Speaker metadata associated with one or more portions of audio encounter information may be received. Selecting at least one of the one or more beams and the one or more nulls may include selecting at least one of the one or more beams and the one or more nulls based upon, at least in part, the speaker metadata associated with the one or more portions of the audio encounter information. Selecting at least one of the one or more beams and the one or more nulls based upon, at least in part, the speaker metadata associated with the one or more portions of the audio encounter information may include selecting the one or more predefined beams configured to receive audio encounter information from the one or more target speaker locations within the acoustic environment based upon, at least in part, the speaker metadata associated with the one or more portions of the audio encounter information. Selecting at least one of the one or more beams and the one or more nulls based upon, at least in part, the speaker metadata associated with the one or more portions of the audio encounter information may include selecting the one or more predefined nulls configured to limit receiving audio encounter information from the one or more target speaker locations within the acoustic environment based upon, at least in part, the speaker metadata associated with the one or more portions of the audio encounter information.

[0010] The details of one or more implementations are set forth in the accompanying drawings and the description below. Other features and advantages will become apparent from the description, the drawings, and the claims.

Brief Description of the Drawings

[0011] FIG. 1 is a diagrammatic view of an automated clinical documentation computer system and an automated clinical documentation process coupled to a distributed computing network;

[0012] FIG. 2 is a diagrammatic view of a modular ACD system incorporating the automated clinical documentation computer system of FIG. 1;

[0013] FIG. 3 is a diagrammatic view of a mixed-media ACD device included within the modular ACD system of FIG. 2;

[0014] FIG. 4 is a diagrammatic view of various modules included within a ACD computer device within the modular ACD system of FIG. 2;

[0015] FIG. 5 is a flow chart of one implementation of the automated clinical documentation process of FIG. 1;

[0016] FIGS. 6-7 are diagrammatic views of a modular ACD system according to various implementations of the automated clinical documentation process of FIG. 1;

[0017] FIG. 8 is a flow chart of one implementation of the automated clinical documentation process of FIG. 1;

[0018] FIGS. 9-10 are diagrammatic views of audio encounter information received by various microphones of a microphone array according to various implementations of the automated clinical documentation process of FIG. 1;

[0019] FIG. 11 is a flow chart of one implementation of the automated clinical documentation process of FIG. 1;

[0020] FIG. 12 is a diagrammatic view of audio encounter information received by various microphones of a microphone array according to various implementations of the automated clinical documentation process of FIG. 1;

[0021] FIG. 13 is a diagrammatic view of a plurality of speaker representations generated according to one implementation of the automated clinical documentation process of FIG. 1;

[0022] FIG. 14 is a diagrammatic view of speaker metadata generated according to various implementations of the automated clinical documentation process of FIG. 1;

[0023] FIG. 15 is a flow chart of one implementation of the automated clinical

documentation process of FIG. 1;

[0024] FIG. 16 is a diagrammatic view of a modular ACD system according to one implementation of the automated clinical documentation process of FIG. 1;

[0025] FIG. 17 is a flow chart of one implementation of the automated clinical documentation process of FIG. 1;

[0026] FIGS. 18-19 are diagrammatic views of the alignment of audio encounter information received by various microphone systems according to various implementations of the automated clinical documentation process of FIG. 1; and

[0027] FIGS. 20A-20B are diagrammatic views of the processing of audio encounter information received by various microphone systems using various speech processing system configurations according to various implementations of the automated clinical documentation process of FIG. 1.

[0028] Like reference symbols in the various drawings indicate like elements.

Detailed Description of the Preferred Embodiments

System Overview:

[0029] Referring to FIG. 1, there is shown automated clinical documentation process 10. As will be discussed below in greater detail, automated clinical documentation process 10 may be configured to automate the collection and processing of clinical encounter information to generate / store / distribute medical records.

[0030] Automated clinical documentation process 10 may be implemented as a server-side process, a client-side process, or a hybrid server-side / client-side process. For example, automated clinical documentation process 10 may be implemented as a purely server-side process via automated clinical documentation process 10s. Alternatively, automated clinical documentation process 10 may be implemented as a purely client-side process via one or more of automated clinical documentation process 10c1, automated

clinical documentation process 10c2, automated clinical documentation process 10c3, and automated clinical documentation process 10c4. Alternatively still, automated clinical documentation process 10 may be implemented as a hybrid server-side / client-side process via automated clinical documentation process 10s in combination with one or more of automated clinical documentation process 10c1, automated clinical documentation process 10c2, automated clinical documentation process 10c3, and automated clinical documentation process 10c4.

[0031] Accordingly, automated clinical documentation process 10 as used in this disclosure may include any combination of automated clinical documentation process 10s, automated clinical documentation process 10c1, automated clinical documentation process 10c2, automated clinical documentation process 10c3, and automated clinical documentation process 10c4.

[0032] Automated clinical documentation process 10s may be a server application and may reside on and may be executed by automated clinical documentation (ACD) computer system 12, which may be connected to network 14 (e.g., the Internet or a local area network). ACD computer system 12 may include various components, examples of which may include but are not limited to: a personal computer, a server computer, a series of server computers, a mini computer, a mainframe computer, one or more Network Attached Storage (NAS) systems, one or more Storage Area Network (SAN) systems, one or more Platform as a Service (PaaS) systems, one or more Infrastructure as a Service (IaaS) systems, one or more Software as a Service (SaaS) systems, a cloud-based computational system, and a cloud-based storage platform.

[0033] As is known in the art, a SAN may include one or more of a personal computer, a server computer, a series of server computers, a mini computer, a mainframe computer, a RAID device and a NAS system. The various components of ACD computer system 12 may execute one or more operating systems, examples of which may include but are not

limited to: Microsoft Windows Server [™]; Redhat Linux [™], Unix, or a custom operating system, for example.

[0034] The instruction sets and subroutines of automated clinical documentation process 10s, which may be stored on storage device 16 coupled to ACD computer system 12, may be executed by one or more processors (not shown) and one or more memory architectures (not shown) included within ACD computer system 12. Examples of storage device 16 may include but are not limited to: a hard disk drive; a RAID device; a random access memory (RAM); a read-only memory (ROM); and all forms of flash memory storage devices.

[0035] Network 14 may be connected to one or more secondary networks (e.g., network 18), examples of which may include but are not limited to: a local area network; a wide area network; or an intranet, for example.

[0036] Various IO requests (e.g. IO request 20) may be sent from automated clinical documentation process 10s, automated clinical documentation process 10c1, automated clinical documentation process 10c2, automated clinical documentation process 10c3 and/or automated clinical documentation process 10c4 to ACD computer system 12. Examples of IO request 20 may include but are not limited to data write requests (i.e. a request that content be written to ACD computer system 12) and data read requests (i.e. a request that content be read from ACD computer system 12).

[0037] The instruction sets and subroutines of automated clinical documentation process 10c1, automated clinical documentation process 10c2, automated clinical documentation process 10c3 and/or automated clinical documentation process 10c4, which may be stored on storage devices 20, 22, 24, 26 (respectively) coupled to ACD client electronic devices 28, 30, 32, 34 (respectively), may be executed by one or more processors (not shown) and one or more memory architectures (not shown) incorporated into ACD client electronic devices 28, 30, 32, 34 (respectively). Storage devices 20, 22, 24, 26 may

include but are not limited to: hard disk drives; optical drives; RAID devices; random access memories (RAM); read-only memories (ROM), and all forms of flash memory storage devices. Examples of ACD client electronic devices 28, 30, 32, 34 may include, but are not limited to, personal computing device 28 (e.g., a smart phone, a personal digital assistant, a laptop computer, a notebook computer, and a desktop computer), audio input device 30 (e.g., a handheld microphone, a lapel microphone, an embedded microphone (such as those embedded within eyeglasses, smart phones, tablet computers and/or watches) and an audio recording device), display device 32 (e.g., a tablet computer, a computer monitor, and a smart television), machine vision input device 34 (e.g., an RGB imaging system, an infrared imaging system, an ultraviolet imaging system, a laser imaging system, a SONAR imaging system, a RADAR imaging system, and a thermal imaging system), a hybrid device (e.g., a single device that includes the functionality of one or more of the above-references devices; not shown), an audio rendering device (e.g., a speaker system, a headphone system, or an earbud system; not shown), various medical devices (e.g., medical imaging equipment, heart monitoring machines, body weight scales, body temperature thermometers, and blood pressure machines; not shown), and a dedicated network device (not shown).

[0038] Users 36, 38, 40, 42 may access ACD computer system 12 directly through network 14 or through secondary network 18. Further, ACD computer system 12 may be connected to network 14 through secondary network 18, as illustrated with link line 44.

[0039] The various ACD client electronic devices (e.g., ACD client electronic devices 28, 30, 32, 34) may be directly or indirectly coupled to network 14 (or network 18). For example, personal computing device 28 is shown directly coupled to network 14 via a hardwired network connection. Further, machine vision input device 34 is shown directly coupled to network 18 via a hardwired network connection. Audio input device 30 is shown wirelessly coupled to network 14 via wireless communication channel 46

established between audio input device 30 and wireless access point (i.e., WAP) 48, which is shown directly coupled to network 14. WAP 48 may be, for example, an IEEE 802.11a, 802.11b, 802.11g, 802.11n, Wi-Fi, and/or Bluetooth device that is capable of establishing wireless communication channel 46 between audio input device 30 and WAP 48. Display device 32 is shown wirelessly coupled to network 14 via wireless communication channel 50 established between display device 32 and WAP 52, which is shown directly coupled to network 14.

[0040] The various ACD client electronic devices (e.g., ACD client electronic devices 28, 30, 32, 34) may each execute an operating system, examples of which may include but are not limited to Microsoft Windows™, Apple Macintosh™, Redhat Linux™, or a custom operating system, wherein the combination of the various ACD client electronic devices (e.g., ACD client electronic devices 28, 30, 32, 34) and ACD computer system 12 may form modular ACD system 54.

[0041] Referring also to FIG. 2, there is shown a simplified example embodiment of modular ACD system 54 that is configured to automate clinical documentation. Modular ACD system 54 may include: machine vision system 100 configured to obtain machine vision encounter information 102 concerning a patient encounter; audio recording system 104 configured to obtain audio encounter information 106 concerning the patient encounter; and a computer system (e.g., ACD computer system 12) configured to receive machine vision encounter information 102 and audio encounter information 106 from machine vision system 100 and audio recording system 104 (respectively). Modular ACD system 54 may also include: display rendering system 108 configured to render visual information 110; and audio rendering system 112 configured to render audio information 114, wherein ACD computer system 12 may be configured to provide visual information 110 and audio information 114 to display rendering system 108 and audio rendering system 112 (respectively).

[0042] Example of machine vision system 100 may include but are not limited to: one or more ACD client electronic devices (e.g., ACD client electronic device 34, examples of which may include but are not limited to an RGB imaging system, an infrared imaging system, a ultraviolet imaging system, a laser imaging system, a SONAR imaging system, a RADAR imaging system, and a thermal imaging system). Examples of audio recording system 104 may include but are not limited to: one or more ACD client electronic devices (e.g., ACD client electronic device 30, examples of which may include but are not limited to a handheld microphone, a lapel microphone, an embedded microphone (such as those embedded within eyeglasses, smart phones, tablet computers and/or watches) and an audio recording device). Examples of display rendering system 108 may include but are not limited to: one or more ACD client electronic devices (e.g., ACD client electronic device 32, examples of which may include but are not limited to a tablet computer, a computer monitor, and a smart television). Examples of audio rendering system 112 may include but are not limited to: one or more ACD client electronic devices (e.g., audio rendering device 116, examples of which may include but are not limited to a speaker system, a headphone system, and an earbud system).

[0043] As will be discussed below in greater detail, ACD computer system 12 may be configured to access one or more datasources 118 (e.g., plurality of individual datasources 120, 122, 124, 126, 128), examples of which may include but are not limited to one or more of a user profile datasource, a voice print datasource, a voice characteristics datasource (e.g., for adapting the automated speech recognition models), a face print datasource, a humanoid shape datasource, an utterance identifier datasource, a wearable token identifier datasource, an interaction identifier datasource, a medical conditions symptoms datasource, a prescriptions compatibility datasource, a medical insurance coverage datasource, and a home healthcare datasource. While in this particular example, five different examples of datasources 118, are shown, this is for illustrative purposes only and

is not intended to be a limitation of this disclosure, as other configurations are possible and are considered to be within the scope of this disclosure.

[0044] As will be discussed below in greater detail, modular ACD system 54 may be configured to monitor a monitored space (e.g., monitored space 130) in a clinical environment, wherein examples of this clinical environment may include but are not limited to: a doctor's office, a medical facility, a medical practice, a medical lab, an urgent care facility, a medical clinic, an emergency room, an operating room, a hospital, a long term care facility, a rehabilitation facility, a nursing home, and a hospice facility. Accordingly, an example of the above-referenced patient encounter may include but is not limited to a patient visiting one or more of the above-described clinical environments (e.g., a doctor's office, a medical facility, a medical practice, a medical lab, an urgent care facility, a medical clinic, an emergency room, an operating room, a hospital, a long term care facility, a rehabilitation facility, a nursing home, and a hospice facility).

[0045] Machine vision system 100 may include a plurality of discrete machine vision systems when the above-described clinical environment is larger or a higher level of resolution is desired. As discussed above, examples of machine vision system 100 may include but are not limited to: one or more ACD client electronic devices (e.g., ACD client electronic device 34, examples of which may include but are not limited to an RGB imaging system, an infrared imaging system, an ultraviolet imaging system, a laser imaging system, a SONAR imaging system, a RADAR imaging system, and a thermal imaging system). Accordingly, machine vision system 100 may include one or more of each of an RGB imaging system, an infrared imaging systems, an ultraviolet imaging systems, a laser imaging system, a SONAR imaging system, a RADAR imaging system, and a thermal imaging system.

[0046] Audio recording system 104 may include a plurality of discrete audio recording systems when the above-described clinical environment is larger or a higher level of

resolution is desired. As discussed above, examples of audio recording system 104 may include but are not limited to: one or more ACD client electronic devices (e.g., ACD client electronic device 30, examples of which may include but are not limited to a handheld microphone, a lapel microphone, an embedded microphone (such as those embedded within eyeglasses, smart phones, tablet computers and/or watches) and an audio recording device). Accordingly, audio recording system 104 may include one or more of each of a handheld microphone, a lapel microphone, an embedded microphone (such as those embedded within eyeglasses, smart phones, tablet computers and/or watches) and an audio recording device.

[0047] Display rendering system 108 may include a plurality of discrete display rendering systems when the above-described clinical environment is larger or a higher level of resolution is desired. As discussed above, examples of display rendering system 108 may include but are not limited to: one or more ACD client electronic devices (e.g., ACD client electronic device 32, examples of which may include but are not limited to a tablet computer, a computer monitor, and a smart television). Accordingly, display rendering system 108 may include one or more of each of a tablet computer, a computer monitor, and a smart television.

[0048] Audio rendering system 112 may include a plurality of discrete audio rendering systems when the above-described clinical environment is larger or a higher level of resolution is desired. As discussed above, examples of audio rendering system 112 may include but are not limited to: one or more ACD client electronic devices (e.g., audio rendering device 116, examples of which may include but are not limited to a speaker system, a headphone system, or an earbud system). Accordingly, audio rendering system 112 may include one or more of each of a speaker system, a headphone system, or an earbud system.

[0049] ACD computer system 12 may include a plurality of discrete computer systems.

As discussed above, ACD computer system 12 may include various components, examples of which may include but are not limited to: a personal computer, a server computer, a series of server computers, a mini computer, a mainframe computer, one or more Network Attached Storage (NAS) systems, one or more Storage Area Network (SAN) systems, one or more Platform as a Service (PaaS) systems, one or more Infrastructure as a Service (IaaS) systems, one or more Software as a Service (SaaS) systems, a cloud-based computational system, and a cloud-based storage platform. Accordingly, ACD computer system 12 may include one or more of each of a personal computer, a server computer, a series of server computers, a mini computer, a mainframe computer, one or more Network Attached Storage (NAS) systems, one or more Storage Area Network (SAN) systems, one or more Platform as a Service (PaaS) systems, one or more Infrastructure as a Service (IaaS) systems, one or more Software as a Service (SaaS) systems, a cloud-based computational system, and a cloud-based storage platform.

[0050] Referring also to FIG. 3, audio recording system 104 may include directional microphone array 200 having a plurality of discrete microphone assemblies. For example, audio recording system 104 may include a plurality of discrete audio acquisition devices (e.g., audio acquisition devices 202, 204, 206, 208, 210, 212, 214, 216, 218) that may form microphone array 200. As will be discussed below in greater detail, modular ACD system 54 may be configured to form one or more audio recording beams (e.g., audio recording beams 220, 222, 224) via the discrete audio acquisition devices (e.g., audio acquisition devices 202, 204, 206, 208, 210, 212, 214, 216, 218) included within audio recording system 104.

[0051] For example, modular ACD system 54 may be further configured to steer the one or more audio recording beams (e.g., audio recording beams 220, 222, 224) toward one or more encounter participants (e.g., encounter participants 226, 228, 230) of the above-described patient encounter. Examples of the encounter participants (e.g., encounter

participants 226, 228, 230) may include but are not limited to: medical professionals (e.g., doctors, nurses, physician's assistants, lab technicians, physical therapists, scribes (e.g., a transcriptionist) and/or staff members involved in the patient encounter), patients (e.g., people that are visiting the above-described clinical environments for the patient encounter), and third parties (e.g., friends of the patient, relatives of the patient and/or acquaintances of the patient that are involved in the patient encounter).

[0052] Accordingly, modular ACD system 54 and/or audio recording system 104 may be configured to utilize one or more of the discrete audio acquisition devices (e.g., audio acquisition devices 202, 204, 206, 208, 210, 212, 214, 216, 218) to form an audio recording beam. For example, modular ACD system 54 and/or audio recording system 104 may be configured to utilize audio acquisition device 210 to form audio recording beam 220, thus enabling the capturing of audio (e.g., speech) produced by encounter participant 226 (as audio acquisition device 210 is pointed to (i.e., directed toward) encounter participant 226). Additionally, modular ACD system 54 and/or audio recording system 104 may be configured to utilize audio acquisition devices 204, 206 to form audio recording beam 222, thus enabling the capturing of audio (e.g., speech) produced by encounter participant 228 (as audio acquisition devices 204, 206 are pointed to (i.e., directed toward) encounter participant 228). Additionally, modular ACD system 54 and/or audio recording system 104 may be configured to utilize audio acquisition devices 212, 214 to form audio recording beam 224, thus enabling the capturing of audio (e.g., speech) produced by encounter participant 230 (as audio acquisition devices 212, 214 are pointed to (i.e., directed toward) encounter participant 230). Further, modular ACD system 54 and/or audio recording system 104 may be configured to utilize null-steering precoding to cancel interference between speakers and/or noise.

[0053] As is known in the art, null-steering precoding is a method of spatial signal processing by which a multiple antenna transmitter may null multiuser interference signals

in wireless communications, wherein null-steering precoding may mitigate the impact of background noise and unknown user interference.

[0054] In particular, null-steering precoding may be a method of beamforming for narrowband signals that may compensate for delays of receiving signals from a specific source at different elements of an antenna array. In general and to improve performance of the antenna array, incoming signals may be summed and averaged, wherein certain signals may be weighted and compensation may be made for signal delays.

[0055] Machine vision system 100 and audio recording system 104 may be stand-alone devices (as shown in FIG. 2). Additionally / alternatively, machine vision system 100 and audio recording system 104 may be combined into one package to form mixed-media ACD device 232. For example, mixed-media ACD device 232 may be configured to be mounted to a structure (e.g., a wall, a ceiling, a beam, a column) within the above-described clinical environments (e.g., a doctor's office, a medical facility, a medical practice, a medical lab, an urgent care facility, a medical clinic, an emergency room, an operating room, a hospital, a long term care facility, a rehabilitation facility, a nursing home, and a hospice facility), thus allowing for easy installation of the same. Further, modular ACD system 54 may be configured to include a plurality of mixed-media ACD devices (e.g., mixed-media ACD device 232) when the above-described clinical environment is larger or a higher level of resolution is desired.

[0056] Modular ACD system 54 may be further configured to steer the one or more audio recording beams (e.g., audio recording beams 220, 222, 224) toward one or more encounter participants (e.g., encounter participants 226, 228, 230) of the patient encounter based, at least in part, upon machine vision encounter information 102. As discussed above, mixed-media ACD device 232 (and machine vision system 100 / audio recording system 104 included therein) may be configured to monitor one or more encounter participants (e.g., encounter participants 226, 228, 230) of a patient encounter.

[0057] Specifically, machine vision system 100 (either as a stand-alone system or as a component of mixed-media ACD device 232) may be configured to detect humanoid shapes within the above-described clinical environments (e.g., a doctor's office, a medical facility, a medical practice, a medical lab, an urgent care facility, a medical clinic, an emergency room, an operating room, a hospital, a long term care facility, a rehabilitation facility, a nursing home, and a hospice facility). And when these humanoid shapes are detected by machine vision system 100, modular ACD system 54 and/or audio recording system 104 may be configured to utilize one or more of the discrete audio acquisition devices (e.g., audio acquisition devices 202, 204, 206, 208, 210, 212, 214, 216, 218) to form an audio recording beam (e.g., audio recording beams 220, 222, 224) that is directed toward each of the detected humanoid shapes (e.g., encounter participants 226, 228, 230).

[0058] As discussed above, ACD computer system 12 may be configured to receive machine vision encounter information 102 and audio encounter information 106 from machine vision system 100 and audio recording system 104 (respectively); and may be configured to provide visual information 110 and audio information 114 to display rendering system 108 and audio rendering system 112 (respectively). Depending upon the manner in which modular ACD system 54 (and/or mixed-media ACD device 232) is configured, ACD computer system 12 may be included within mixed-media ACD device 232 or external to mixed-media ACD device 232.

[0059] As discussed above, ACD computer system 12 may execute all or a portion of automated clinical documentation process 10, wherein the instruction sets and subroutines of automated clinical documentation process 10 (which may be stored on one or more of e.g., storage devices 16, 20, 22, 24, 26) may be executed by ACD computer system 12 and/or one or more of ACD client electronic devices 28, 30, 32, 34.

The Automated Clinical Documentation Process:

[0060] In some implementations consistent with the present disclosure, systems and

methods may be provided for speaker diarization and distant speech recognition using spatial and spectral information. For example, distant conversation automated speech recognition (DCASR) may include multiple microphone systems configured to record and recognize speech of one or more speakers (e.g., a microphone array and a single mobile microphone carried by a speaker). Conventional approaches to DCASR are subject to various challenges when utilizing signals from multiple microphones. For example, conventional beamforming techniques combine the multiple microphone signals via spatial filtering, but these are mostly limited to single microphone array devices (i.e. they typically do not provide a mechanism for combining various microphones not part of an array, nor different microphone arrays deployed in a room operating as independent devices). Conventional techniques also allow for single channel noise reduction and de-reverberation techniques. However, such techniques do not account for the spatial information that maybe available. Additionally, conventional approaches to DCASR are unable to manage the combination of enhanced signals from various devices in the room. In some implementations, conventional approaches to DCASR may encounter issues when locating where the speakers are at any given time, tracking their movements, and then identifying their identify. This may result in less accurately diarized text transcriptions (i.e., text transcriptions with speaker labels to define who said what and when in a conversation). In addition, an ASR system, that converts the audio to text, may be unable to process the various metadata and audio signals from multiple sources (e.g., multiple microphone systems).

[0061] Accordingly, implementations of the present disclosure may address these challenges experienced by conventional approaches to DCASR by predefining beamforming configurations based upon, at least in part, information associated with an acoustic environment; selecting particular beam patterns and/or null patterns for particular locations (e.g., where beams and nulls are selected to point to certain locations in a room,

some which may have a high probability of being occupied by a particular speaker); voice activity detection and localization using spatial and spectral information; speaker identification with spatial and spectral information; and alignment of audio streams from multiple microphone systems based upon, at least in part, voice activity detection. For example, automated clinical documentation process 10 may include various hardware and/or software modules configured to perform various DCASR functions within the scope of the present disclosure.

[0062] Referring also to FIG. 4 and in some implementations, machine vision system 100 and audio recording system 104 may be combined into one package to form mixed-media ACD device 232. In this example, ACD computer system 12 may be configured to receive machine vision encounter information 102 and audio encounter information 106 from machine vision system 100 and audio recording system 104 (respectively). Specifically, ACD computer system 12 may receive the audio encounter information at a voice activity detection (VAD) module (e.g., VAD module 400); at an acoustic beamforming module (e.g., acoustic beamforming module 402); and/or at a speaker identification and tracking module (e.g., speaker identification and tracking module 404). In some implementations, ACD computer system 12 may also receive machine vision encounter information 102 at speaker identification and tracking module 404. In some implementations, VAD module 400 may communicate with sound localization module 406 to output metadata for speaker identification and tracking module 404. In some implementations, acoustic beamforming module 402 may be configured to provide a plurality of beams and/or a plurality of nulls to a beam and null selection module (e.g., beam and null selection module 408). Speaker identification and tracking module 404 may be configured to provide speaker identification and tracking information to beam and null selection module 408.

[0063] In some implementations, beam and null selection module 408 may be

configured to provide an aligned audio stream from audio recording system 104 to a device selection and weighting module (e.g., device selection and weighting module 410). In addition, beam and null selection module 408 may be configured to provide the aligned audio stream to an alignment module (e.g., alignment module 412). Alignment module 412 may be configured to receive audio encounter information from a VAD module associated with a second microphone system (e.g., VAD module 414 associated with mobile electronic device 416). In some implementations, alignment module 412 may be configured to provide an aligned audio stream from mobile electronic device 416 to device selection and weighting module 410. In some implementations, device selection and weighting module 410 may provide the aligned audio stream from audio recording system 104, the aligned audio stream from mobile electronic device 416, and metadata from the other modules to one or more speech processing systems (e.g., speech processing system 418). In some implementations, speech processing system 418 may be an automated speech recognition (ASR) system.

[0064] As will be discussed in greater detail below, the combination of the various modules of ACD computer system 12 / automated clinical documentation process 10 may be configured to improve DCASR by providing aligned audio encounter information from multiple microphone systems. In this manner, automated clinical documentation process 10 may improve distant speech recognition using spatial and spectral information determined by the various modules of ACD computer system 12 / automated clinical documentation process 10.

[0065] Referring also at least to FIGS. 5-7, automated clinical documentation process 10 may receive 500 information associated with an acoustic environment. A plurality of filters may be predefined 502 to produce a plurality of beams based upon, at least in part, the information associated with the acoustic environment. The plurality of filters may be predefined 504 to produce a plurality of nulls based upon, at least in part, the information

associated with the acoustic environment. Audio encounter information may be obtained 506, via one or more microphone arrays, using the plurality of beams and the plurality of nulls produced by the plurality of predefined filters.

[0066] As will be discussed in greater detail below, automated clinical documentation process 10 may predefine or pre-compute a plurality of filters to produce beams and nulls for targeting specific speakers and/or noise sources to improve DCASR. As is known in the art, beamformers may use a set of filters associated with a set of input signals to generate a spatially filtered sensitivity pattern formed from beams and nulls with characteristics determined by the filters. As will be discussed in greater detail below, automated clinical documentation process 10 may receive 500 information associated with an acoustic environment and may predefine filters to produce beams and/or nulls to target or isolate specific speakers and/or noise sources. In this manner, automated clinical documentation process 10 may utilize acoustic environment information to predefine beams and nulls that may be selected for various situations (i.e., selecting certain beams and nulls when a patient is speaking to a medical professional and selecting other beams and nulls when the medical professional is speaking to the patient).

[0067] In some implementations, automated clinical documentation process 10 may receive 500 information associated with an acoustic environment. An acoustic environment may represent the layout and acoustic properties of a room or other space where multiple microphone systems may be deployed. For example, information associated with the acoustic environment may describe the type / dimensions of a room, the activity zones in the room where the participants are likely to be or operate, the acoustic properties of the room (e.g., the types of noise expected, range of reverberation, etc.), position of the microphone system(s) within the acoustic environment, etc. In some implementations, automated clinical documentation process 10 may provide a user interface to receive information associated with the acoustic environment. Accordingly, a

user and/or automated clinical documentation process 10 may provide the information associated with an acoustic environment (e.g., via the user interface). However, it will be appreciated that the information associated with the acoustic environment may be received in various ways (e.g., default information for a default acoustic environment, automatically defined by automated clinical documentation process 10, etc.).

[0068] In some implementations, the information associated with the acoustic environment may indicate one or more target speaker locations within the acoustic environment. Referring also to FIG. 6 and in some implementations, suppose that acoustic environment 600 includes one or more encounter participants (e.g., encounter participants 226, 228, 230) of the above-described patient encounter. As described above, examples of the encounter participants (e.g., encounter participants 226, 228, 230) may include but are not limited to: medical professionals (e.g., doctors, nurses, physician's assistants, lab technicians, physical therapists, scribes (e.g., a transcriptionist) and/or staff members involved in the patient encounter), patients (e.g., people that are visiting the above-described clinical environments for the patient encounter), and third parties (e.g., friends of the patient, relatives of the patient and/or acquaintances of the patient that are involved in the patient encounter).

[0069] In some implementations, automated clinical documentation process 10 may receive 500 information associated with acoustic environment 600 that indicates the location of encounter participants 226, 228, 230. In some implementations, the information associated with acoustic environment 600 may indicate locations within acoustic environment where the one or more target speakers are likely to be when speaking. For example, suppose automated clinical documentation process 10 determines that an examination table is positioned at approximately e.g., 45° from the base of microphone array 200. In this example, automated clinical documentation process 10 may determine that a patient is most likely to speak while sitting on or near the examination table at

approximately 45° from the base of microphone array 200. Further suppose that automated clinical documentation process 10 determines that a doctor's desk is positioned at approximately e.g., 90° from the base of microphone array 200. In this example, automated clinical documentation process 10 may determine that a doctor is most likely to speak while sitting on or near the desk at approximately 90° from the base of microphone array 200. Additionally, suppose automated clinical documentation process 10 determines that a waiting area is positioned at approximately e.g., 120° from the base of microphone array 200. In this example, automated clinical documentation process 10 may determine that other patients or other third parties are most likely to speak from at approximately 120° from the base of microphone array 200. While three examples of relative locations for target speakers within an acoustic environment have been provided, it will be appreciated that the information associated with the acoustic environment may include any number of target speaker locations or probability-based target speaker locations for any number of target speakers within the scope of the present disclosure.

[0070] In some implementations, automated clinical documentation process 10 may predefine 502 a plurality of filters to produce a plurality of beams based upon, at least in part, the information associated with the acoustic environment. A beam may generally include a pattern of constructive interference among microphones of a microphone array that is generated by modifying the phase and/or amplitude of the signal at each microphone of the microphone array via a plurality of filters. The pattern of constructive interference may improve the signal processing performance of the microphone array. Predefining 502 a plurality of filters to produce a plurality of beams may generally include defining the plurality of beams at any point in time prior to obtaining 506 audio encounter information via the one or more microphone arrays. The plurality of filters may include a plurality of finite impulse response (FIR) filters configured to adjust the phase and/or amplitude of a signal at each microphone of a microphone array. Accordingly, automated clinical

documentation process 10 may predefine the plurality of filters (e.g., the plurality of FIR filters) to produce a plurality of beams that are configured to “look” or orient in a particular direction based upon, at least in part, the information associated with the acoustic environment. In some implementations, automated clinical documentation process 10 may predefine 502 the plurality of filters to produce a plurality of beams by adjusting the phase and/or amplitude of the signal at each microphone of the microphone array.

[0071] For instance, automated clinical documentation process 10 may receive 500 information associated with acoustic environment and may determine how the acoustic properties of the acoustic environment impact beams deployed within the acoustic environment. For example, suppose that the information associated with the acoustic environment indicates that a particular reverberation level may be present at particular frequencies and at varying amplitudes at different parts of the acoustic environment. In this example, automated clinical documentation process 10 may predefine 502 the plurality of filters to produce a plurality of beams to account for the layout, acoustic properties, etc. of the acoustic environment. As will be discussed in greater detail below, by predefining 502 the plurality of filters to produce a plurality of beams, automated clinical documentation process 10 may allow for the selection of particular beams for various situations.

[0072] In some implementations, predefining 502 the plurality of filters to produce the plurality of beams based upon, at least in part, the information associated with the acoustic environment may include predefining 508 the plurality of filters to produce one or more beams configured to receive audio encounter information from the one or more target speaker locations within the acoustic environment. For example and referring again to FIG. 6, automated clinical documentation process 10 may predefine 508 the plurality of filters to produce one or more beams configured to receive audio encounter information from the one or more target speaker locations within the acoustic environment. Continuing

with the above example, suppose that automated clinical documentation process 10 determines from the information associated with acoustic environment 600, that a patient is most likely to speak while sitting on or near the examination table at approximately 45° from the base of microphone array 200; a doctor is most likely to speak while sitting on or near the desk at approximately 90° from the base of microphone array 200; and that other patients or other third parties are most likely to speak from at approximately 120° from the base of microphone array 200. In this example, automated clinical documentation process 10 may predefine 508 the plurality of filters to produce beam 220 for receiving audio encounter information from a patient (e.g., encounter participant 228); beam 222 for receiving audio encounter information from a doctor (e.g., encounter participant 226); and beam 224 for receiving audio encounter information from another patient / third party (e.g., encounter participant 230). As will be discussed in greater detail below, automated clinical documentation process 10 may select particular beams for receiving audio encounter information when a specific participant is speaking.

[0073] In some implementations, predefining 502 the plurality of filters to produce a plurality of beams based upon, at least in part, the information associated with the acoustic environment may include predefining 510 the plurality of filters to produce a plurality of frequency-independent beams based upon, at least in part, the information associated with the acoustic environment. In some implementations, the spatial sensitivity to audio encounter information of a beam may be frequency-dependent. For example, for receiving high frequency signals, the beam of sensitivity may be narrow while the beam of sensitivity for low frequency signals may be wide. Accordingly, automated clinical documentation process 10 may predefine 510 the plurality of filters to produce a plurality of frequency-independent beams based upon, at least in part, the information associated with the acoustic environment. For example and as discussed above, automated clinical documentation process 10 may determine particular locations within an acoustic environment where

speakers may be located when they speak based upon, at least in part, the information associated with the acoustic environment. Additionally, automated clinical documentation process 10 may receive acoustic properties of the acoustic environment. For example, automated clinical documentation process 10 may determine the location(s) and frequency characteristics of one or more noise sources within the acoustic environment. In this manner, automated clinical documentation process 10 may predefine 510 the plurality of filters to produce a plurality of frequency-independent beams that account for the acoustic properties of the acoustic environment while providing sufficient microphone sensitivity for frequency variations.

[0074] For example, automated clinical documentation process 10 may predefine 502 the plurality of filters to produce beams 220, 222, 224 by modifying the phase and/or amplitude of the signals of each microphone such that the beam is sufficiently sensitive to receive audio encounter information for target speakers within the acoustic environment regardless of the frequency. In this manner, automated clinical documentation process 10 may predefine 510 the plurality of filters to produce a plurality of frequency-independent beams (e.g., frequency-independent beams 220, 222, 224) based upon, at least in part, the information associated with the acoustic environment (e.g., acoustic environment 600).

[0075] In some implementations, automated clinical documentation process 10 may predefine 504 the plurality of filters to produce a plurality of nulls based upon, at least in part, the information associated with the acoustic environment. A null may generally include a pattern of destructive interference among microphones of a microphone array that is generated by modifying the phase and/or amplitude of the signal at each microphone of the microphone array. The pattern of destructive interference may limit the reception of signals by the microphone array. Predefining 504 the plurality of filters to produce a plurality of nulls may generally include defining the plurality of nulls at any point in time prior to obtaining 506 audio encounter information via one or more microphone arrays. As

discussed above, the plurality of filters may include a plurality of finite impulse response (FIR) filters configured to adjust the phase and/or amplitude of a signal at each microphone of a microphone array. Accordingly, automated clinical documentation process 10 may predefine the plurality of filters (e.g., the plurality of FIR filters) to produce a plurality of nulls that are configured to “look” or orient in a particular direction based upon, at least in part, the information associated with the acoustic environment. In contrast to a beam, a null may limit or attenuate reception at the target direction. In some implementations, automated clinical documentation process 10 may predefine 504 the plurality of filters to produce a plurality of nulls by adjusting the phase and/or amplitude of the signal at each microphone of the microphone array.

[0076] For instance, automated clinical documentation process 10 may receive 500 information associated with acoustic environment and may determine how the acoustic properties of the acoustic environment impact beams deployed within the acoustic environment. For example and as discussed above, suppose that the information associated with the acoustic environment indicates that a particular noise signal (e.g., sound of an air-conditioning system) may be present at particular frequencies and at varying amplitudes at different parts of the acoustic environment. In this example, automated clinical documentation process 10 may predefine 504 the plurality of filters to produce a plurality of nulls to limit the reception of noise signals. As will be discussed in greater detail below, by predefining 504 the plurality of filters to produce the plurality of nulls, automated clinical documentation process 10 may allow for the selection of particular nulls for various situations.

[0077] In some implementations, automated clinical documentation process 10 may predefine 504 the plurality of filters to produce one or more nulls configured to limit receiving noise signals from one or more noise sources based upon, at least in part, the information associated with the acoustic environment. For example and referring again to

FIG. 6, suppose that automated clinical documentation process 10 determines from the information associated with acoustic environment 600 that a noise source (e.g., fan 244) is positioned e.g., approximately 70° from the base of microphone array 200 and that a second noise source (e.g., a doorway to a busy hallway 246) is positioned e.g., approximately 150° from the base of microphone array 200. In this example, automated clinical documentation process 10 may predefine 512 the plurality of filters to produce null 248 to limit receiving noise signals from a first noise source (e.g., fan 244) and null 250 to limit receiving noise signals from a second noise source (e.g., a doorway to a busy hallway 246). While an example of predefining 504 the plurality of filters to produce two nulls for two noise sources has been described, it will be appreciated that automated clinical documentation process 10 may predefine 504 the plurality of filters to produce any number of nulls for any number of noise sources within the scope of the present disclosure. In some implementations, automated clinical documentation process 10 may select particular nulls to limit receiving noise signals for various situations.

[0078] In some implementations, predefining 504 the plurality of filters to produce the plurality of nulls based upon, at least in part, the information associated with the acoustic environment may include predefining 512 the plurality of filters to produce one or more nulls to limit receiving audio encounter information from the one or more target speaker locations within the acoustic environment. For example and referring again to FIG. 6, automated clinical documentation process 10 may predefine 512 the plurality of filters to produce one or more nulls to limit receiving audio encounter information from the one or more target speaker locations within the acoustic environment. Continuing with the above example, suppose that automated clinical documentation process 10 determines from the information associated with acoustic environment 600, that a patient is most likely to speak while sitting on or near the examination table at approximately 45° from the base of microphone array 200; a doctor is most likely to speak while sitting on or near the desk at

approximately 90° from the base of microphone array 200; and that other patients or other third parties are most likely to speak from at approximately 120° from the base of microphone array 200.

[0079] In this example and referring also to FIG. 7, automated clinical documentation process 10 may predefine 504 the plurality of filters to produce null 252 to limit receiving audio encounter information from a patient (e.g., encounter participant 228); null 254 to limit receiving audio encounter information from a doctor (e.g., encounter participant 226); and null 256 to limit receiving audio encounter information from another patient / third party (e.g., encounter participant 230). As will be discussed in greater detail below, automated clinical documentation process 10 may select particular nulls to limit audio encounter information from other speakers to focus on audio encounter information from a particular target speaker.

[0080] However, in another example, suppose a patient (e.g., encounter participant 228) and a doctor (e.g., encounter participant 226) are talking over each other for a short duration (i.e., “overlapped speech”). Automated clinical documentation process 10 may form two output signals where a first signal is the sum a beam pattern oriented toward the doctor (e.g., encounter participant 226) and a null oriented toward the patient (e.g., encounter participant 228); and a second signal is the sum a beam pattern oriented toward the patient (e.g., encounter participant 228) and a null oriented toward the doctor (e.g., encounter participant 226). In this manner, the combination of the two signals formed from the beams and nulls defined for each signal may have a better chance of accurately transcribing what the doctor and the patient said.

[0081] In some implementations, automated clinical documentation process 10 may obtain 506 audio encounter information, via one or more microphone arrays, using the plurality of beams and the plurality of nulls produced by the plurality of predefined filters. As discussed above, in the example of audio signals, one or more microphone arrays may

utilize beams and nulls to audio encounter information from particular speakers and to limit receiving audio encounter information from other speakers or other sound sources. For example, automated clinical documentation process 10 may obtain audio encounter information 106 by combining the plurality of discrete microphone elements (e.g., audio acquisition devices 202, 204, 206, 208, 210, 212, 214, 216, 218) in the microphone array (e.g., microphone array 200) with the plurality of predefined filters in such a way that signals at particular angles experience constructive interference while others experience destructive interference.

[0082] In some implementations, obtaining 506 the audio encounter information, via one or more microphone arrays, using the plurality of beams and the plurality of nulls produced by the plurality of predefined filters may include adaptively steering 514 the plurality of beams and the plurality of nulls. For example, automated clinical documentation process 10 may allow for the plurality of beams and/or the plurality of nulls to be steered toward a target speaker within the acoustic environment. Returning to the above example, suppose a patient (e.g., encounter participant 228) is speaking while moving within acoustic environment 600. Further suppose the patient (e.g., encounter participant 228) moves from approximately e.g. 45° relative to the base of microphone array 200 to approximately e.g., 35° relative to the base of microphone array 200. In this example, automated clinical documentation process 10 may adaptively steer 514 a beam (e.g., beam 220) to follow the patient (e.g., encounter participant 228) as the patient moves within acoustic environment 600. Further, suppose that while the patient (e.g., encounter participant 228) is speaking, the doctor (e.g., encounter participant 226) begins moving around within acoustic environment 600. In this example, automated clinical documentation process 10 may adaptively steer 514 a null (e.g., null 254) to follow the doctor (e.g., encounter participant 226) to limit any noise signals or audio encounter information from the doctor (e.g., encounter participant 226) while the patient (e.g.,

encounter participant 228) is speaking. While an example of adaptively steering one beam and one null has been provided, it will be appreciated that any number of beams and/or nulls may be adaptively steered for various purposes within the present disclosure.

[0083] In some implementations, obtaining 506 audio encounter information from particular speakers and to limit receiving audio encounter information from other speakers may include one or more of: selecting 516 one or more beams from the plurality of beams; and selecting 518 one or more nulls from the plurality of nulls. For example and as discussed above, automated clinical documentation process 10 may predefine the plurality of filters to produce the plurality of beams and/or the plurality of nulls based upon, at least in part, information associated with the acoustic environment. Referring again to FIG. 4, these actions may be generally performed by or associated with acoustic beamforming module 402. In this manner, automated clinical documentation process 10 may provide the predefined filters to beam and null selection module 508. As will be discussed in greater detail below, beam and null selection module 508 may be configured to select 516 one or more beams from the plurality of beams and/or select 518 one or more nulls from the plurality of nulls for various situations using the plurality of predefined filters.

[0084] For example and as discussed above, suppose the patient (e.g., encounter participant 228) begins speaking. Automated clinical documentation process 10 may detect the patient's audio encounter information and may select 516 one or more beams (e.g., beam 220) to receive audio encounter information from the patient (e.g., encounter participant 228). Additionally, automated clinical documentation process 10 may detect speech from another participant (e.g., participant 230). In one example, automated clinical documentation process 10 may select 516 one or more beams (e.g., beam 224) to receive audio encounter information from patient 230 or may select 518 one or more nulls (e.g., null 256) to limit receiving audio encounter information from participant 230. In this manner, automated clinical documentation process 10 may select, from the beams and nulls

produced by the predefined plurality of filters, which beams and nulls to utilize when audio encounter information is detected from particular speakers.

[0085] Referring also at least to FIGS. 8-10, automated clinical documentation process 10 may receive 800 audio encounter information from a microphone array. Speech activity within one or more portions of the audio encounter information may be identified 802 based upon, at least in part, a correlation among the audio encounter information received from the microphone array. Location information for the one or more portions of the audio encounter information may be determined 804 based upon, at least in part, the correlation among the signals received by each microphone of the microphone array. The one or more portions of the audio encounter information may be labeled 806 with the speech activity and the location information.

[0086] In some implementations, automated clinical documentation process 10 may receive 800 audio encounter information from a microphone array. Referring again to FIG. 4 and in some implementations, a microphone array (e.g., microphone array 200) may include a plurality of microphones (e.g., audio acquisition devices 202, 204, 206, 208, 210, 212, 214, 216, 218) configured to receive audio encounter information (e.g., audio encounter information 106). As discussed above, audio encounter information 106 may include speech signals or other signals recorded by the plurality of audio acquisition devices. In some implementations, automated clinical documentation process 10 may receive audio encounter information 106 at a voice activity detection (VAD) module (e.g., VAD module 400) and/or a sound localization module (e.g., sound localization module 406) of ACD computer system 12.

[0087] Referring also to FIG. 9, suppose automated clinical documentation process 10 receives audio encounter information 106 at a plurality of microphones (e.g., audio acquisition devices 202, 204, 206, 208) of microphone array 200. While only four audio acquisition devices are shown, it will be appreciated that this is for ease of explanation and

that any number of audio acquisition devices of a microphone array may record audio encounter information within the scope of the present disclosure. In this example, suppose each microphone receives an audio signal (e.g., audio encounter information 106). Based on the orientation of each microphone, the properties of each microphone, the properties of the audio signals, etc., each microphone may receive a different version of the signal. For example, audio encounter information received by microphone 202 may have distinct signal components (i.e., amplitude and phase) relative to the audio encounter information received by microphones 204, 206, and/or 208. As will be discussed in greater detail below, automated clinical documentation process 10 may utilize correlation between the audio encounter information to detect speech and to determine location information associated with audio encounter information.

[0088] In some implementations, automated clinical documentation process 10 may identify 802 speech activity within one or more portions of the audio encounter information based upon, at least in part, a correlation among the audio encounter information received by the microphone array. Referring again to the example of FIG. 9 and in some implementations, audio encounter information 106 may include a plurality of portions or frames of audio information (e.g., portions 900, 902, 904, 906, 908, 910, 912, 914, 916, 918, 920, 922, 924, 926). While an example with e.g., 14 portions of audio encounter information 106 have been described, it will be appreciated that this is for example purposes only and that audio encounter information 106 may include or be defined as any number of portions within the scope of the present disclosure. In one example, each portion or frame may represent audio encounter information 106 over a predefined amount of time (e.g., 20 milliseconds).

[0089] In some implementations, automated clinical documentation process 10 may determine a correlation among the audio encounter information received by microphone array 200. For example, automated clinical documentation process 10 may compare the

one or more portions (e.g., portions 900, 902, 904, 906, 908, 910, 912, 914, 916, 918, 920, 922, 924, 926) of audio encounter information 106 to determine a degree of correlation between the audio encounter information present in each portion across the plurality of microphones of the microphone array. In some implementations, automated clinical documentation process 10 may perform various cross-correlation processes known in the art to determine a degree of similarity between the one or more portions (e.g., portions 900, 902, 904, 906, 908, 910, 912, 914, 916, 918, 920, 922, 924, 926) of audio encounter information 106 across the plurality of microphones of microphone array 200.

[0090] For example, suppose automated clinical documentation process 10 receives just ambient noise (i.e., no speech and no directional noise sources). Automated clinical documentation process 10 may determine that the spectrum observed in each microphone channel is different (i.e. uncorrelated at each microphone). However, suppose automated clinical documentation process 10 receives a speech or other ‘directional’ signal within the audio encounter information. In this example, automated clinical documentation process 10 may determine that one or more portions of the audio encounter information (e.g., the portions of audio encounter information with speech components) are highly correlated at each microphone in the microphone array.

[0091] In some implementations, automated clinical documentation process 10 may identify 802 speech activity within one or more portions of the audio encounter information based upon, at least in part, determining a threshold amount or degree of correlation among the audio encounter information received from the microphone array. For example and as discussed above, various thresholds may be defined (e.g., user-defined, default thresholds, automatically defined via automated clinical documentation process 10, etc.) to determine when portions of audio encounter information are sufficiently correlated. Accordingly, in response to determining at least a threshold degree of correlation among the portions of audio encounter information across the plurality of microphones, automated clinical

documentation process 10 may determine or identify speech activity within the one or more portions of the audio encounter information. In combination with determining a threshold correlation among the portions of audio encounter information, automated clinical documentation process 10 may identify speech activity using other approaches known in the art for voice activity detection (VAD) such as filtering, noise reduction, the application of classification rules, etc. In this manner, conventional VAD techniques may be used in combination with the determination of a threshold correlation among the one or more portions of the audio encounter information to identify 802 speech activity within the one or more portions of the audio encounter information.

[0092] Referring again to the example of FIG. 9, suppose automated clinical documentation process 10 determines at least a threshold amount or degree of correlation among portions 900, 902, 904, 906, 908, 910, 912 of audio encounter information 106 across microphones 202, 204, 206, 208 and a lack of correlation among portions 914, 916, 918, 920, 922, 924, 926 of audio encounter information 106. In this example, automated clinical documentation process 10 may identify 802 speech activity within portions 900, 902, 904, 906, 908, 910, 912 of audio encounter information 106 based upon, at least in part, a threshold on the correlation among portions 900, 902, 904, 906, 908, 910, 912 of audio encounter information 106. In some implementations, identifying 802 speech activity within the one or more portions of the audio encounter information may include generating timestamps indicating portions of audio encounter information 106 that include speech activity (e.g., start and end times for each portion). In some implementations, automated clinical documentation process 10 may generate a vector of start and end times for each portion with detected speech activity. In some implementations and as will be discussed in greater detail below, automated clinical documentation process 10 may label speech activity as a time domain label (i.e., a set of samples of the signal include or are speech) or as a set of frequency domain labels (i.e., a vector that gives the likelihood that

a particular frequency bin in a certain time frame includes or is speech).

[0093] In some implementations, automated clinical documentation process 10 may determine 804 location information for the one or more portions of the audio encounter information based upon, at least in part, the correlation among the signals received by each microphone of the microphone array. Location information may generally include a relative location or position of the origin of audio encounter information within an acoustic environment. For example, automated clinical documentation process 10 may determine 804 location information associated with the origin of a signal received by microphone array 200. Referring again to the FIG. 6 and in some implementations, automated clinical documentation process 10 may receive audio encounter information from various sources (e.g., encounter participants 226, 228, 230; noise sources 244, 250; etc.). In some implementations, automated clinical documentation process 10 may utilize information associated with the microphone array (e.g., the spacing between microphones of the microphone array; the orientation of the microphones within the microphone array; etc.) to determine 804 location information associated with the audio encounter information received by the microphone array.

[0094] As discussed above and in some implementations, automated clinical documentation process 10 may determine a correlation among the audio encounter information received from microphone array 200. For example, automated clinical documentation process 10 may compare the one or more portions (e.g., portions 900, 902, 904, 906, 908, 910, 912, 914, 916, 918, 920, 922, 924, 926) of audio encounter information 106 received from the various microphones of microphone array 200 to determine a degree of correlation between the audio encounter information present in each portion across the plurality of microphones of the microphone array.

[0095] In some implementations, determining 804 the location information for the one or more portions of the audio encounter information may include determining 808 a time

difference of arrival between each pair of microphones of the microphone array for the one or more portions of the audio encounter information. As is known in the art, a time difference of arrival (TDOA) between a pair of microphones may include locating a signal origin based upon, at least in part, the different arrival times of the signal at various receivers. In some implementations, automated clinical documentation process 10 may determine 808 the time difference of arrival (TDOA) between each pair of microphones for the one or more portions of the audio encounter information based upon, at least in part, the correlation among the audio encounter information received by the microphone array.

[0096] For example and in some implementations, a highly correlated signal (e.g., signal with at least a threshold degree of correlation) may allow for accurate determination of the time difference between microphones. As discussed above, suppose automated clinical documentation process 10 receives just ambient noise (i.e., no speech and no directional noise sources). Automated clinical documentation process 10 may determine that the spectrum observed in each microphone channel will be different (i.e. uncorrelated at each microphone). Accordingly, the time differences between pairs of microphones may be difficult to determine and/or may be inaccurate based upon the lack of correlation among the microphone channels of the microphone array. However, suppose automated clinical documentation process 10 receives a speech or other ‘directional’ signal within audio encounter information. In this example, automated clinical documentation process 10 may determine that one or more portions of the audio encounter information (e.g., the portions of audio encounter information with speech components) are highly correlated at each microphone in the microphone array. Accordingly, automated clinical documentation process 10 may more accurately determine the time difference of arrival between pairs of microphones using the highly correlated audio encounter information across the microphone channels of the microphone array.

[0097] In some implementations, identifying 802 the speech activity and determining

804 the location information may be performed jointly for the one or more portions of the audio encounter information. For example and as discussed above, correlation among portions of audio encounter information may identify the presence of speech activity within the one or more portions of the audio encounter information received by a microphone array. In addition and as discussed above, correlation among portions of audio encounter information may allow for a more accurate determination of location information (e.g., time difference of arrival) between pairs of microphones of the microphone array. Accordingly, automated clinical documentation process 10 may determine a correlation between the audio encounter information across the microphones or microphone channels of the microphone array and may utilize the determined correlation (e.g., by comparing the correlation to a threshold for speech activity detection and a threshold for time difference of arrival determination) to identify 802 speech activity within the one or more audio portions and to determine 804 the time difference of arrival between microphone pairs jointly.

[0098] In some implementations, jointly identifying 802 the speech activity and determining 804 the location information for the one or more portions of the audio encounter information may be performed using a machine learning model. In some implementations, automated clinical documentation process 10 may train a machine learning model to “learn” how to jointly identify speech activity and determine location information for the one or more portions of the audio encounter information. As is known in the art, a machine learning model may generally include an algorithm or combination of algorithms that has been trained to recognize certain types of patterns. For example, machine learning approaches may be generally divided into three categories, depending on the nature of the signal available: supervised learning, unsupervised learning, and reinforcement learning. As is known in the art, supervised learning may include presenting a computing device with example inputs and their desired outputs, given by a “teacher”,

where the goal is to learn a general rule that maps inputs to outputs. With unsupervised learning, no labels are given to the learning algorithm, leaving it on its own to find structure in its input. Unsupervised learning can be a goal in itself (discovering hidden patterns in data) or a means towards an end (feature learning). As is known in the art, reinforcement learning may generally include a computing device interacting in a dynamic environment in which it must perform a certain goal (such as driving a vehicle or playing a game against an opponent). As it navigates its problem space, the program is provided feedback that's analogous to rewards, which it tries to maximize. While three examples of machine learning approaches have been provided, it will be appreciated that other machine learning approaches are possible within the scope of the present disclosure.

[0099] Accordingly, automated clinical documentation process 10 may utilize a machine learning model (e.g., machine learning model 420) to identify speech activity within and location information for the one or more portions of the audio encounter information. For example, automated clinical documentation process 10 may provide time domain waveform data or frequency domain features as input to machine learning model 420. In some implementations, the phase spectrum features may provide spatial information for identifying speech activity and location information for the one or more portions of the audio encounter information.

[00100] For example, automated clinical documentation process 10 may provide training data associated with various portions of audio encounter information pre-labeled with speech activity information and/or location information to machine learning model 420. Automated clinical documentation process 10 may train machine learning model 420 to identify 802 speech activity and to determine 804 location information for the one or more portions of the audio encounter information based upon, at least in part, a degree of correlation among the portions of the audio encounter information and the training data. In this manner, machine learning model 420 may be configured to learn how correlation

between portions of audio encounter information maps to speech activity within and accurate location information for the portions of the audio encounter information. Accordingly, and as shown in FIG. 4, automated clinical documentation process 10 may implement VAD module 400 and sound localization module 406 with machine learning model 420.

[00101] In some implementations, automated clinical documentation process 10 may receive 810 information associated with an acoustic environment. As discussed above, an acoustic environment may represent the layout and acoustic properties of a room or other space where multiple microphone systems may be deployed. For example, information associated with the acoustic environment may describe the type / dimensions of a room, the activity zones in the room where the participants operate, the acoustic properties of the room (e.g., the types of noise expected, range of reverberation, etc.), position of the microphone system(s) within the acoustic environment, etc. In some implementations, automated clinical documentation process 10 may provide a user interface to receive information associated with the acoustic environment. Accordingly, a user and/or automated clinical documentation process 10 may provide the information associated with an acoustic environment via the user interface. However, it will be appreciated that the information associated with the acoustic environment may be received in various ways (e.g., default information for a default acoustic environment, automatically defined by automated clinical documentation process 10, etc.).

[00102] In some implementations, the information associated with the acoustic environment may indicate one or more target speaker locations within the acoustic environment. Referring again to FIG. 6 and in some implementations, suppose that acoustic environment 600 includes one or more encounter participants (e.g., encounter participants 226, 228, 230) of the above-described patient encounter. As described above, examples of the encounter participants (e.g., encounter participants 226, 228, 230) may

include but are not limited to: medical professionals (e.g., doctors, nurses, physician's assistants, lab technicians, physical therapists, scribes (e.g., a transcriptionist) and/or staff members involved in the patient encounter), patients (e.g., people that are visiting the above-described clinical environments for the patient encounter), and third parties (e.g., friends of the patient, relatives of the patient and/or acquaintances of the patient that are involved in the patient encounter).

[00103] In some implementations, automated clinical documentation process 10 may receive 810 information associated with acoustic environment 600 that may indicate locations within acoustic environment 600 where encounter participants 226, 228, 230 may be positioned. In some implementations, the information associated with acoustic environment 600 may indicate locations within acoustic environment where the one or more target speakers are likely to be when speaking. For example, suppose automated clinical documentation process 10 determines that an examination table is positioned at approximately e.g., 45° from the base of microphone array 200. In this example, automated clinical documentation process 10 may determine that a patient is most likely to speak while sitting on or near the examination table at approximately 45° from the base of microphone array 200. Further suppose that automated clinical documentation process 10 determines that a doctor's desk is positioned at approximately e.g., 90° from the base of microphone array 200. In this example, automated clinical documentation process 10 may determine that a doctor is most likely to speak while sitting on or near the desk at approximately 90° from the base of microphone array 200. Additionally, suppose automated clinical documentation process 10 determines that a waiting area is positioned at approximately e.g., 120° from the base of microphone array 200. In this example, automated clinical documentation process 10 may determine that other patients or other third parties are most likely to speak from at approximately 120° from the base of microphone array 200. While three examples of relative locations for target speakers

within an acoustic environment have been provided, it will be appreciated that the information associated with the acoustic environment may include any number of target speaker locations or probability-based target speaker locations for any number of target speakers within the scope of the present disclosure.

[00104] In some implementations, identifying 802 the speech activity within the one or more portions of the audio encounter information may be based upon, at least in part, the location information for the one or more portions of the audio encounter information and the information associated with the acoustic environment. For example, automated clinical documentation process 10 may utilize the information associated with the acoustic environment to more accurately identify speech activity within the one or more portions of the audio encounter information. In one example, suppose that automated clinical documentation process 10 receives information associated with the acoustic environment that indicates locations within the acoustic environment where speakers are likely to be. In this example, automated clinical documentation process 10 may, based upon the location information for the one or more portions of the audio encounter information, determine whether the one or more portions of the audio encounter information are within the potential speaker locations defined by the acoustic environment information. In response to determining that the one or more portions of the audio encounter information originate from a potential speaker location within the acoustic environment, automated clinical documentation process 10 may determine a higher probability that the one or more portions of the audio encounter information include speech activity. Accordingly, automated clinical documentation process 10 may utilize the location information and the acoustic environment information to identify 802 speech activity within the one or more portions of the audio encounter information.

[00105] In some implementations, automated clinical documentation process 10 may label 806 the one or more portions of the audio encounter information with the speech

activity and the location information. Labeling one or more portions of the audio encounter information with the speech activity and the location information may generally include generating metadata for the one or more portions of the audio encounter information with the speech activity and location information associated with each respective portion of the audio encounter information. Referring also to FIG. 10 and in some implementations, automated clinical documentation process 10 may label 806 portions 900, 902, 904, 906, 908, 910, 912 with the speech activity and location information associated with each respective portion by generating acoustic metadata for each portion of audio encounter information 106 (e.g., acoustic metadata 1000, 1002, 1004, 1006, 1008, 1010, 1012). In some implementations, the acoustic metadata associated with the audio encounter information may include the speech activity and the location information associated with the audio encounter information. In some implementations, automated clinical documentation process 10 may generate acoustic metadata 1000, 1002, 1004, 1006, 1008, 1010, 1012 that identifies portions of audio encounter information 106 that includes speech activity (e.g., portions 900, 902, 904, 906, 908, 910, 912). In one example, automated clinical documentation process 10 may generate acoustic metadata with timestamps indicating portions of audio encounter information 106 that include speech activity (e.g., start and end times for each portion). In some implementations, automated clinical documentation process 10 may label 806 speech activity as a time domain label (i.e., a set of samples of the signal include or are speech) or as a set of frequency domain labels (i.e., a vector that gives the likelihood that a particular frequency bin in a certain time frame includes or is speech).

[00106] Referring again to FIG. 4 and as will be discussed in greater detail below, automated clinical documentation process 10 may provide the acoustic metadata (e.g., acoustic metadata 1014) to speaker identification and tracking module 404 to identify a speaker from the audio encounter information and/or to track the position of a speaker

within an acoustic environment based upon, at least in part, acoustic metadata 1014.

[00107] Referring also at least to FIGS. 11-14, automated clinical documentation process 10 may receive 1100 information associated with an acoustic environment. Acoustic metadata associated with audio encounter information received by a first microphone system may be received 1102. One or more speaker representations may be defined 1104 based upon, at least in part, the acoustic metadata associated with the audio encounter information and the information associated with the acoustic environment. One or more portions of the audio encounter information may be labeled 1106 with the one or more speaker representations and a speaker location within the acoustic environment.

[00108] In some implementations, automated clinical documentation process 10 may receive 1100 information associated with an acoustic environment. As discussed above, an acoustic environment may represent the layout and acoustic properties of a room or other space where multiple microphone systems may be deployed. For example, information associated with the acoustic environment may describe the type / dimensions of a room, the activity zones in the room where the participants operate, the acoustic properties of the room (e.g., the types of noise expected, range of reverberation, etc.), position of the microphone system(s) within the acoustic environment, etc. In some implementations, automated clinical documentation process 10 may provide a user interface to receive information associated with the acoustic environment. Accordingly, a user and/or automated clinical documentation process 10 may provide the information associated with an acoustic environment via the user interface. However, it will be appreciated that the information associated with the acoustic environment may be received in various ways (e.g., default information for a default acoustic environment, automatically defined by automated clinical documentation process 10, etc.).

[00109] In some implementations, the information associated with the acoustic environment may indicate one or more target speaker locations within the acoustic

environment. Referring again to FIG. 6 and in some implementations, suppose that acoustic environment 600 includes one or more encounter participants (e.g., encounter participants 226, 228, 230) of the above-described patient encounter. As described above, examples of the encounter participants (e.g., encounter participants 226, 228, 230) may include but are not limited to: medical professionals (e.g., doctors, nurses, physician's assistants, lab technicians, physical therapists, scribes (e.g., a transcriptionist) and/or staff members involved in the patient encounter), patients (e.g., people that are visiting the above-described clinical environments for the patient encounter), and third parties (e.g., friends of the patient, relatives of the patient and/or acquaintances of the patient that are involved in the patient encounter).

[00110] In some implementations, automated clinical documentation process 10 may receive 1100 information associated with acoustic environment 600 that may indicate locations within acoustic environment 600 where encounter participants 226, 228, 230 may be positioned. In some implementations, the information associated with acoustic environment 600 may indicate locations within acoustic environment where the one or more target speakers are likely to be when speaking. For example, suppose automated clinical documentation process 10 determines that an examination table is positioned at approximately e.g., 45° from the base of microphone array 200. In this example, automated clinical documentation process 10 may determine that a patient is most likely to speak while sitting on or near the examination table at approximately 45° from the base of microphone array 200. Further suppose that automated clinical documentation process 10 determines that a doctor's desk is positioned at approximately e.g., 90° from the base of microphone array 200. In this example, automated clinical documentation process 10 may determine that a doctor is most likely to speak while sitting on or near the desk at approximately 90° from the base of microphone array 200.

[00111] Additionally, suppose automated clinical documentation process 10

determines that a waiting area is positioned at approximately e.g., 120° from the base of microphone array 200. In this example, automated clinical documentation process 10 may determine that other patients or other third parties are most likely to speak from at approximately 120° from the base of microphone array 200. While three examples of relative locations for target speakers within an acoustic environment have been provided, it will be appreciated that the information associated with the acoustic environment may include any number of target speaker locations or probability-based target speaker locations for any number of target speakers within the scope of the present disclosure.

[00112] In some implementations, automated clinical documentation process 10 may receive 1102 acoustic metadata associated with audio encounter information received by a first microphone system. Referring again to FIG. 4 and in some implementations, VAD module 400 and sound localization module 406 may generate acoustic metadata associated with audio encounter information. For example and in some implementations, the acoustic metadata associated with the audio encounter information may include speech activity information and signal location information associated with the audio encounter information. As discussed above and in some implementations, automated clinical documentation process 10 may identify portions of the audio encounter information (e.g., audio encounter information 106) with speech components. Automated clinical documentation process 10 may associate or label the portions of audio encounter information 106 as having speech activity. In some implementations, automated clinical documentation process 10 may generate acoustic metadata that identifies portions of audio encounter information 106 that includes speech activity.

[00113] In one example, automated clinical documentation process 10 may generate timestamps indicating portions of audio encounter information 106 that include speech activity (e.g., start and end times for each portion). Referring also to FIG. 12, automated clinical documentation process 10 may receive acoustic metadata 1200 for one or more

portions of audio encounter information 106 (e.g., portions 1202, 1204, 1206, 1208, 1210, 1212, 1214, 1216, 1218, 1220, 1222, 1224, 1226, 1228). As discussed above, automated clinical documentation process 10 may generate acoustic metadata with speech activity and signal location information associated with each portion of the audio encounter information (e.g., represented in FIG. 12 as acoustic metadata 1230, 1232, 1234, 1236, 1238, 1240, 1242, 1244, 1246, 1248, 1250, 1252, 1254, 1256 for portions 1202, 1204, 1206, 1208, 1210, 1212, 1214, 1216, 1218, 1220, 1222, 1224, 1226, 1228, respectively) of audio encounter information 106).

[00114] Additionally, automated clinical documentation process 10 may determine location information for audio encounter information 106. For example, automated clinical documentation process 10 may determine spatial information from the microphone array (e.g., time difference of arrival (TDOA) between the microphones). As is known in the art, once a signal is received at two reference points, the difference in arrival time may be used to calculate the difference in distances between the target and the two reference points. In this example, automated clinical documentation process 10 may determine the time difference of arrival (TDOA) between microphones of a microphone array. In some implementations, automated clinical documentation process 10 may generate acoustic metadata with the TDOA for a particular microphone array.

[00115] In some implementations, automated clinical documentation process 10 may define 1104 one or more speaker representations based upon, at least in part, the acoustic metadata associated with the audio encounter information and the information associated with the acoustic environment. A speaker representation may generally include a cluster of data associated with a unique speaker within an acoustic environment. For example, automated clinical documentation process 10 may cluster spatial information and spectral information into separate speaker representations to account for the combination of spatial and spectral data that reference unique speakers within an acoustic environment.

[00116] Referring again to FIG. 6 and in some implementations, suppose automated clinical documentation process 10 receives 1100 information associated with acoustic environment 600 that may indicate locations within acoustic environment 600 where encounter participants 226, 228, 230 may be positioned. Further suppose that automated clinical documentation process 10 receives 1102 acoustic metadata (e.g., acoustic metadata 1200) associated with audio encounter information 106 received by microphone array 200. In this example, suppose that acoustic metadata 1200 includes spatial information (e.g., location information for audio encounter information 106) and spectral information (e.g., speech activity information). In some implementations, the spectral information may include acoustic features (e.g., Mel frequency cepstral coefficients (MFCCs)) associated with particular speakers (e.g., encounter participants 226, 228, 230).

[00117] Referring also to FIG. 13 and in some implementations, automated clinical documentation process 10 may define 1104 speaker representations for encounter participants 226, 228, 230 based upon, at least in part, the acoustic metadata associated with the audio encounter information and the information associated with the acoustic environment. For example, automated clinical documentation process 10 may use a dynamic model of speaker movements and the information associated with the acoustic environment to cluster the spatial information (e.g., TDOA) and spectral information (e.g., acoustic features like MFCCs) into separate speaker representations (e.g., speaker representation 1300 for encounter participant 226; speaker representation 1302 for encounter participant 228; and speaker representation 1304 for encounter participant 230). While the above example includes defining e.g., three speaker representations, it will be appreciated that any number of speaker representations may be defined by automated clinical documentation process 10 within the scope of the present disclosure.

[00118] In some implementations, automated clinical documentation process 10 may receive 1108 visual metadata associated with one or more encounter participants

within the acoustic environment. For example, automated clinical documentation process 10 may be configured to track the movement and/or interaction of humanoid shapes within the monitored space (e.g., acoustic environment 600) during the patient encounter (e.g., a visit to a doctor's office). Accordingly and referring again to FIG. 2, automated clinical documentation process 10 may process the machine vision encounter information (e.g., machine vision encounter information 102) to identify one or more humanoid shapes. As discussed above, examples of machine vision system 100 generally (and ACD client electronic device 34 specifically) may include but are not limited to one or more of an RGB imaging system, an infrared imaging system, an ultraviolet imaging system, a laser imaging system, a SONAR imaging system, a RADAR imaging system, and a thermal imaging system.

[00119] When ACD client electronic device 34 includes a visible light imaging system (e.g., an RGB imaging system), ACD client electronic device 34 may be configured to monitor various objects within acoustic environment 600 by recording motion video in the visible light spectrum of these various objects. When ACD client electronic device 34 includes an invisible light imaging systems (e.g., a laser imaging system, an infrared imaging system and/or an ultraviolet imaging system), ACD client electronic device 34 may be configured to monitor various objects within acoustic environment 600 by recording motion video in the invisible light spectrum of these various objects. When ACD client electronic device 34 includes an X-ray imaging system, ACD client electronic device 34 may be configured to monitor various objects within acoustic environment 600 by recording energy in the X-ray spectrum of these various objects. When ACD client electronic device 34 includes a SONAR imaging system, ACD client electronic device 34 may be configured to monitor various objects within acoustic environment 600 by transmitting soundwaves that may be reflected off of these various objects. When ACD client electronic device 34 includes a RADAR imaging system, ACD client electronic

device 34 may be configured to monitor various objects within acoustic environment 600 by transmitting radio waves that may be reflected off of these various objects. When ACD client electronic device 34 includes a thermal imaging system, ACD client electronic device 34 may be configured to monitor various objects within acoustic environment 600 by tracking the thermal energy of these various objects.

[00120] As discussed above, ACD computer system 12 may be configured to access one or more datasources 118 (e.g., plurality of individual datasources 120, 122, 124, 126, 128), wherein examples of which may include but are not limited to one or more of a user profile datasource, a voice print datasource, a voice characteristics datasource (e.g., for adapting the automated speech recognition models), a face print datasource, a humanoid shape datasource, a humanoid shape datasource, an utterance identifier datasource, a wearable token identifier datasource, an interaction identifier datasource, a medical conditions symptoms datasource, a prescriptions compatibility datasource, a medical insurance coverage datasource, and a home healthcare datasource.

[00121] Accordingly and when processing the machine vision encounter information (e.g., machine vision encounter information 102) to identify one or more humanoid shapes, automated clinical documentation process 10 may be configured to compare the humanoid shapes defined within one or more datasources 118 to potential humanoid shapes within the machine vision encounter information (e.g., machine vision encounter information 102).

[00122] When processing the machine vision encounter information (e.g., machine vision encounter information 102) to identify one or more humanoid shapes, automated clinical documentation process 10 may track the movement of the one or more humanoid shapes within the monitored space (e.g., acoustic environment 600). For example and when tracking the movement of the one or more humanoid shapes within acoustic environment 600, automated clinical documentation process 10 may add a new humanoid

shape to the one or more humanoid shapes when the new humanoid shape enters the monitored space (e.g., acoustic environment 600) and/or may remove an existing humanoid shape from the one or more humanoid shapes when the existing humanoid shape leaves the monitored space (e.g., acoustic environment 600).

[00123] Also and when tracking the movement of the one or more humanoid shapes within acoustic environment 600, automated clinical documentation process 10 may monitor the trajectories of the various humanoid shapes within acoustic environment 600. Accordingly, assume that when leaving acoustic environment 600, encounter participant 242 walks in front of (or behind) encounter participant 226. As automated clinical documentation process 10 is monitoring the trajectories of (in this example) encounter participant 242 (who is e.g., moving from left to right) and encounter participant 226 (who is e.g., stationary), when encounter participant 242 passes in front of (or behind) encounter participant 226, the identities of these two humanoid shapes may not be confused by automated clinical documentation process 10.

[00124] Automated clinical documentation process 10 may be configured to obtain the encounter information of the patient encounter (e.g., a visit to a doctor's office), which may include machine vision encounter information 102 (in the manner described above) and/or audio encounter information 106. In some implementations, automated clinical documentation process 10 may generate visual metadata associated with machine vision encounter information 102. For example and as discussed above, automated clinical documentation process 10 may generate visual metadata 1306 indicative of the direction or position of speakers within acoustic environment 600, the number of speakers within acoustic environment 600, and/or the identity of the speakers within acoustic environment 600. As shown in the example of FIG. 13, visual metadata 1306 may be defined for each portion of audio encounter information 106 (e.g., portions 1202, 1204, 1206, 1208, 1210, 1212, 1214, 1216, 1218, 1220, 1222, 1224, 1226, 1228 of audio encounter information

106) and represented as separate visual metadata for each portion (e.g., visual metadata 1308, 1310, 1312, 1314, 1316, 1318, 1320, 1322, 1324, 1326, 1328, 1330, 1332, 1334).

[00125] In some implementations, defining 1104 the one or more speaker representations may include defining 1110 the one or more speaker representations based upon, at least in part, the visual metadata associated with the one or more encounter participants within the acoustic environment. For example and as discussed above, machine vision system 100 may detect and track location estimates for particular speaker representations by tracking the humanoid shapes within the acoustic environment. Automated clinical documentation process 10 may “fuse” the location estimates from the visual metadata with the acoustic localization information to cluster the spatial-spectral features of the visual metadata and the acoustic metadata into the plurality of speaker representations.

[00126] Referring again to the example of FIG. 6, automated clinical documentation process 10 may receive 1108 visual metadata 1306 associated with encounter participants 226, 228, 230. In this example, visual metadata 1306 may indicate the relative position of each speaker within acoustic environment 600; the number of speakers; and/or the identity of the speakers within acoustic environment 600. For example and as discussed above, machine vision system 100 may be configured to identify one or more humanoid shapes and track the movement of the one or more humanoid shapes within the acoustic environment. Automated clinical documentation process 10 may receive visual metadata 1306 associated with the identity and/or location of encounter participants 226, 228, 230. Suppose that visual metadata 1306 includes an identity of each participant (e.g., based upon, at least in part, comparing the humanoid shapes defined within one or more datasources 118 to potential humanoid shapes within the machine vision encounter information (e.g., machine vision encounter information 102)).

[00127] Further suppose that automated clinical documentation process 10 receives

1102 acoustic metadata 1200 associated with audio encounter information 106. Automated clinical documentation process 10 may define 1110 the one or more speaker representations (e.g., speaker representations 1300, 1302, 1304 for encounter participants 226, 228, 230, respectively) based upon, at least in part, visual metadata 1306 associated with encounter participants 226, 228, 230 within acoustic environment 600 and acoustic metadata 1200 associated with the audio encounter information 106. In this example, automated clinical documentation process 10 may combine the location estimates from visual metadata 1306 with the location information of acoustic metadata 1200 to cluster the spatial-spectral features of visual metadata 1306 and acoustic metadata 1200 into speaker representations 1300, 1302, 1304 for encounter participants 226, 228, 230, respectively.

[00128] In some implementations, automated clinical documentation process 10 may receive 1112 weighting metadata associated with audio encounter information received by a second microphone system. As will be discussed in greater detail below, device selection and weighting module 410 of ACD computer system 12 may be configured to weight multiple audio streams (e.g., from different microphone systems) based upon, at least in part, a signal-to-noise ratio for each audio stream, thus defining a weight for each audio stream. In some implementations, the weight may be defined for each audio stream such that the estimated speech processing system performance is maximized. For example, automated clinical documentation process 10 may train the device selection and weighting module based upon, at least in part, the signal-to-noise (SNR) ratio for each portion or frame of each audio stream. However, it will be appreciated that other metrics or properties associated with each audio stream may be utilized to select and weight each audio stream. For example, automated clinical documentation process 10 may train the device selection and weighting module based upon, at least in part, a reverberation level (e.g., a C50 ratio) for each portion or frame of each audio stream. While two examples of particular metrics that may be used to select and weight each audio stream have been provided, it will be

appreciated that any metric or property may be used to select and/or weigh various audio streams within the scope of the present disclosure. In some implementations, device selection and weighting module 410 may provide a previously processed or weighted portion of the audio encounter information to define weighting for each audio stream from multiple microphone systems. In this manner, automated clinical documentation process 10 may utilize audio encounter information from multiple microphone systems when tracking and/or identifying a speaker.

[00129] In some implementations, defining 1104 the one or more speaker representations may include defining 1114 the one or more speaker representations based upon, at least in part, the weighting metadata associated with the audio encounter information received by the second microphone system. For example, automated clinical documentation process 10 may utilize weighting metadata (e.g., weighting metadata 1336) associated with audio encounter information received by a second microphone system (e.g., second microphone system 416) to help identify which speaker is speaking at a given time within the acoustic environment (e.g., acoustic environment 600).

[00130] For example, suppose that encounter participant 226 has a second microphone system (e.g., mobile electronic device 416) nearby (e.g., in a pocket). In some implementations, the second microphone system (e.g., mobile electronic device 416) may receive audio encounter information 106 from encounter participants 226, 228, 230. As will be discussed in greater detail below, automated clinical documentation process 10 may apply a weight to portions of the audio encounter information received by the first microphone system (e.g., microphone array 200) and a weight to portions of the audio encounter information received by the second microphone system (e.g., mobile electronic device 416). In some implementations, automated clinical documentation process 10 may generate weighting metadata (e.g., weighting metadata 1336) associated with the audio encounter information received by the second microphone system (e.g., mobile electronic

device 416). For example, weighting metadata 1336 may indicate that, based on audio encounter information received by mobile electronic device 416, encounter participant 226 is speaking. Accordingly, automated clinical documentation process 10 may define 1114 the speaker representation for encounter participant 226 based upon, at least in part, weighting metadata 1336 associated with the audio encounter information received by the second microphone system.

[00131] In some implementations, defining 1104 the one or more speaker representations may include defining 1116 one or more of: at least one known speaker representation and at least one unknown speaker representation. For example and as discussed above, ACD computer system 12 may be configured to access one or more datasources 118 (e.g., plurality of individual datasources 120, 122, 124, 126, 128), examples of which may include but are not limited to one or more of a user profile datasource, a voice print datasource, a voice characteristics datasource (e.g., for adapting the automated speech recognition models), a face print datasource, a humanoid shape datasource, an utterance identifier datasource, a wearable token identifier datasource, an interaction identifier datasource, a medical conditions symptoms datasource, a prescriptions compatibility datasource, a medical insurance coverage datasource, and a home healthcare datasource.

[00132] In some implementations, process 10 may compare the data included within the user profile (defined within the user profile datasource) to at least a portion of the audio encounter information and/or machine vision encounter information. The data included within the user profile may include voice-related data (e.g., a voice print that is defined locally within the user profile or remotely within the voice print datasource), language use patterns, user accent identifiers, user-defined macros, and user-defined shortcuts, for example. Specifically and when attempting to associate at least a portion of the audio encounter information with at least one known encounter participant, automated clinical

documentation process 10 may compare one or more voice prints (defined within the voice print datasource) to one or more voices defined within the audio encounter information.

[00133] As discussed above and for this example, assume: that encounter participant 226 is a medical professional that has a voice print / profile; that encounter participant 228 is a patient that has a voice print / profile; and that encounter participant 230 is a third party (the acquaintance of encounter participant 228) and, therefore, does not have a voice print / profile. Accordingly and for this example: assume that automated clinical documentation process 10 will be successful and identify encounter participant 226 when comparing audio encounter information 106A to the various voice prints / profiles included within voice print datasource; assume that automated clinical documentation process 10 will be successful and identify encounter participant 228 when comparing audio encounter information 106B to the various voice prints / profiles included within voice print datasource; and assume that automated clinical documentation process 10 will be unsuccessful and not identify encounter participant 230 when comparing audio encounter information 106C to the various voice prints / profiles included within voice print datasource.

[00134] Accordingly and when processing the encounter information (e.g., machine vision encounter information 102 and/or audio encounter information 106), automated clinical documentation process 10 may associate audio encounter information 106A with the voice print / profile of Doctor Susan Jones and may identify encounter participant 226 as “Doctor Susan Jones”. Automated clinical documentation process 10 may further associate audio encounter information 106B with the voice print / profile of Patient Paul Smith and may identify encounter participant 228 as “Patient Paul Smith”. Further, automated clinical documentation process 10 may not be able to associate audio encounter information 106C with any voice prints / profiles and may identify encounter participant 230 as “Unknown Participant”.

[00135] As described above, automated clinical documentation process 10 may define 1116 a known speaker representation for “Doctor Susan Jones” (e.g., speaker representation 1300 for encounter participant 226) and a known speaker representation for “Patient Paul Smith” (e.g., speaker representation 1302 for encounter participant 228) based upon, at least in part, acoustic metadata 1200 associated with audio encounter information 106 and visual metadata 1306 associated with machine vision encounter information 102. Similarly, automated clinical documentation process 10 may define an unknown speaker representation for participant 230 (e.g., speaker representation 1304 for encounter participant 230).

[00136] In some implementations, automated clinical documentation process 10 may label 1106 one or more portions of the audio encounter information with the one or more speaker representations and a speaker location within the acoustic environment. For example, labeling 1106 one or more portions of the audio encounter information with the one or more speaker representations and a speaker location within the acoustic environment may generally include associating a speaker representation and speaker location information within the acoustic environment with a particular portion (e.g., a segment or frame) of the audio encounter information. For example, automated clinical documentation process 10 may generate labels for each portion of the audio encounter information with the speaker representation and a speaker location within the acoustic environment associated with each portion.

[00137] Referring also to FIG. 14 and in some implementations, automated clinical documentation process 10 may label one or more portions of the audio encounter information (e.g., portions 1202, 1204, 1206, 1208, 1210, 1212, 1214, 1216, 1218, 1220, 1222, 1224, 1226, 1228 of audio encounter information 106) with the one or more speaker representations and a speaker location within the acoustic environment. For example, automated clinical documentation process 10 may, for each portion of audio encounter

information 106, generate a “label” or speaker metadata 1400 (e.g., represented as speaker metadata 1402, 1404, 1406, 1408, 1410, 1412, 1414, 1416, 1418, 1420, 1422, 1424, 1426, 1428 defined for portions 1202, 1204, 1206, 1208, 1210, 1212, 1214, 1216, 1218, 1220, 1222, 1224, 1226, 1228, respectively) with the speaker representation and speaker location information associated with that portion of audio encounter information 106.

[00138] Continuing with the above example, suppose portions 1202, 1204, 1206, 1208, 1210 include speech from encounter participant 226; portions 1212, 1214, 1216, 1218, 1220 include speech from encounter participant 228; and that portions 1222, 1224, 1226, 1228 include speech from encounter participant 230. In this example, automated clinical documentation process 10 may label portions 1202, 1204, 1206, 1208, 1210 with speaker representation 1300 and a speaker location within the acoustic environment associated with speaker 226 during the associated portion of the audio encounter information; may label portions 1212, 1214, 1216, 1218, 1220 with speaker representation 1302 and a speaker location within the acoustic environment associated with encounter participant 228 during the associated portions of the audio encounter information; and may label portions 1222, 1224, 1226, 1228 with speaker representation 1304 and a speaker location within the acoustic environment associated with encounter participant 230 during the associated portions of the audio encounter information. As will be discussed in greater detail below, speaker metadata 1400 may be provided to a beam and null selection module (e.g., beam and null selection module 406) of ACD computing system 12 and may allow for the selection and/or combination of particular beams and/or nulls based upon, at least in part, the speaker representation and speaker location defined within speaker metadata 1400.

[00139] Referring also at least to FIGS. 15-16, automated clinical documentation process 10 may receive 1500 a plurality of predefined beams associated with a microphone array. A plurality of predefined nulls associated with the microphone array may be received

1502. One or more predefined beams from the plurality of predefined beams or one or more predefined nulls from the plurality of predefined nulls may be selected 1504. A microphone array may obtain 1506 audio encounter information, via the microphone array, using at least one of the one or more selected beams and the one or more selected nulls.

[00140] In some implementations, automated clinical documentation process 10 may receive 1500 a plurality of predefined beams associated with a microphone array. As discussed above, automated clinical documentation process 10 may predefine a plurality of predefined beams associated with a microphone array. For example, automated clinical documentation process 10 may predefine a plurality of beams based upon, at least in part, information associated with acoustic environment. As discussed above, a beam may generally include a pattern of constructive interference among microphones of a microphone array that is generated by modifying the phase and/or amplitude of the signal at each microphone of the microphone array. The pattern of constructive interference may improve the signal processing performance of the microphone array. In some implementations, automated clinical documentation process 10 may predefine the plurality of beams by adjusting the phase and/or amplitude of the signal at each microphone of the microphone array. In some implementations, automated clinical documentation process 10 may provide the predefined beams to beam and null selection module 408. In one example, automated clinical documentation process 10 may receive 1500 the plurality of predefined beams as a vector of phases and/or amplitudes for each signal of each microphone channel of the microphone array to achieve a desired sensitivity pattern. In another example and as discussed above, automated clinical documentation process 10 may receive 1500 the plurality of predefined beams as a plurality of predefined filters that produce the plurality of predefined beams.

[00141] In some implementations, the plurality of predefined beams may include one or more predefined beams configured to receive audio encounter information from one

or more target speaker locations within an acoustic environment. For example and as discussed above, automated clinical documentation process 10 may receive information associated with an acoustic environment. In some implementations, the acoustic environment information may indicate one or more target speaker locations where a speaker is likely to speak from within the acoustic environment. Accordingly, automated clinical documentation process 10 may receive 1500 predefined beams configured to receive audio encounter information from the one or more target speaker locations within the acoustic environment.

[00142] For example and referring again to FIG. 6, automated clinical documentation process 10 may receive 1500 a plurality of predefined beams configured to receive audio encounter information from the one or more target speaker locations within the acoustic environment. Suppose that automated clinical documentation process 10 determines from the information associated with acoustic environment 600, that a patient is most likely to speak while sitting on or near the examination table at approximately 45° from the base of microphone array 200; a doctor is most likely to speak while sitting on or near the desk at approximately 90° from the base of microphone array 200; and that other patients or other third parties are most likely to speak from at approximately 120° from the base of microphone array 200. In this example, automated clinical documentation process 10 may receive 1500 predefined beam 220 configured for receiving audio encounter information from a patient (e.g., encounter participant 228); beam 222 configured for receiving audio encounter information from a doctor (e.g., encounter participant 226); and beam 224 configured for receiving audio encounter information from another patient / third party (e.g., encounter participant 230).

[00143] In some implementations, automated clinical documentation process 10 may receive 1502 a plurality of predefined nulls associated with the microphone array. As discussed above, a null may generally include a pattern of destructive interference among

microphones of a microphone array that is generated by modifying the phase and/or amplitude of the signal at each microphone of the microphone array. The pattern of destructive interference may limit the reception of signals by the microphone array. In some implementations, automated clinical documentation process 10 may predefine the plurality of nulls by adjusting the phase and/or amplitude of the signal at each microphone of the microphone array via a plurality of filters (e.g., a plurality of FIR filters). In one example, automated clinical documentation process 10 may receive 1502 the plurality of predefined nulls as a vector of phases and/or amplitudes for each signal of each microphone channel of the microphone array to achieve a desired sensitivity pattern. In another example and as discussed above, automated clinical documentation process 10 may receive 1502 the plurality of predefined nulls as a plurality of predefined filters that produce the plurality of predefined nulls.

[00144] In some implementations, the plurality of predefined nulls may include one or more predefined nulls configured to limit receiving audio encounter information from one or more target speaker locations within an acoustic environment. For example and as discussed above, automated clinical documentation process 10 may receive information associated with an acoustic environment. In some implementations, the acoustic environment information may indicate one or more target speaker locations where a speaker is likely to speak from within the acoustic environment. Accordingly, automated clinical documentation process 10 may receive 1502 predefined beams configured to limit receiving audio encounter information from the one or more target speaker locations within the acoustic environment.

[00145] For example and referring again to FIG. 6, automated clinical documentation process 10 may receive 1502 a plurality of predefined nulls configured to receive audio encounter information from the one or more target speaker locations within the acoustic environment. Suppose that automated clinical documentation process 10

determines from the information associated with acoustic environment 600, that a patient is most likely to speak while sitting on or near the examination table at approximately 45° from the base of microphone array 200; a doctor is most likely to speak while sitting on or near the desk at approximately 90° from the base of microphone array 200; and that other patients or other third parties are most likely to speak from at approximately 120° from the base of microphone array 200. In this example and referring also to FIG. 7, automated clinical documentation process 10 may predefine null 252 to limit receiving audio encounter information from a patient (e.g., encounter participant 228); null 254 to limit receiving audio encounter information from a doctor (e.g., encounter participant 226); and null 256 to limit receiving audio encounter information from another patient / third party (e.g., encounter participant 230).

[00146] In some implementations, automated clinical documentation process 10 may select 1504 at least one of: one or more predefined beams from the plurality of predefined beams, thus defining one or more selected beams; and one or more predefined nulls from the plurality of predefined nulls, thus defining one or more selected nulls. Selecting a predefined beam and/or a predefined null may include selecting a pattern or combination of predefined beams and/or predefined nulls to achieve a particular microphone array sensitivity. For example and as will be described in greater detail below, suppose automated clinical documentation process 10 determines that the doctor (e.g., encounter participant 226) is speaking. In this example, automated clinical documentation process 10 may select 1504 one or more beams and/or nulls to achieve a beamforming pattern that enables microphone array 200 to receive audio encounter information from doctor (e.g., encounter participant 226). In some implementations, the one or more selected beams may enable audio encounter information to be received from doctor (e.g., encounter participant 226) and the one or more selected nulls may limit the reception of audio encounter information from other speakers or noise sources.

[00147] In some implementations, automated clinical documentation process 10 may receive 1508 speaker metadata associated with one or more portions of audio encounter information. For example and as discussed above, automated clinical documentation process 10 may, for each portion of audio encounter information 106, generate a “label” or speaker metadata with the speaker representation and speaker location information associated with that portion of the audio encounter information. As discussed above, a speaker representation may include a cluster of data associated with a unique speaker within an acoustic environment and the speaker location information may include information indicating the location of the speaker within the acoustic environment. In some implementations, automated clinical documentation process 10 may utilize the speaker representation and speaker location information to select particular beams and/or nulls for beamforming via a microphone array.

[00148] In some implementations, selecting 1504 at least one of the one or more beams and the one or more nulls may include selecting 1510 at least one of the one or more beams and the one or more nulls based upon, at least in part, the speaker metadata associated with the one or more portions of the audio encounter information. For example, automated clinical documentation process 10 may determine a speaker location within the acoustic environment and a speaker identity from the speaker metadata. In one instance, suppose that speaker metadata 1400 indicates that the doctor (e.g., encounter participant 226) is speaking (e.g., based upon, at least in part, speaker metadata including a speaker identity associated with the doctor (e.g., encounter participant 226) and speaker location information indicating that audio encounter information is being received from near the doctor’s desk). In this example, automated clinical documentation process 10 may select one or more beams and/or one or more nulls based upon, at least in part, the speaker metadata indicating that the doctor (e.g., encounter participant 226) is speaking.

[00149] In some implementations, selecting 1510 at least one of the one or more

beams and the one or more nulls based upon, at least in part, the speaker metadata associated with the one or more portions of the audio encounter information may include selecting 1512 the one or more predefined beams configured to receive audio encounter information from the one or more target speaker locations within the acoustic environment based upon, at least in part, the speaker metadata associated with the one or more portions of the audio encounter information. Continuing with the above example where speaker metadata 1400 indicates that the doctor (e.g., encounter participant 226) is speaking, automated clinical documentation process 10 may determine which beam(s) provide microphone sensitivity for the speaker location included in speaker metadata 1400. For example, suppose automated clinical documentation process 10 determines that beam 222 (as shown in FIG. 6) provides microphone sensitivity at or adjacent to the speaker location included in speaker metadata 1400. In this example, automated clinical documentation process 10 may select 1512 beam 222 for receiving audio encounter information from the doctor (e.g., encounter participant 226). While an example of selecting a single beam has been provided, it will be appreciated that any number of beams may be selected within the scope of the present disclosure.

[00150] In some implementations, selecting 1504 at least one of the one or more beams and the one or more nulls based upon, at least in part, the speaker metadata associated with the one or more portions of the audio encounter information may include selecting 1514 the one or more predefined nulls configured to limit receiving audio encounter information from the one or more target speaker locations within the acoustic environment based upon, at least in part, the speaker metadata associated with the one or more portions of the audio encounter information. Continuing with the above example where speaker metadata 1400 indicates that the doctor (e.g., encounter participant 226) is speaking, automated clinical documentation process 10 may determine which null(s) limit microphone sensitivity for other speaker locations and/or noise sources included in speaker

metadata 1400. For example, suppose automated clinical documentation process 10 determines that null 252 (as shown in FIG. 7) limits microphone sensitivity at or adjacent to a speaker location associated with encounter participant 228 and that null 256 (as shown in FIG. 7) limits microphone sensitivity at or adjacent to a speaker location associated with encounter participants 230, 242.

[00151] In addition, suppose that automated clinical documentation process 10 determines that null 248 limits microphone sensitivity at or adjacent to a first noise source (e.g., fan 244) and that null 250 limits microphone sensitivity at or adjacent to a second noise source (e.g., doorway 246). In this example, automated clinical documentation process 10 may select 1514 nulls 248, 250, 252, 254 to limit receiving audio encounter information from other encounter participants and noise sources. While an example of selecting four nulls has been provided, it will be appreciated that any number of nulls may be selected within the scope of the present disclosure.

[00152] In some implementations, automated clinical documentation process 10 may select 1504 at least one of the one or more beams and the one or more nulls based upon, at least in part, the information associated with the acoustic environment. As discussed above and in some implementations, automated clinical documentation process 10 may receive information associated with the acoustic environment. In some implementations, automated clinical documentation process 10 may select particular beams and/or nulls to use when obtaining 1506 audio encounter information with a microphone array based upon, at least in part, the acoustic properties of the acoustic environment. For example, suppose the acoustic environment information indicates that acoustic environment 600 includes a particular reverberation level. In this example, automated clinical documentation process 10 may select particular beams and/or nulls to account for and/or to minimize signal degradation associated with the reverberation level of acoustic environment 600. Accordingly, automated clinical documentation process 10

may dynamically select beams and/or nulls based upon, at least in part, various acoustic properties of the acoustic environment.

[00153] In some implementations, automated clinical documentation process 10 may obtain 1506 audio encounter information, via the microphone array, using at least one of the one or more selected beams and the one or more selected nulls. In some implementations, automated clinical documentation process 10 may utilize the predefined plurality of beams and plurality of nulls to obtain 1506 audio encounter information from particular speakers and to limit receiving audio encounter information from other speakers. Continuing with the above example and as shown in FIG. 16, automated clinical documentation process 10 may obtain 1506 audio encounter information 106A from the doctor (e.g., encounter participant 226) via microphone array 200 with the one or more selected beams (e.g., beam 222) and the one or more nulls (e.g., nulls 248, 250, 252, 254).

[00154] In some implementations and referring again to FIG. 4, automated clinical documentation process 10 may provide the audio encounter information obtained 1506 with the selected beams and/or selected nulls to a device selection and weighting module (e.g., device selection and weighting module 410). As will be discussed in greater detail below, automated clinical documentation process 10 may provide audio encounter information received from the microphone array (e.g., a first microphone system) to device selection and weighting module 410 to determine which audio stream (i.e., stream of audio encounter information) to process with a speech processing system (e.g., speech processing system 418). As will be discussed in greater detail below, device selection and weighting module may select from a plurality of audio streams (e.g., an audio stream from the first microphone system and an audio stream from a second microphone system). In this manner, automated clinical documentation process 10 may select an audio stream (or from portions of an audio stream) from the plurality of audio streams to provide to a speech processing system.

[00155] Referring also at least to FIGS. 17-20B, automated clinical documentation process 10 may receive 1700 audio encounter information from a first microphone system, thus defining a first audio stream. Audio encounter information may be received 1702 from a second microphone system, thus defining a second audio stream. Speech activity may be detected 1704 in one or more portions of the first audio stream, thus defining one or more speech portions of the first audio stream. Speech activity may be detected 1706 in one or more portions of the second audio stream, thus defining one or more speech portions of the second audio stream. The first audio stream and the second audio stream may be aligned 1708 based upon, at least in part, the one or more speech portions of the first audio stream and the one or more speech portions of the second audio stream.

[00156] In some implementations, automated clinical documentation process 10 may receive 1700 audio encounter information from a first microphone system, thus defining a first audio stream. As discussed above and in some implementations, the first microphone system may be a microphone array. For example and as shown in FIG. 4, automated clinical documentation process 10 may receive audio encounter information (e.g., audio encounter information 106) from a first microphone system (e.g., microphone array 200 with audio acquisition devices 202, 204, 206, 208, 210, 212, 214, 216, 218). Referring again to FIG. 6, audio encounter information 106 may be received 1700 with one or more beams and/or one or more nulls generated by microphone array 200. As shown in FIG. 4, the solid lines between beam and null selection module 408 and device selection and weighting module 410; and between beam and null selection module 408 and alignment module 412 may represent first audio stream received from microphone array 200. Referring also to FIG. 18 and in some implementations, automated clinical documentation process 10 may receive audio encounter information with one or more portions (e.g., portions 1800, 1802, 1804, 1806, 1808, 1810, 1812, 1814, 1816, 1818, 1820, 1822, 1824, 1826) from a first microphone system (e.g., microphone array 200), thus

defining a first audio stream.

[00157] In some implementations, automated clinical documentation process 10 may receive 1702 audio encounter information from a second microphone system, thus defining a second audio stream. In some implementations, the second microphone system may be a mobile electronic device. For example and as shown in FIG. 4, automated clinical documentation process 10 may receive audio encounter information from a second microphone system (e.g., mobile electronic device 416). As shown in FIG. 4, the lines with dashes and dots between mobile electronic device 516 and VAD module 414; between VAD module 414 and alignment module 412; and between alignment module 412 and device selection and weighting module 410 may represent first audio stream received from microphone array 200. Referring again to FIG. 18 and in some implementations, automated clinical documentation process 10 may receive audio encounter information with one or more portions (e.g., portions 1828, 1830, 1832, 1834, 1836, 1838, 1840, 1842, 1844, 1846, 1848, 1850, 1852, 1854) from a second microphone system (e.g., mobile electronic device 416), thus defining a second audio stream.

[00158] In some implementations, automated clinical documentation process 10 may detect 1704 speech activity in one or more portions of the first audio stream, thus defining one or more speech portions of the first audio stream. As discussed above and in some implementations, automated clinical documentation process 10 may identify speech activity within one or more portions of the audio encounter information based upon, at least in part, a correlation among the audio encounter information received from the microphone array. Referring also to the example of FIG. 9 and in some implementations, audio encounter information 106 may include a plurality of portions or frames of audio information. In some implementations, automated clinical documentation process 10 may determine a correlation among the audio encounter information received from microphone array 200. For example, automated clinical documentation process 10 may compare the

one or more portions of audio encounter information 106 to determine a degree of correlation between the audio encounter information present in each portion across the plurality of microphones of the microphone array. In some implementations and as known in the art, automated clinical documentation process 10 may perform various cross-correlation processes to determine a degree of similarity between the one or more portions of audio encounter information 106 across the plurality of microphones of microphone array 200.

[00159] For example, suppose automated clinical documentation process 10 receives just ambient noise (i.e., no speech and no directional noise sources). Automated clinical documentation process 10 may determine that the spectrum observed in each microphone channel will be different (i.e. uncorrelated at each microphone). However, suppose automated clinical documentation process 10 receives a speech or other ‘directional’ signal within audio encounter information. In this example, automated clinical documentation process 10 may determine that one or more portions of the audio encounter information (e.g., the portions of audio encounter information with speech components) are highly correlated at each microphone in the microphone array.

[00160] In some implementations, automated clinical documentation process 10 may identify speech activity within one or more portions of the audio encounter information based upon, at least in part, determining a threshold amount or degree of correlation among the audio encounter information received from the microphone array. For example and as discussed above, various thresholds may be defined (e.g., user-defined, default thresholds, automatically defined via automated clinical documentation process 10, etc.) to determine when portions of audio encounter information are sufficiently correlated. Accordingly, in response to determining at least a threshold degree of correlation among the portions of audio encounter information across the plurality of microphones, automated clinical documentation process 10 may determine or identify speech activity within the one

or more portions of the audio encounter information. In combination with determining a threshold correlation among the portions of audio encounter information, automated clinical documentation process 10 may identify speech activity using other approaches known in the art for voice activity detection (VAD) such as filtering, noise reduction, the application of classification rules, etc. In this manner, conventional VAD techniques may be used in combination with the determination of a threshold correlation among the one or more portions of the audio encounter information to identify speech activity within the one or more portions of the audio encounter information. Referring again to FIG. 18, automated clinical documentation process 10 may detect speech activity in portions 1800, 1802, 1804, 1806, 1808, 1810, 1812.

[00161] In some implementations, automated clinical documentation process 10 may detect 1706 speech activity in one or more portions of the second audio stream, thus defining one or more speech portions of the second audio stream. For example and in some implementations, automated clinical documentation process 10 may perform various known voice activity detection (VAD) processes to determine which portions of audio encounter information received by mobile electronic device 416 include speech activity. In this manner, automated clinical documentation process 10 may detect 1706 speech activity in the one or more portions of the second audio stream. Referring again to FIG. 18, automated clinical documentation process 10 may detect speech activity in portions 1830, 1832, 1834, 1836, 1838, 1840, 1842 using various known VAD processes.

[00162] In some implementations, automated clinical documentation process 10 may align 1708 the first audio stream and the second audio stream based upon, at least in part, the one or more speech portions of the first audio stream and the one or more speech portions of the second audio stream. In some implementations, automated clinical documentation process 10 may utilize an echo cancellation-based adaptive filtering approach where the delay between the first audio stream and the second audio stream is

estimated from a filter bulk delay and the sparseness of the filter is used to ensure that the two audio streams can be aligned. Referring also to FIG. 19 and in some implementations, automated clinical documentation process 10 may align first audio stream 1900 and second audio stream 1902 based upon, at least in part, the one or more speech portions of the first audio stream (e.g., speech portions 1800, 1802, 1804, 1806, 1808, 1810, 1812 of first audio stream 1900) and the one or more speech portions of the second audio stream (e.g., speech portions 1830, 1832, 1834, 1836, 1838, 1840, 1842 of second audio stream 1902). In the example of FIG. 19, while each audio stream may have different amplitude values at a given point in time, the first and second audio streams may be temporally-aligned.

[00163] In some implementations and in response to aligning the first audio stream and the second audio stream, automated clinical documentation process 10 may process 1710 the first audio stream and the second audio stream with one or more speech processing systems. Referring also to FIGS. 20A-20B, automated clinical documentation process 10 may provide the first audio stream and the second audio stream from the device selection and weighting module (e.g., device selection and weighting module 410) to one or more speech processing systems (e.g., speech processing system 418, 2000, 2002). Examples of speech processing systems may generally include automated speech recognition (ASR) systems, voice biometric systems, emotion detection systems, medical symptom detection systems, hearing enhancement systems, etc. As will be discussed in greater detail below, because different portions of the audio stream may be processed more accurately by the one or more speech processing systems based upon, at least in part, the signal quality of each audio stream. Accordingly, automated clinical documentation process 10 may selectively process 1701 particular portions from each audio stream with the one or more speech processing systems.

[00164] In some implementations, processing 1710 the first audio stream and the second audio stream with one or more speech processing systems may include weighting

1712 the first audio stream and the second audio stream based upon, at least in part, a signal-to-noise ratio for the first audio stream and a signal-to-noise ratio for the second audio stream, thus defining a first audio stream weight and a second audio stream weight. In some implementations, device selection and weighting module 410 may be a machine learning system or model (e.g., a neural network-based model) configured to be trained to weight each audio stream such that the estimated speech processing system performance is maximized. For example, automated clinical documentation process 10 may train the device selection and weighting model based upon, at least in part, the signal-to-noise (SNR) ratio for each portion or frame of each audio stream. In some implementations and as will be discussed in greater detail below, the machine learning model of the device selection and weighting module may be trained jointly in an end-to-end manner to weight and select particular portions of the audio stream for processing with one or more speech processing systems.

[00165] Referring again to the example of FIG. 6, suppose a doctor (e.g., doctor 226) has mobile electronic device 416 near them as they are speaking to a patient (e.g., patient 228). In this example, a first microphone system (e.g., microphone array 200) and a second microphone system (e.g., mobile electronic device 416) may receive audio encounter information while the doctor (e.g., encounter participant 226) is speaking. As discussed above, automated clinical documentation process 10 may detect speech activity within each audio stream and may align the audio streams based upon, at least in part, the speech activity within each audio stream. Automated clinical documentation process 10 may weight 1712 each portion or frame of each audio stream based upon, at least in part, a SNR ratio for each audio stream. In this example, suppose the doctor (e.g., encounter participant 226) is speaking when near the doctor's desk. Suppose the doctor (e.g., encounter participant 226) moves quickly to the examination table to assist the patient (e.g., patient 228). During this period, the doctor (e.g., encounter participant 226) may be speaking. As

such, while the doctor (e.g., encounter participant 226) may move outside of beam 222 towards null 248, microphone array 200 may receive lower quality audio encounter information than mobile electronic device 416 (i.e., the audio encounter information received by microphone array 200 may have a lower SNR ratio). Accordingly, automated clinical documentation process 10 may weight 1712 the portions of the audio encounter information received from microphone array 200 while the doctor is moving with a lower weight than the audio encounter information received from mobile electronic device 416. Additionally / alternatively, automated clinical documentation process 10 may weight 1712 the portions of audio encounter information received from mobile electronic device 416 while the doctor is moving with a higher weight than the audio encounter information received from microphone array 200. In this manner, automated clinical documentation process 10 may utilize the weighting of portions of each audio stream to automated clinical documentation process 101 each audio stream with one or more speech processing systems.

[00166] In some implementations, processing 1710 the first audio stream and the second audio stream with one or more speech processing systems may include processing 1714 the first audio stream and the second audio stream with a single speech processing system based upon, at least in part, the first audio stream weight and the second audio stream weight. Referring again to FIG. 20A and in some implementations, automated clinical documentation process 10 may automated clinical documentation process 101 the first audio stream (e.g., represented as a solid line between device selection weighting module 410 and speech processing system 418) and the second audio stream (e.g., represented as a dashed and dotted line between device selection weighting module 410 and speech processing system 418) with a single speech processing system (e.g., speech processing system 418) based upon, at least in part, the first audio stream weight and the second audio stream weight (e.g., where both audio stream weights are represented as the dashed line between the device selection weighting module 410 and speech processing

system 418). In some implementations, automated clinical documentation process 10 may, via speech processing system 418, select particular portions of either audio stream to process based upon, at least in part, the first audio stream weight and the second audio stream weight.

[00167] In some implementations, processing 1710 the first audio stream and the second audio stream with one or more speech processing systems may include processing 1716 the first audio stream with a first speech processing system, thus defining a first speech processing output; processing 1718 the second audio stream with a second speech processing system, thus defining a second speech processing output; and combining 1720 the first speech processing output with the second speech processing output based upon, at least in part, the first audio stream weight and the second audio stream weight. Referring again to FIG. 20B and in some implementations, automated clinical documentation process 10 may process 1716 the first audio stream (e.g., represented as a solid line between device selection weighting module 410 and speech processing system 418) with a first speech processing system to generate a first speech processing output (e.g., represented as a solid line between speech processing system 418 and speech processing system 2002). Automated clinical documentation process 10 may process 1718 the second audio stream (e.g., represented as a dashed and dotted line between device selection weighting module 410 and speech processing system 2000) with a second speech processing system (e.g., speech processing system 2000) to generate a second speech processing output (e.g., represented as a solid line between speech processing system 2000 and speech processing system 2002). In some implementations, automated clinical documentation process 10 may combine 1720, via a third speech processing system (e.g., speech processing system 2002), the first speech processing output with the second speech processing output based upon, at least in part, the first audio stream weight and the second audio stream weight (e.g., where both audio stream weights are represented as the dashed line between the

device selection weighting module 410 and speech processing system 2002). In some implementations, automated clinical documentation process 10 may, via speech processing system 2002, select particular portions of either audio stream to process or output based upon, at least in part, the first audio stream weight and the second audio stream weight. While an example of e.g., two audio streams has been provided, it will be appreciated that any number of audio streams may be used within the scope of the present disclosure.

[00168] In some implementations, automated clinical documentation process 10 may generate an encounter transcript (e.g., encounter transcript 234) with output of the one or more speech processing systems, wherein at least a portion of the encounter transcript (e.g., encounter transcript 234) may be processed to populate at least a portion of a medical record (e.g., medical record 236) associated with the patient encounter (e.g., a visit to a doctor's office). For example, automated clinical documentation process 10 may utilize speaker representation to identify speakers in audio encounter information when generating the encounter transcript. For example, automated clinical documentation process 10 may generate a diarized encounter transcript (e.g., encounter transcript 234) that identifies the verbal comments and utterances made by particular speakers based upon, at least in part, the speaker representations defined for each encounter participant. In the above example, automated clinical documentation process 10 may generate diarized encounter transcript 234 with verbal comments and utterances may by "Doctor Susan Jones" (e.g., encounter participant 226), "Patient Paul Smith" (e.g., encounter participant 228), and "Unknown Participant" (e.g., encounter participant 230).

General:

[00169] As will be appreciated by one skilled in the art, the present disclosure may be embodied as a method, a system, or a computer program product. Accordingly, the present disclosure may take the form of an entirely hardware embodiment, an entirely software embodiment (including firmware, resident software, micro-code, etc.) or an

embodiment combining software and hardware aspects that may all generally be referred to herein as a “circuit,” “module” or “system.” Furthermore, the present disclosure may take the form of a computer program product on a computer-usable storage medium having computer-usable program code embodied in the medium.

[00170] Any suitable computer usable or computer readable medium may be utilized. The computer-usable or computer-readable medium may be, for example but not limited to, an electronic, magnetic, optical, electromagnetic, infrared, or semiconductor system, apparatus, device, or propagation medium. More specific examples (a non-exhaustive list) of the computer-readable medium may include the following: an electrical connection having one or more wires, a portable computer diskette, a hard disk, a random access memory (RAM), a read-only memory (ROM), an erasable programmable read-only memory (EPROM or Flash memory), an optical fiber, a portable compact disc read-only memory (CD-ROM), an optical storage device, a transmission media such as those supporting the Internet or an intranet, or a magnetic storage device. The computer-usable or computer-readable medium may also be paper or another suitable medium upon which the program is printed, as the program can be electronically captured, via, for instance, optical scanning of the paper or other medium, then compiled, interpreted, or otherwise processed in a suitable manner, if necessary, and then stored in a computer memory. In the context of this document, a computer-usable or computer-readable medium may be any medium that can contain, store, communicate, propagate, or transport the program for use by or in connection with the instruction execution system, apparatus, or device. The computer-usable medium may include a propagated data signal with the computer-usable program code embodied therewith, either in baseband or as part of a carrier wave. The computer usable program code may be transmitted using any appropriate medium, including but not limited to the Internet, wireline, optical fiber cable, RF, etc.

[00171] Computer program code for carrying out operations of the present

disclosure may be written in an object oriented programming language such as Java, Smalltalk, C++ or the like. However, the computer program code for carrying out operations of the present disclosure may also be written in conventional procedural programming languages, such as the “C” programming language or similar programming languages. The program code may execute entirely on the user’s computer, partly on the user’s computer, as a stand-alone software package, partly on the user’s computer and partly on a remote computer or entirely on the remote computer or server. In the latter scenario, the remote computer may be connected to the user’s computer through a local area network / a wide area network / the Internet (e.g., network 14).

[00172] The present disclosure is described with reference to flowchart illustrations and/or block diagrams of methods, apparatus (systems) and computer program products according to embodiments of the disclosure. It will be understood that each block of the flowchart illustrations and/or block diagrams, and combinations of blocks in the flowchart illustrations and/or block diagrams, may be implemented by computer program instructions. These computer program instructions may be provided to a processor of a general purpose computer / special purpose computer / other programmable data processing apparatus, such that the instructions, which execute via the processor of the computer or other programmable data processing apparatus, create means for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks.

[00173] These computer program instructions may also be stored in a computer-readable memory that may direct a computer or other programmable data processing apparatus to function in a particular manner, such that the instructions stored in the computer-readable memory produce an article of manufacture including instruction means which implement the function/act specified in the flowchart and/or block diagram block or blocks.

[00174] The computer program instructions may also be loaded onto a computer or other programmable data processing apparatus to cause a series of operational steps to be performed on the computer or other programmable apparatus to produce a computer implemented process such that the instructions which execute on the computer or other programmable apparatus provide steps for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks.

[00175] The flowcharts and block diagrams in the figures may illustrate the architecture, functionality, and operation of possible implementations of systems, methods and computer program products according to various embodiments of the present disclosure. In this regard, each block in the flowchart or block diagrams may represent a module, segment, or portion of code, which comprises one or more executable instructions for implementing the specified logical function(s). It should also be noted that, in some alternative implementations, the functions noted in the block may occur out of the order noted in the figures. For example, two blocks shown in succession may, in fact, be executed substantially concurrently, or the blocks may sometimes be executed in the reverse order, not at all, or in any combination with any other flowcharts depending upon the functionality involved. It will also be noted that each block of the block diagrams and/or flowchart illustrations, and combinations of blocks in the block diagrams and/or flowchart illustrations, may be implemented by special purpose hardware-based systems that perform the specified functions or acts, or combinations of special purpose hardware and computer instructions.

[00176] The terminology used herein is for the purpose of describing particular embodiments only and is not intended to be limiting of the disclosure. As used herein, the singular forms “a”, “an” and “the” are intended to include the plural forms as well, unless the context clearly indicates otherwise. It will be further understood that the terms “comprises” and/or “comprising,” when used in this specification, specify the presence of

stated features, integers, steps, operations, elements, and/or components, but do not preclude the presence or addition of one or more other features, integers, steps, operations, elements, components, and/or groups thereof.

[00177] The corresponding structures, materials, acts, and equivalents of all means or step plus function elements in the claims below are intended to include any structure, material, or act for performing the function in combination with other claimed elements as specifically claimed. The description of the present disclosure has been presented for purposes of illustration and description, but is not intended to be exhaustive or limited to the disclosure in the form disclosed. Many modifications and variations will be apparent to those of ordinary skill in the art without departing from the scope and spirit of the disclosure. The embodiment was chosen and described in order to best explain the principles of the disclosure and the practical application, and to enable others of ordinary skill in the art to understand the disclosure for various embodiments with various modifications as are suited to the particular use contemplated.

[00178] A number of implementations have been described. Having thus described the disclosure of the present application in detail and by reference to embodiments thereof, it will be apparent that modifications and variations are possible without departing from the scope of the disclosure defined in the appended claims.

What Is Claimed Is:

1. A computer-implemented method, executed on a computing device, comprising:
 - receiving a plurality of predefined beams associated with a microphone array;
 - receiving a plurality of predefined nulls associated with the microphone array;
 - selecting at least one of:
 - one or more predefined beams from the plurality of predefined beams, thus defining one or more selected beams; and
 - one or more predefined nulls from the plurality of predefined nulls, thus defining one or more selected nulls; and
 - obtaining audio encounter information, via the microphone array, using at least one of the one or more selected beams and the one or more selected nulls.
2. The computer-implemented method of claim 1, wherein the plurality of predefined beams include one or more predefined beams configured to receive audio encounter information from one or more target speaker locations within an acoustic environment.
3. The computer-implemented method of claim 2, wherein the plurality of predefined nulls include one or more predefined nulls configured to limit receiving audio encounter information from the one or more target speaker locations within the acoustic environment.
4. The computer-implemented method of claim 3, further comprising:
 - receiving speaker metadata associated with one or more portions of audio encounter information.

5. The computer-implemented method of claim 4, wherein selecting at least one of the one or more beams and the one or more nulls includes selecting at least one of the one or more beams and the one or more nulls based upon, at least in part, the speaker metadata associated with the one or more portions of the audio encounter information.

6. The computer-implemented method of claim 5, wherein selecting at least one of the one or more beams and the one or more nulls based upon, at least in part, the speaker metadata associated with the one or more portions of the audio encounter information includes selecting the one or more predefined beams configured to receive audio encounter information from the one or more target speaker locations within the acoustic environment based upon, at least in part, the speaker metadata associated with the one or more portions of the audio encounter information.

7. The computer-implemented method of claim 5, wherein selecting at least one of the one or more beams and the one or more nulls based upon, at least in part, the speaker metadata associated with the one or more portions of the audio encounter information includes selecting the one or more predefined nulls configured to limit receiving audio encounter information from the one or more target speaker locations within the acoustic environment based upon, at least in part, the speaker metadata associated with the one or more portions of the audio encounter information.

8. A computer program product residing on a non-transitory computer readable medium having a plurality of instructions stored thereon which, when executed by a processor, cause the processor to perform operations comprising:

receiving a plurality of predefined beams associated with a microphone array;

receiving a plurality of predefined nulls associated with the microphone array;

selecting at least one of:

one or more predefined beams from the plurality of predefined beams, thus defining one or more selected beams; and

one or more predefined nulls from the plurality of predefined nulls, thus defining one or more selected nulls; and

obtaining audio encounter information, via the microphone array, using at least one of the one or more selected beams and the one or more selected nulls.

9. The computer program product of claim 8, wherein the plurality of predefined beams include one or more predefined beams configured to receive audio encounter information from one or more target speaker locations within an acoustic environment.

10. The computer program product of claim 9, wherein the plurality of predefined nulls include one or more predefined nulls configured to limit receiving audio encounter information from the one or more target speaker locations within the acoustic environment.

11. The computer program product of claim 10, wherein the operations further comprise:

receiving speaker metadata associated with one or more portions of audio encounter information.

12. The computer program product of claim 11, wherein selecting at least one of the one or more beams and the one or more nulls includes selecting at least one of the one or more beams and the one or more nulls based upon, at least in part, the speaker metadata

associated with the one or more portions of the audio encounter information.

13. The computer program product of claim 12, wherein selecting at least one of the one or more beams and the one or more nulls based upon, at least in part, the speaker metadata associated with the one or more portions of the audio encounter information includes selecting the one or more predefined beams configured to receive audio encounter information from the one or more target speaker locations within the acoustic environment based upon, at least in part, the speaker metadata associated with the one or more portions of the audio encounter information.

14. The computer program product of claim 12, wherein selecting at least one of the one or more beams and the one or more nulls based upon, at least in part, the speaker metadata associated with the one or more portions of the audio encounter information includes selecting the one or more predefined nulls configured to limit receiving audio encounter information from the one or more target speaker locations within the acoustic environment based upon, at least in part, the speaker metadata associated with the one or more portions of the audio encounter information.

15. A computing system comprising:
a memory; and
a processor configured to receive a plurality of predefined beams associated with a microphone array, wherein the processor is further configured to receive a plurality of predefined nulls associated with the microphone array, wherein the processor is further configured to select at least one of: one or more predefined beams from the plurality of predefined beams, thus defining one or more selected beams and one or more predefined nulls from the plurality of predefined nulls, thus

defining one or more selected nulls, and wherein the processor is further configured to obtain audio encounter information, via the microphone array, using at least one of the one or more selected beams and the one or more selected nulls.

16. The computing system of claim 15, wherein the plurality of predefined beams include one or more predefined beams configured to receive audio encounter information from one or more target speaker locations within an acoustic environment.

17. The computing system of claim 16, wherein the plurality of predefined nulls include one or more predefined nulls configured to limit receiving audio encounter information from the one or more target speaker locations within the acoustic environment.

18. The computing system of claim 17, wherein the processor is further configured to:
receive speaker metadata associated with one or more portions of audio encounter information.

19. The computing system of claim 18, wherein selecting at least one of the one or more beams and the one or more nulls includes selecting at least one of the one or more beams and the one or more nulls based upon, at least in part, the speaker metadata associated with the one or more portions of the audio encounter information.

20. The computing system of claim 19, wherein selecting at least one of the one or more beams and the one or more nulls based upon, at least in part, the speaker metadata associated with the one or more portions of the audio encounter information includes selecting the one or more predefined beams configured to receive audio encounter information from the one or more target speaker locations within the acoustic environment

based upon, at least in part, the speaker metadata associated with the one or more portions of the audio encounter information.

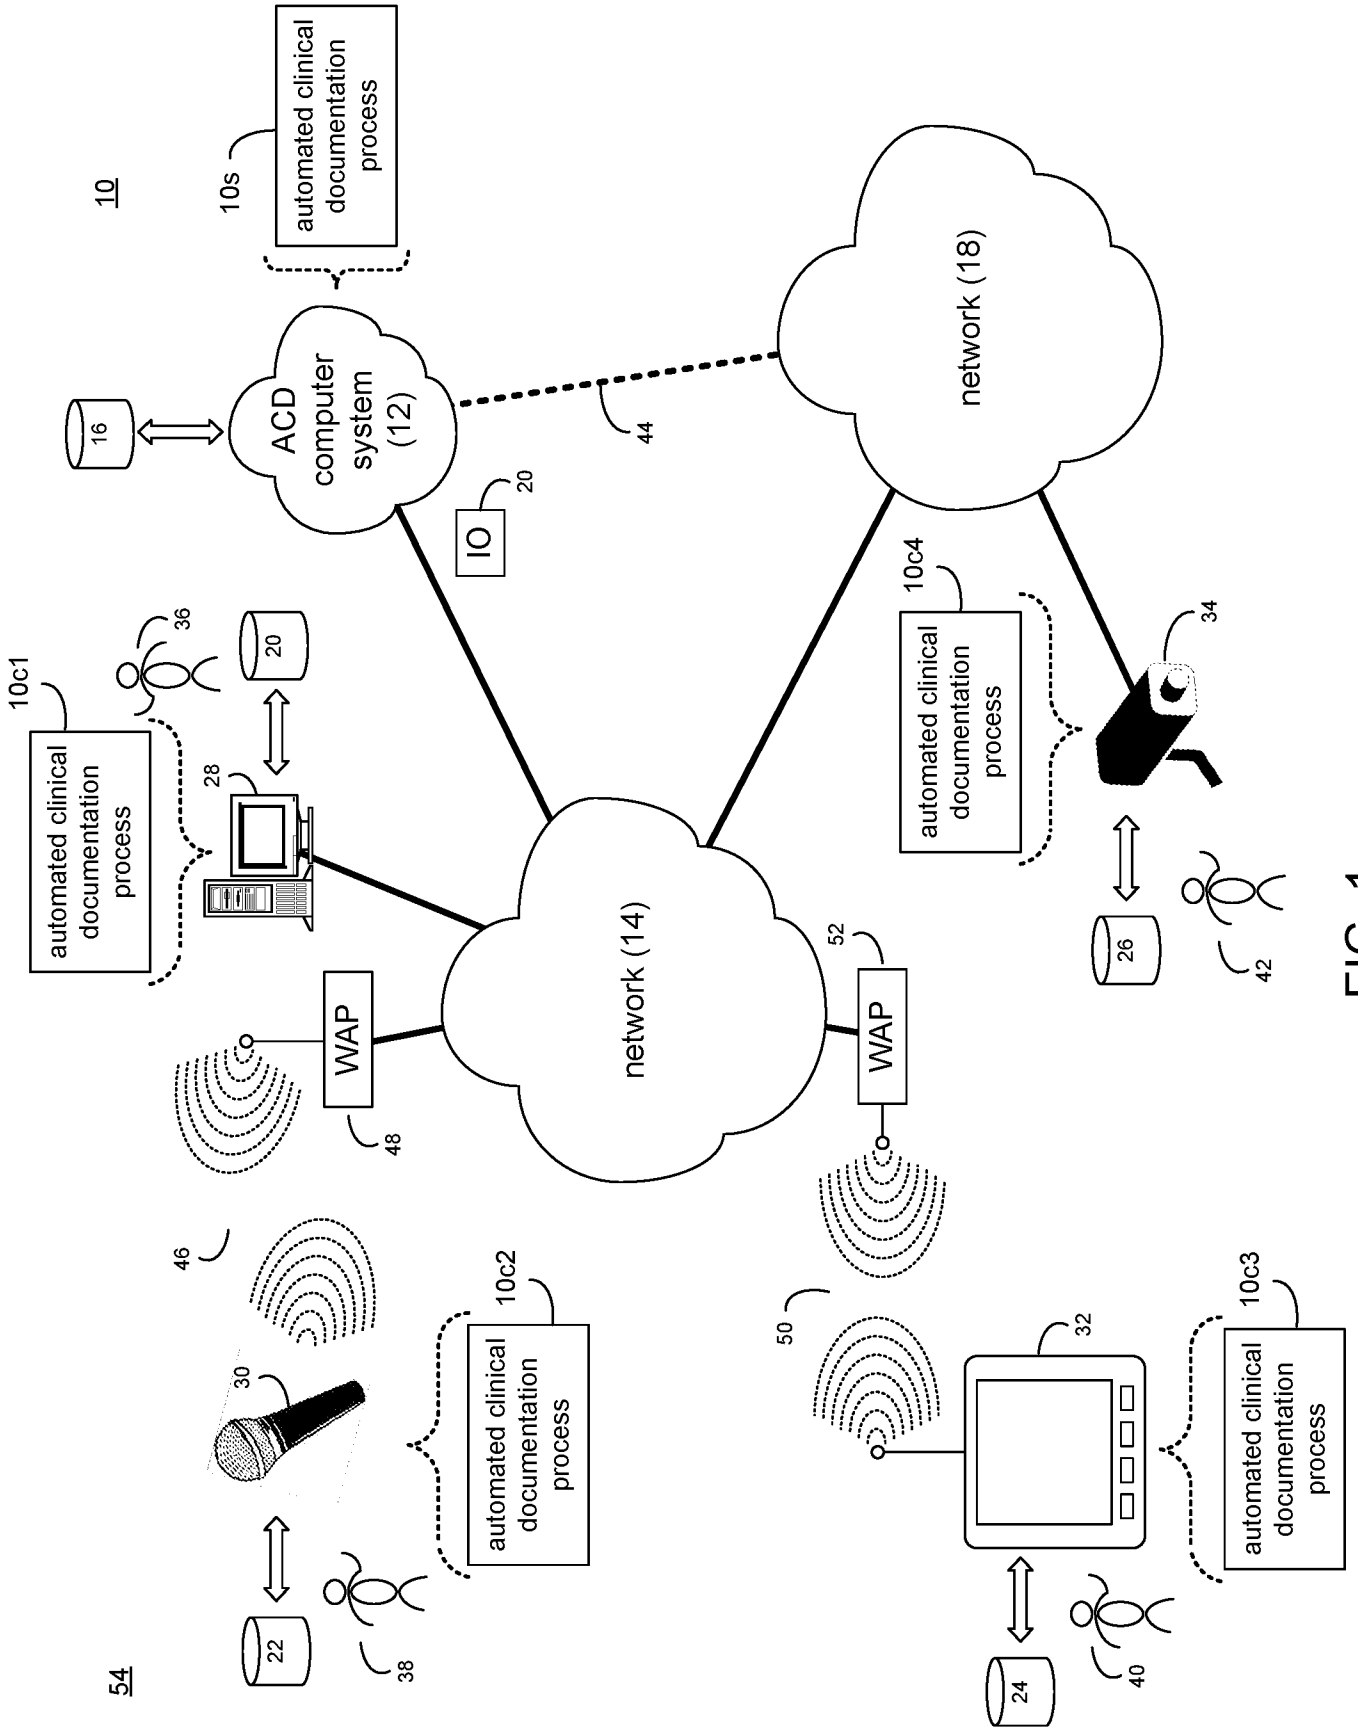


FIG. 1

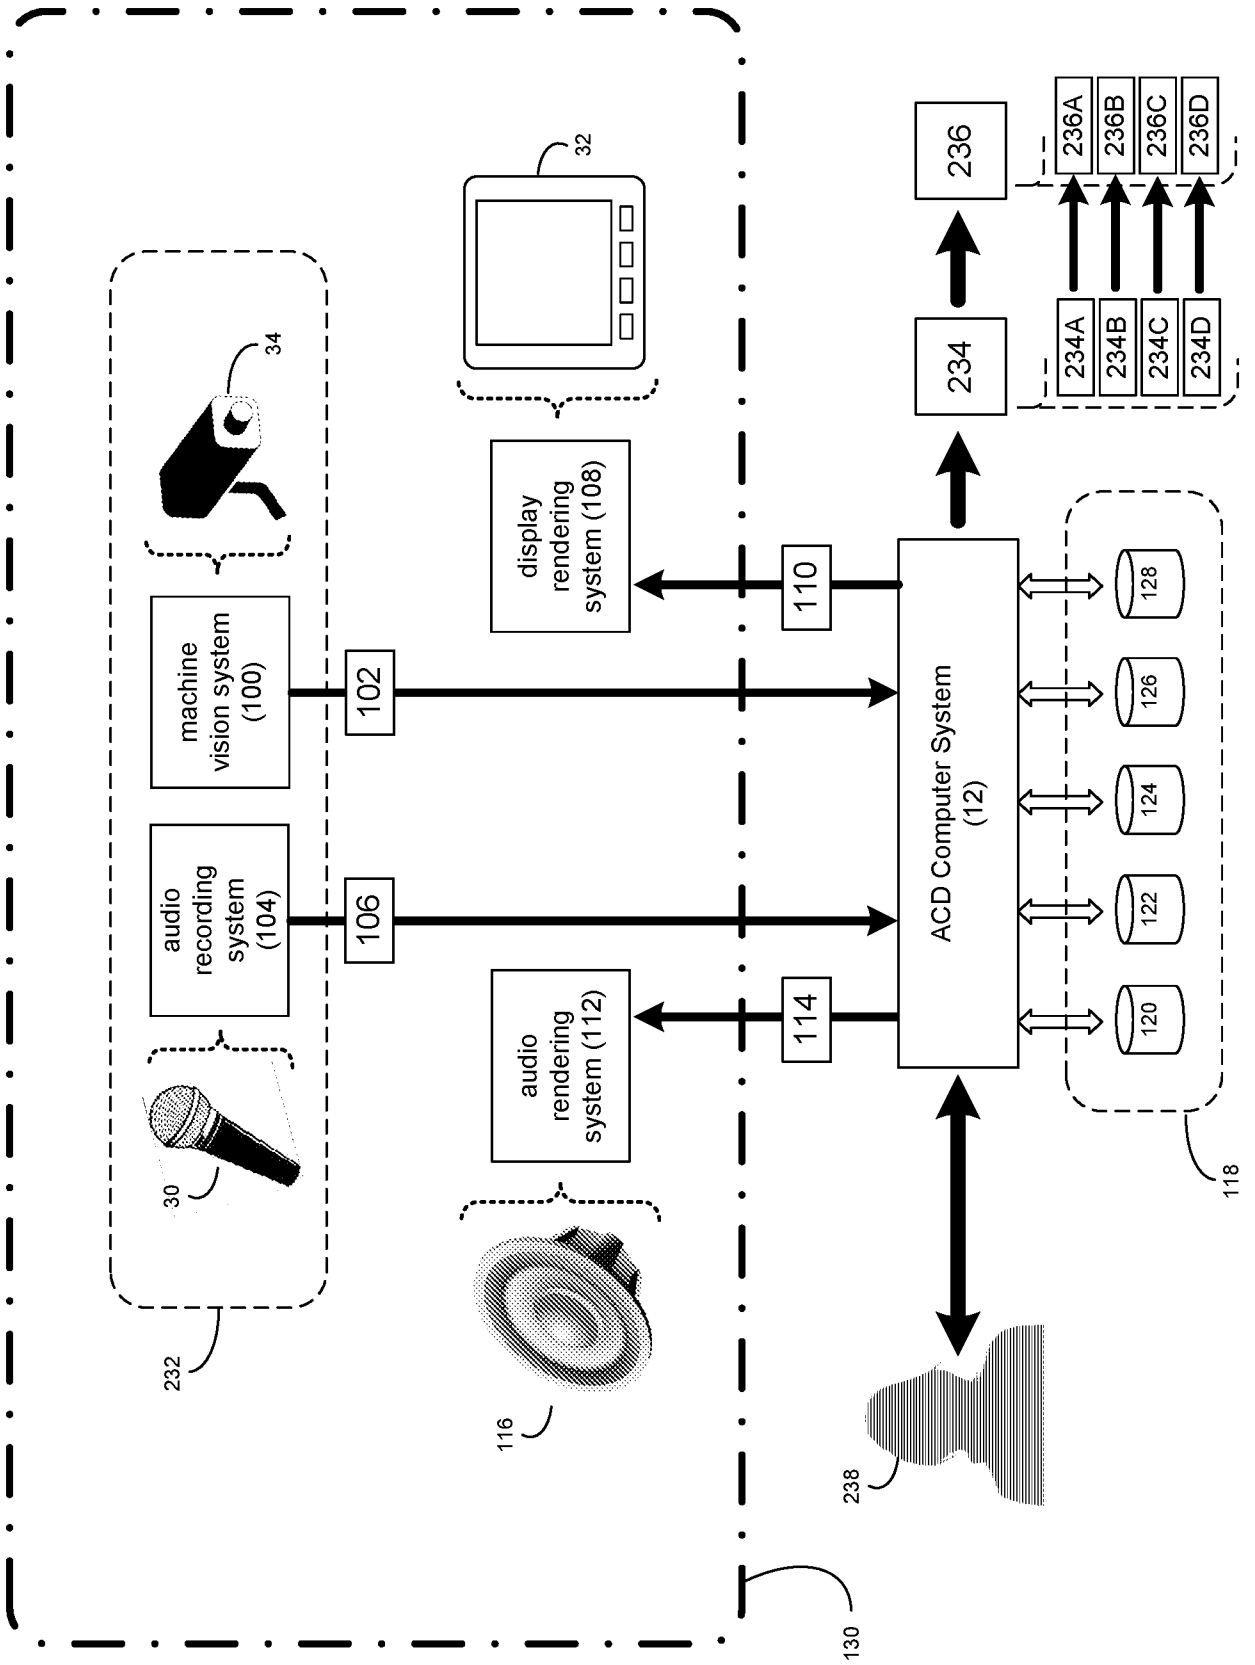


FIG. 2




FIG. 3

10

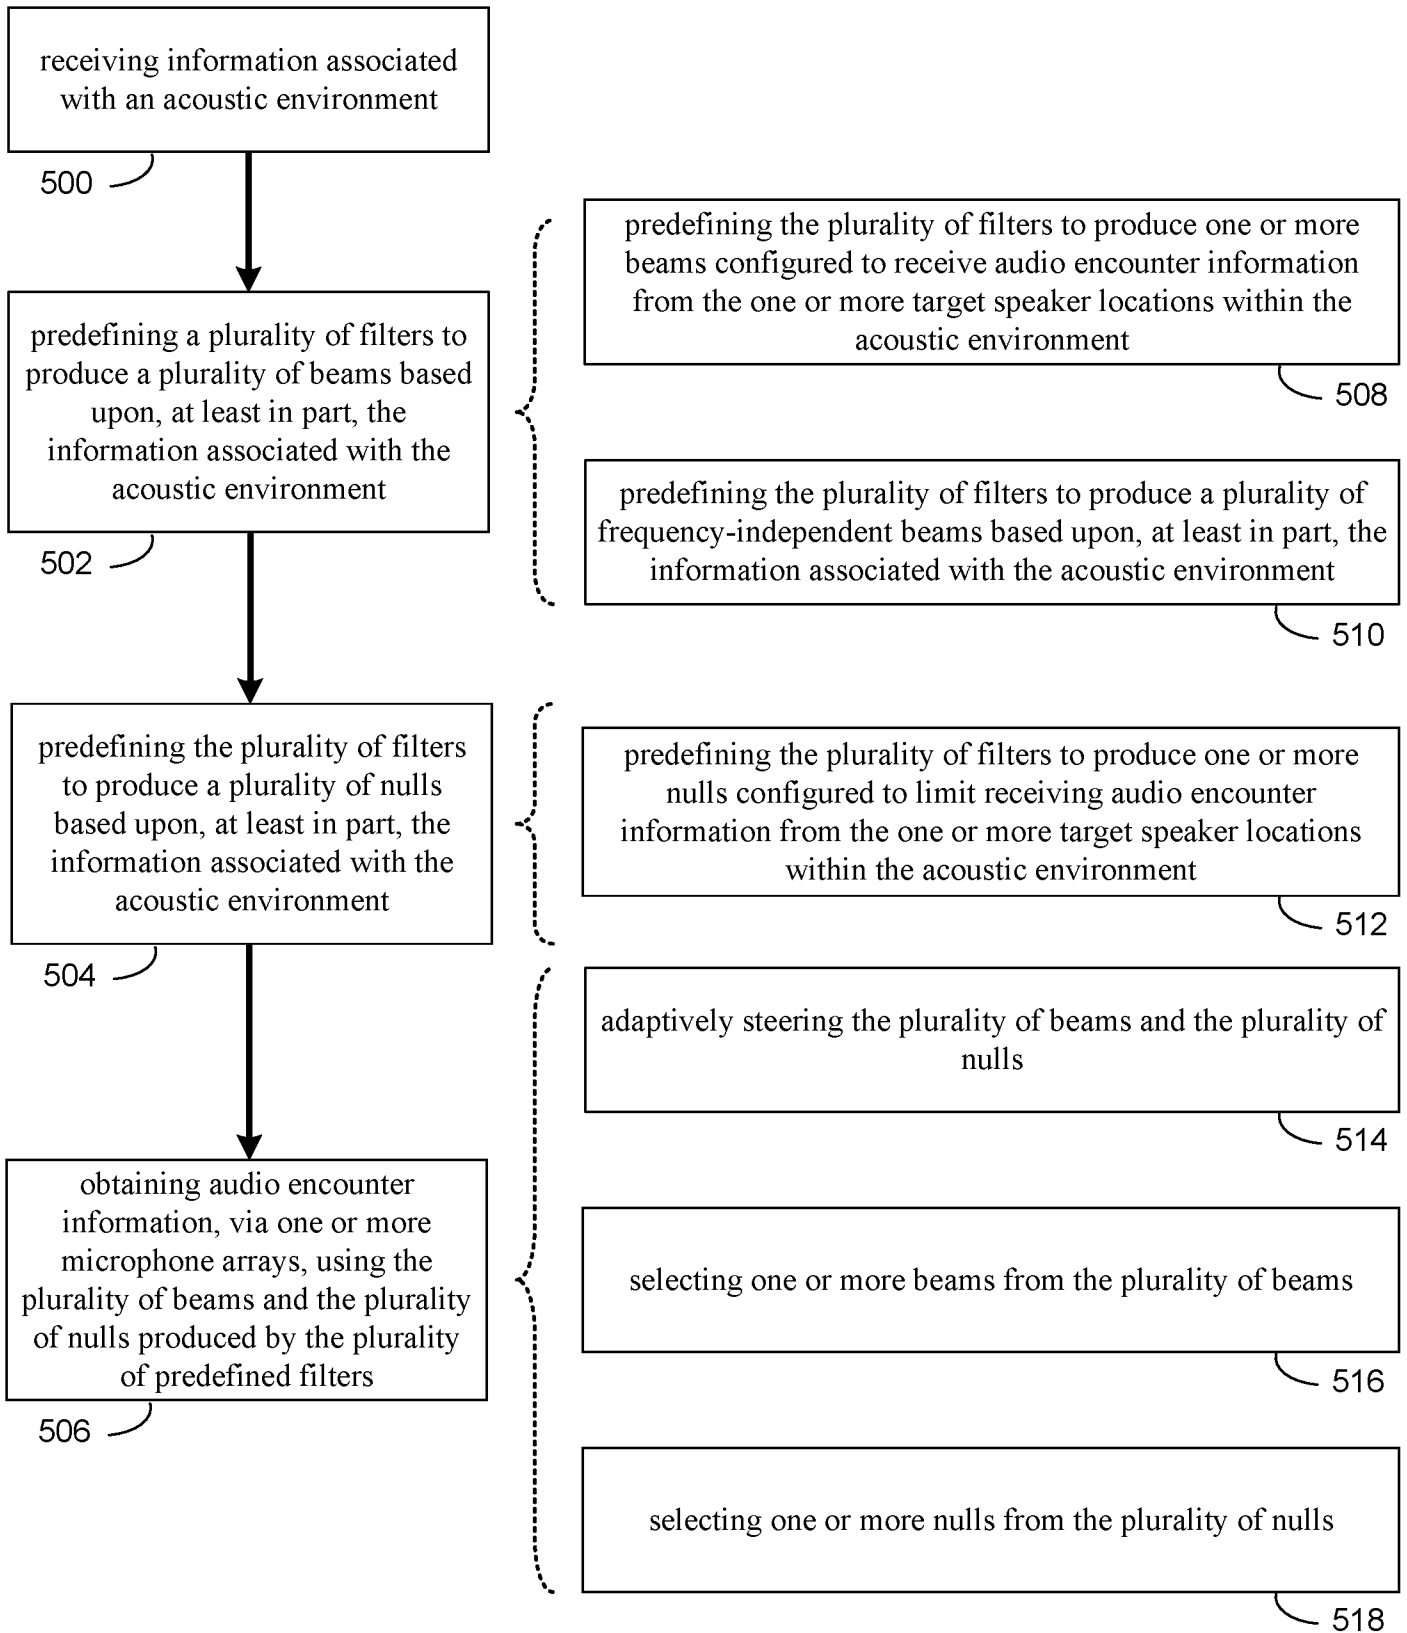


FIG. 5




FIG. 6




FIG. 7

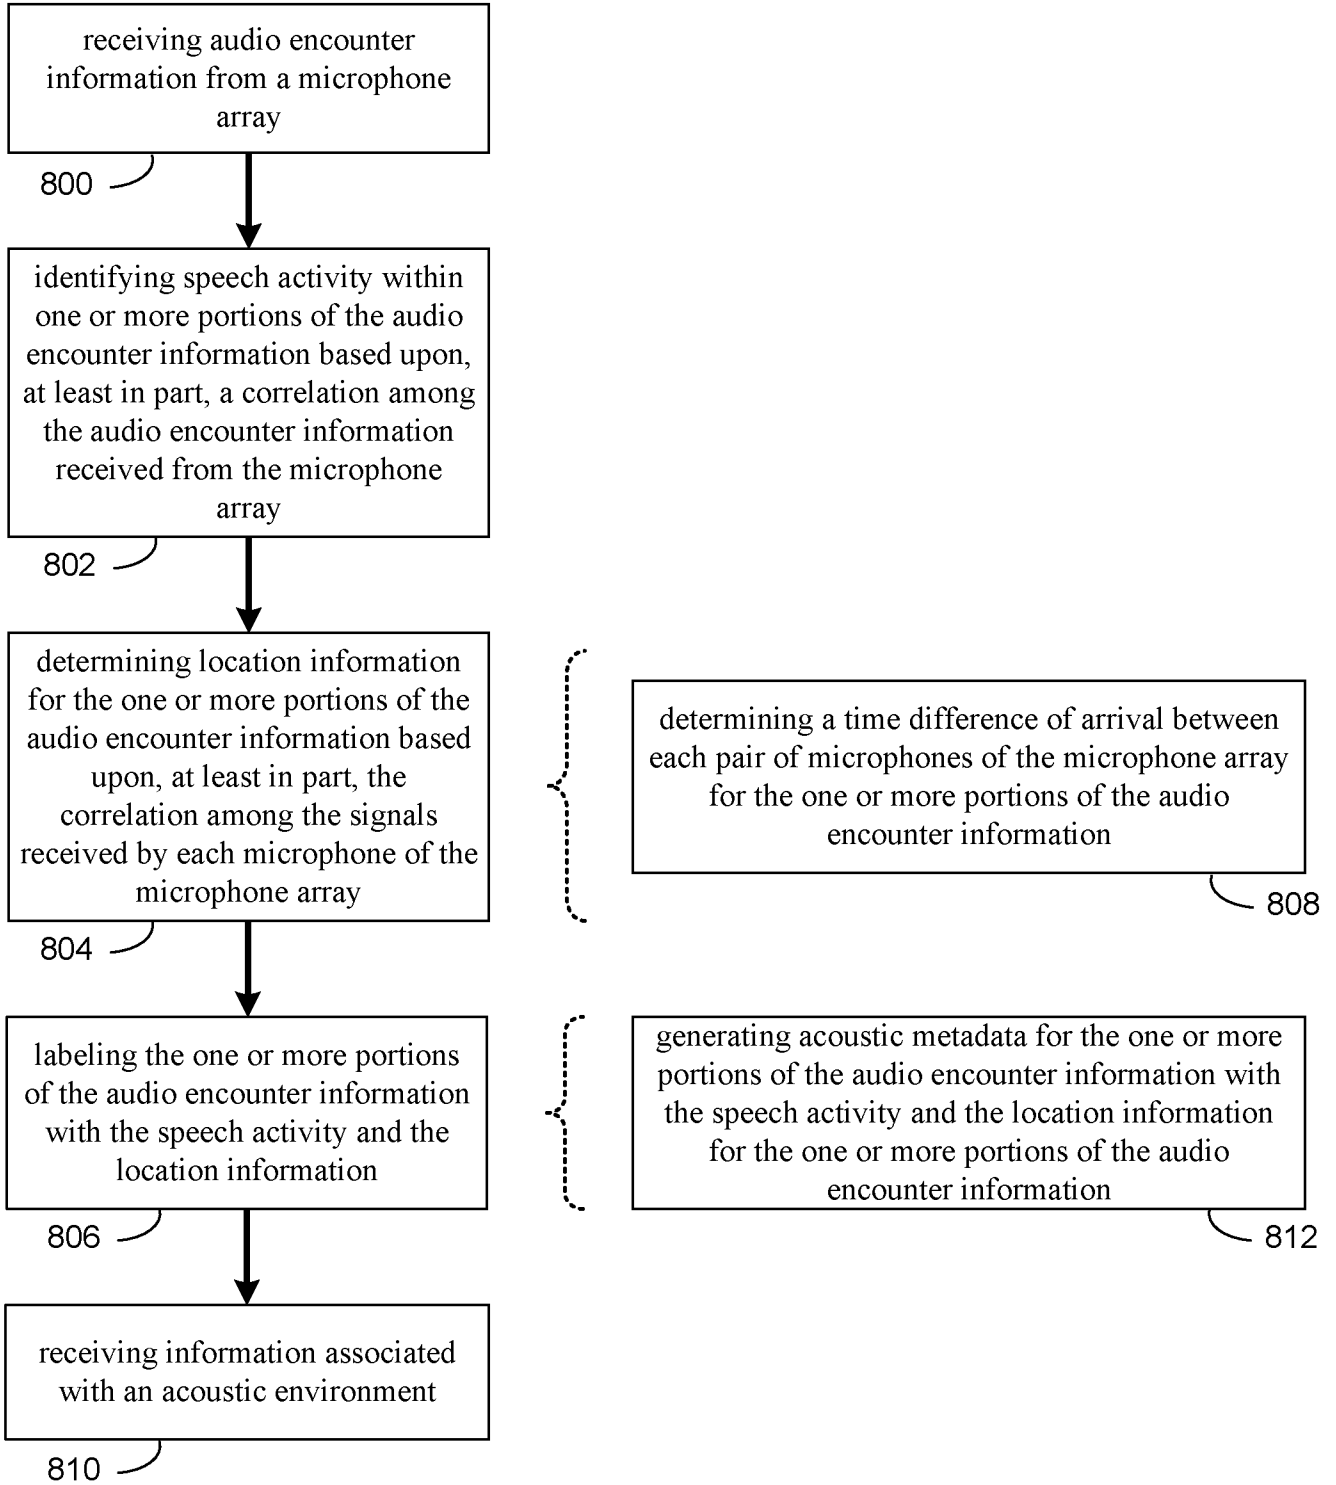


FIG. 8

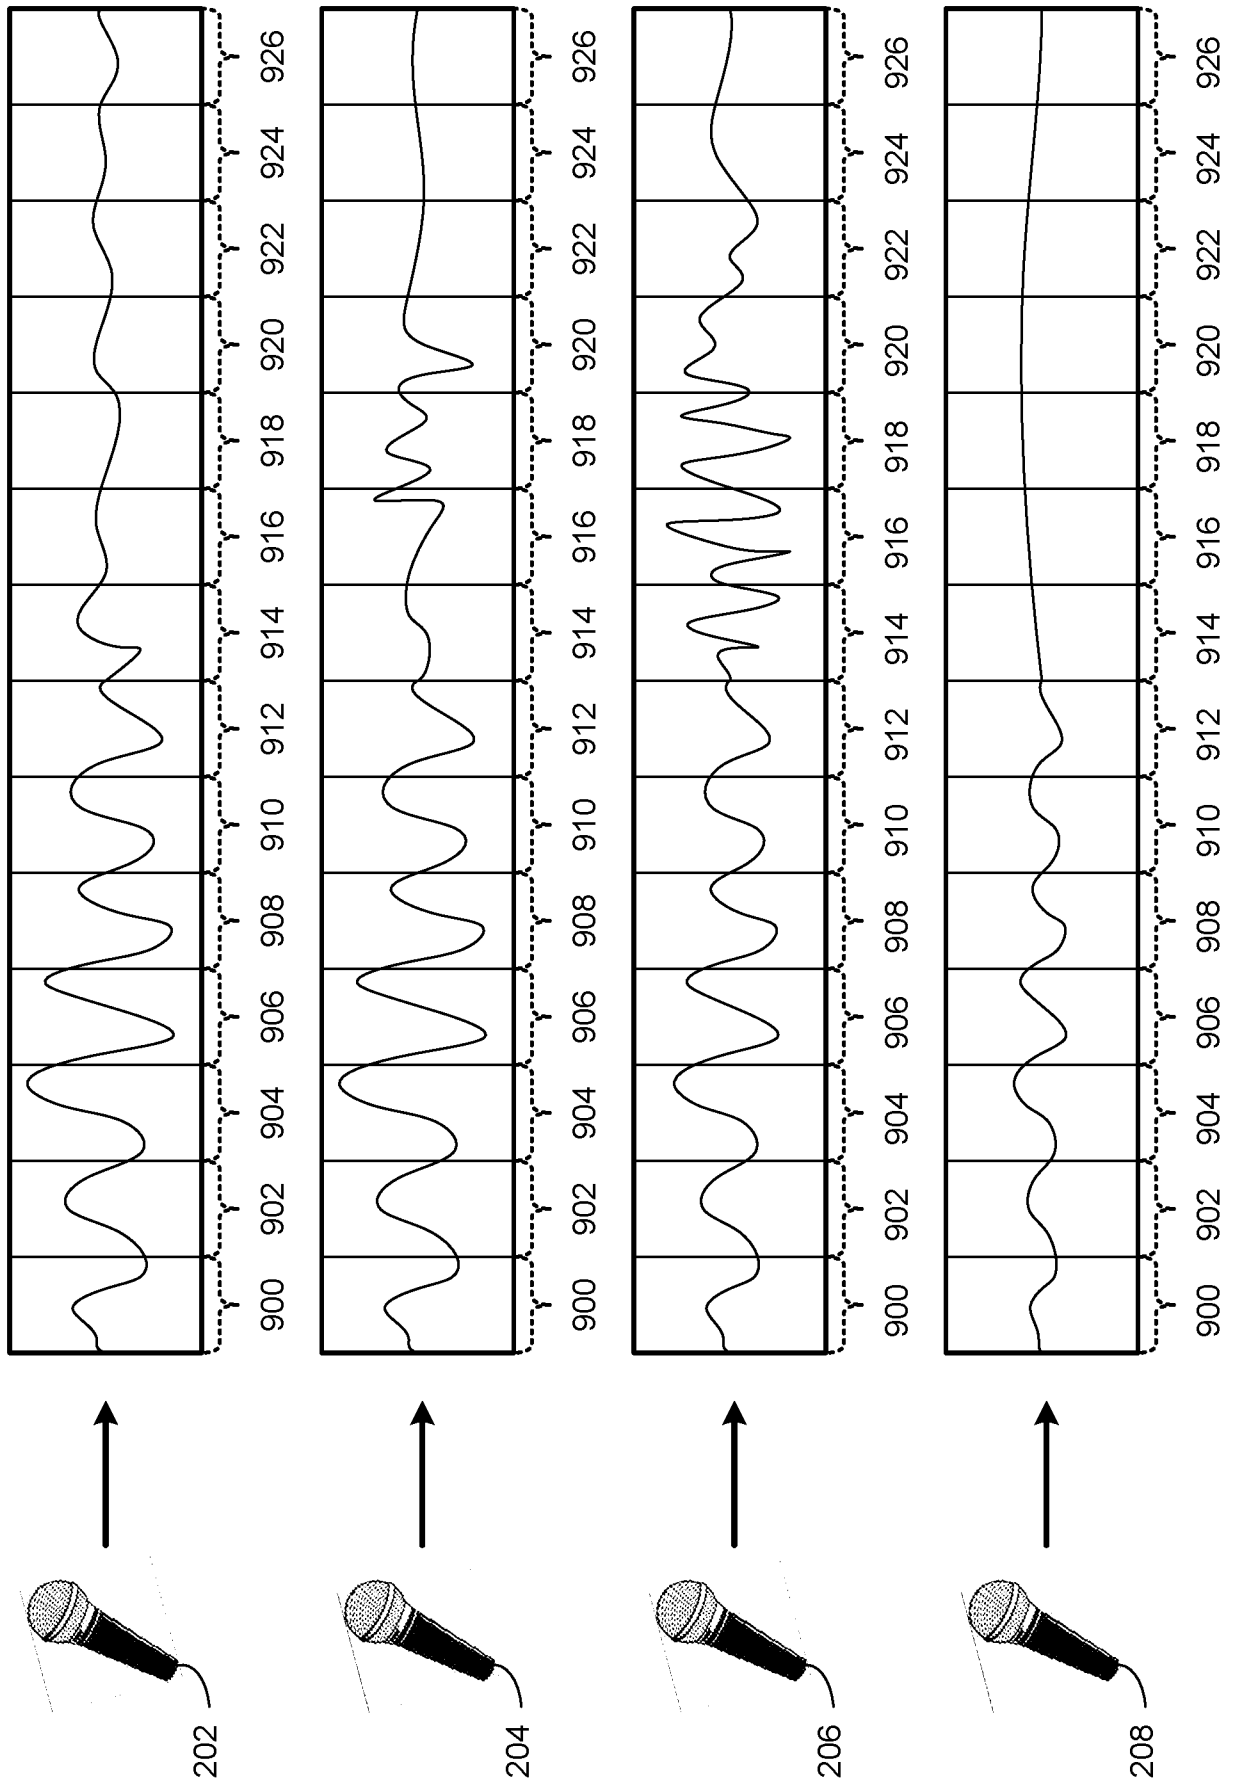


FIG. 9




FIG. 10

10

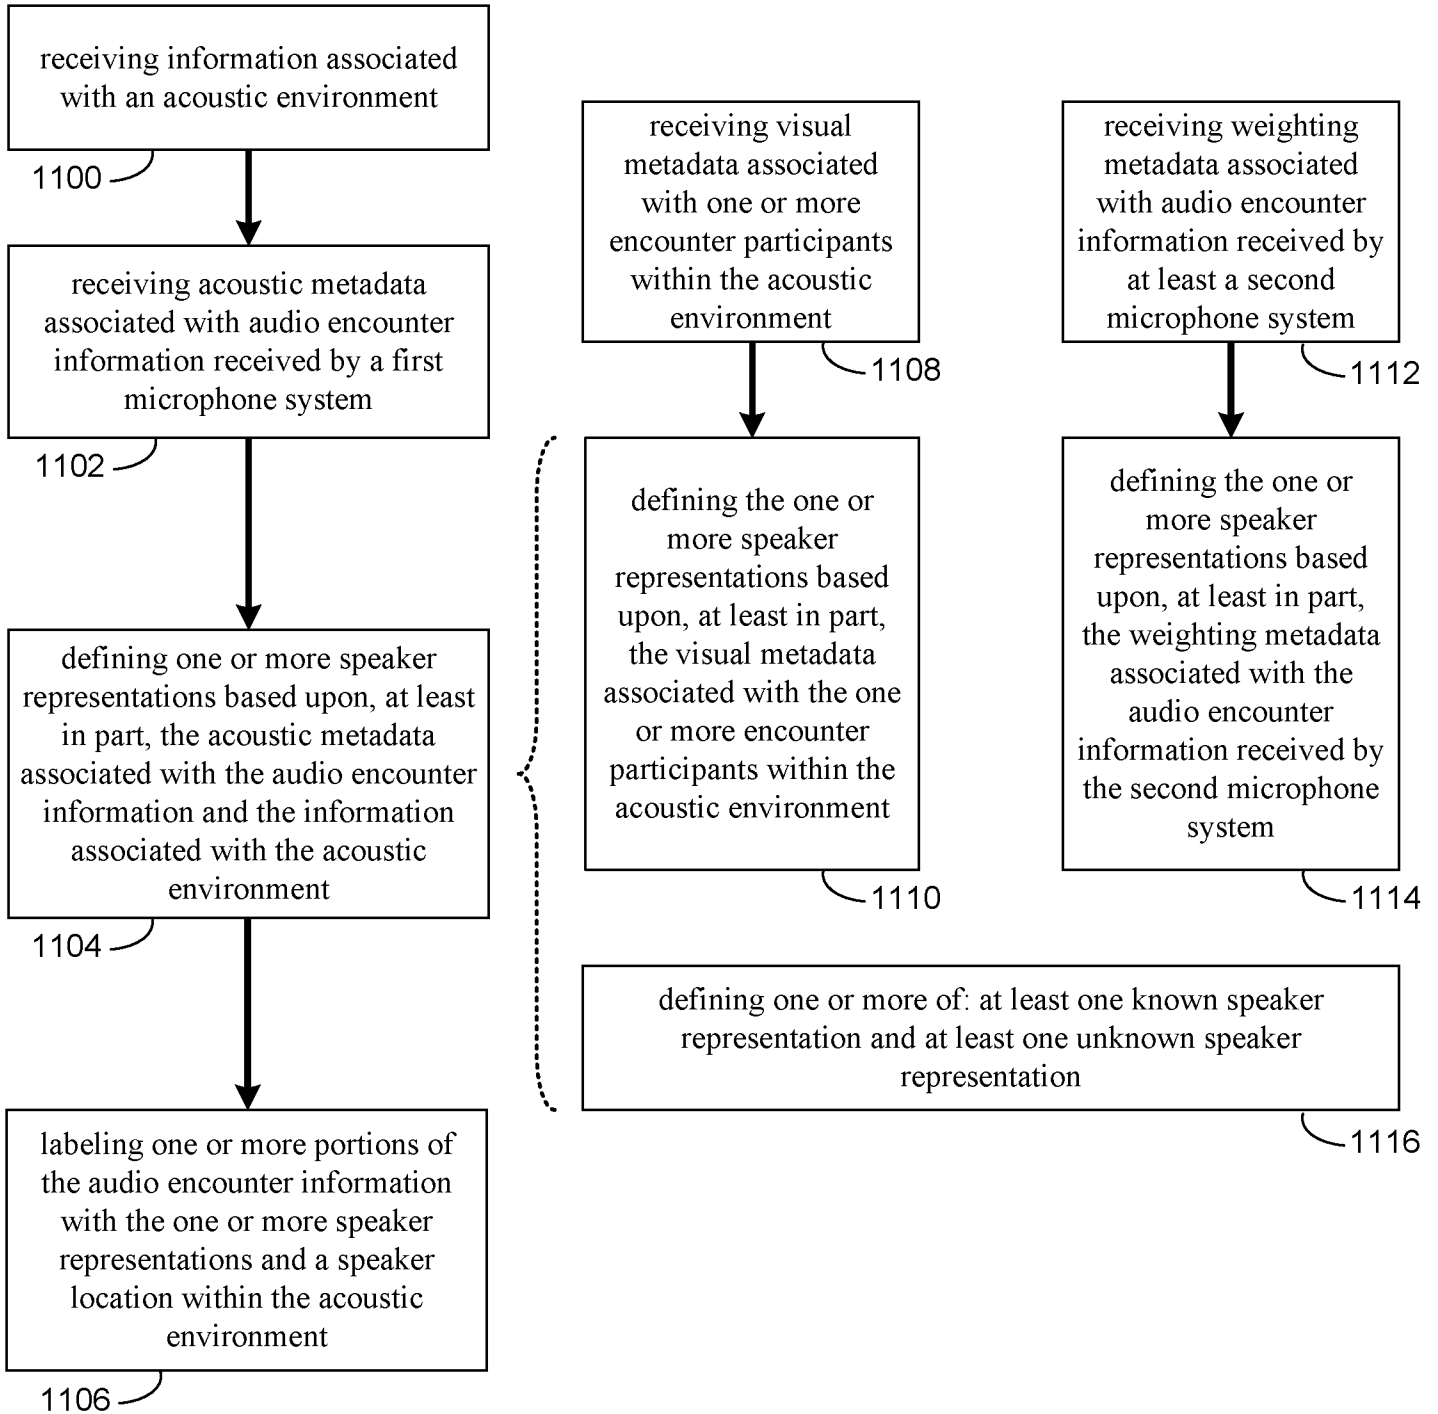


FIG. 11




FIG. 12

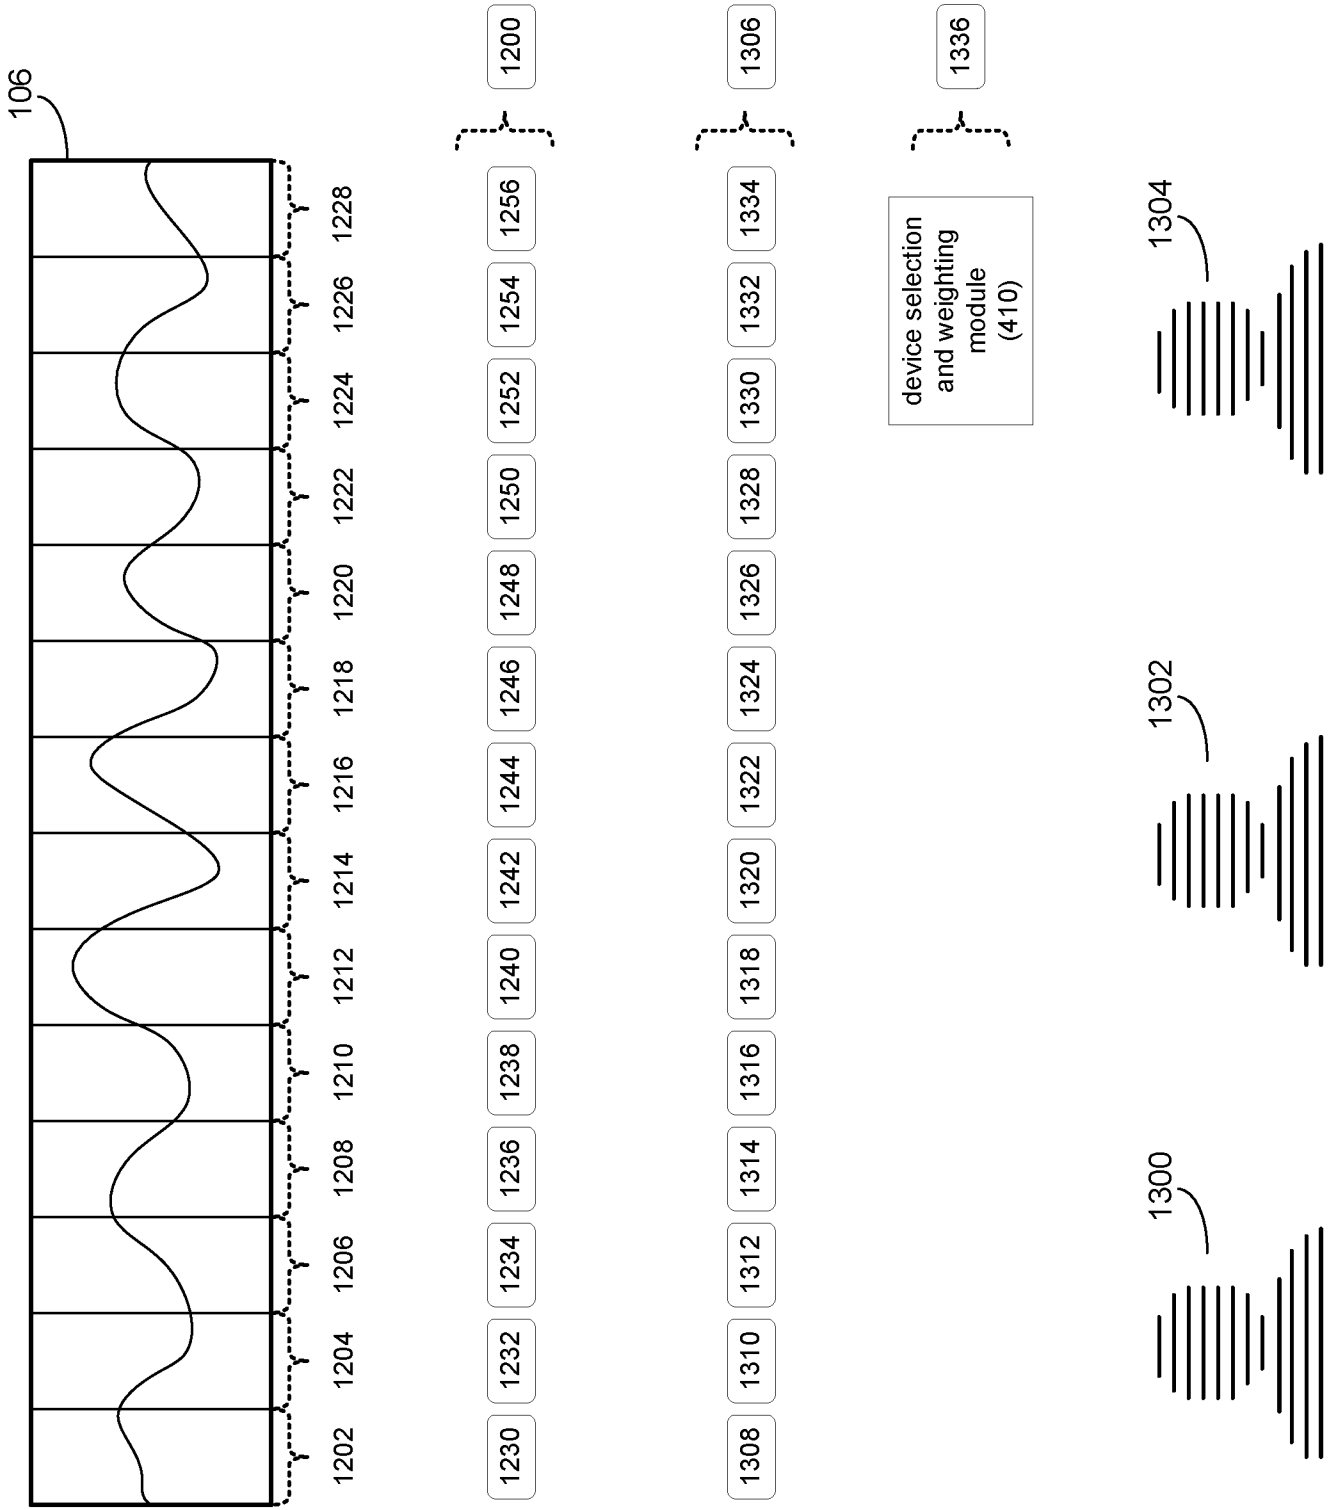


FIG. 13




FIG. 14

10

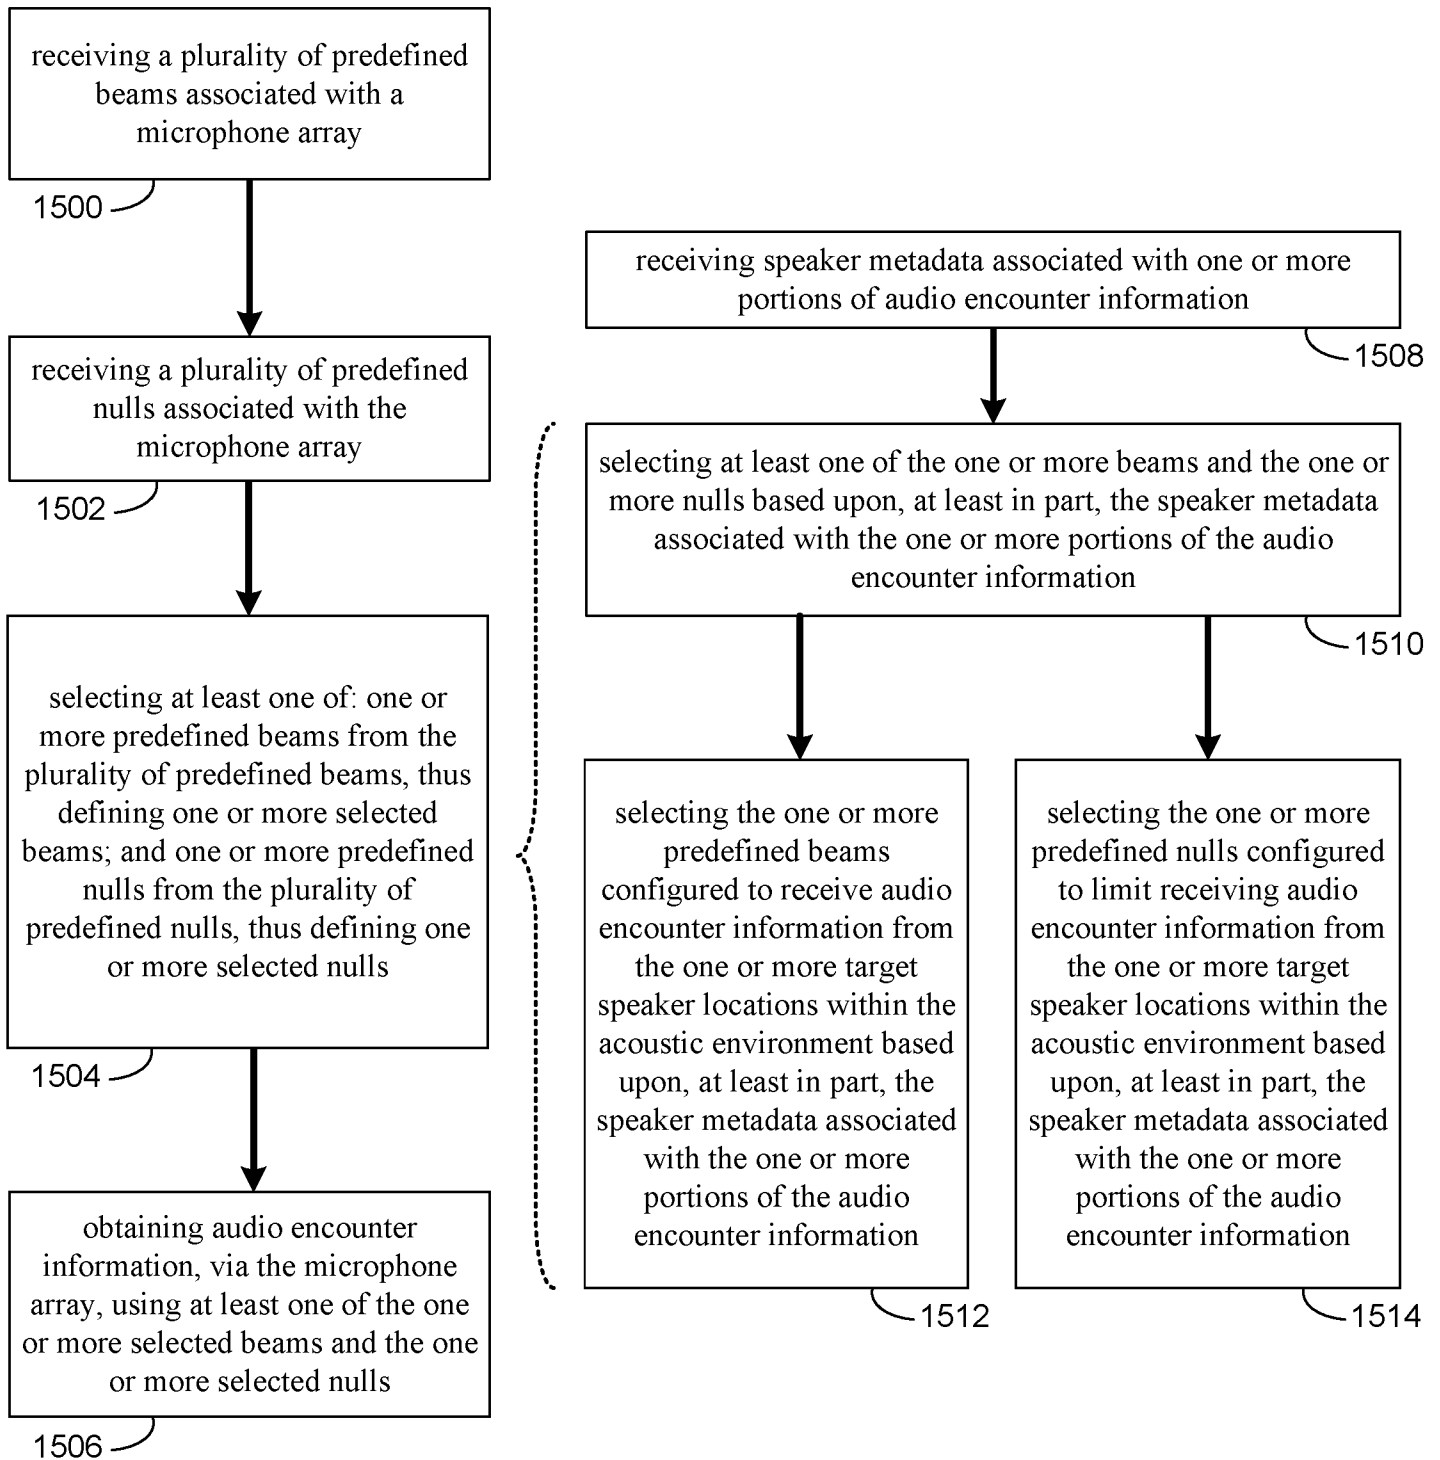


FIG. 15

10

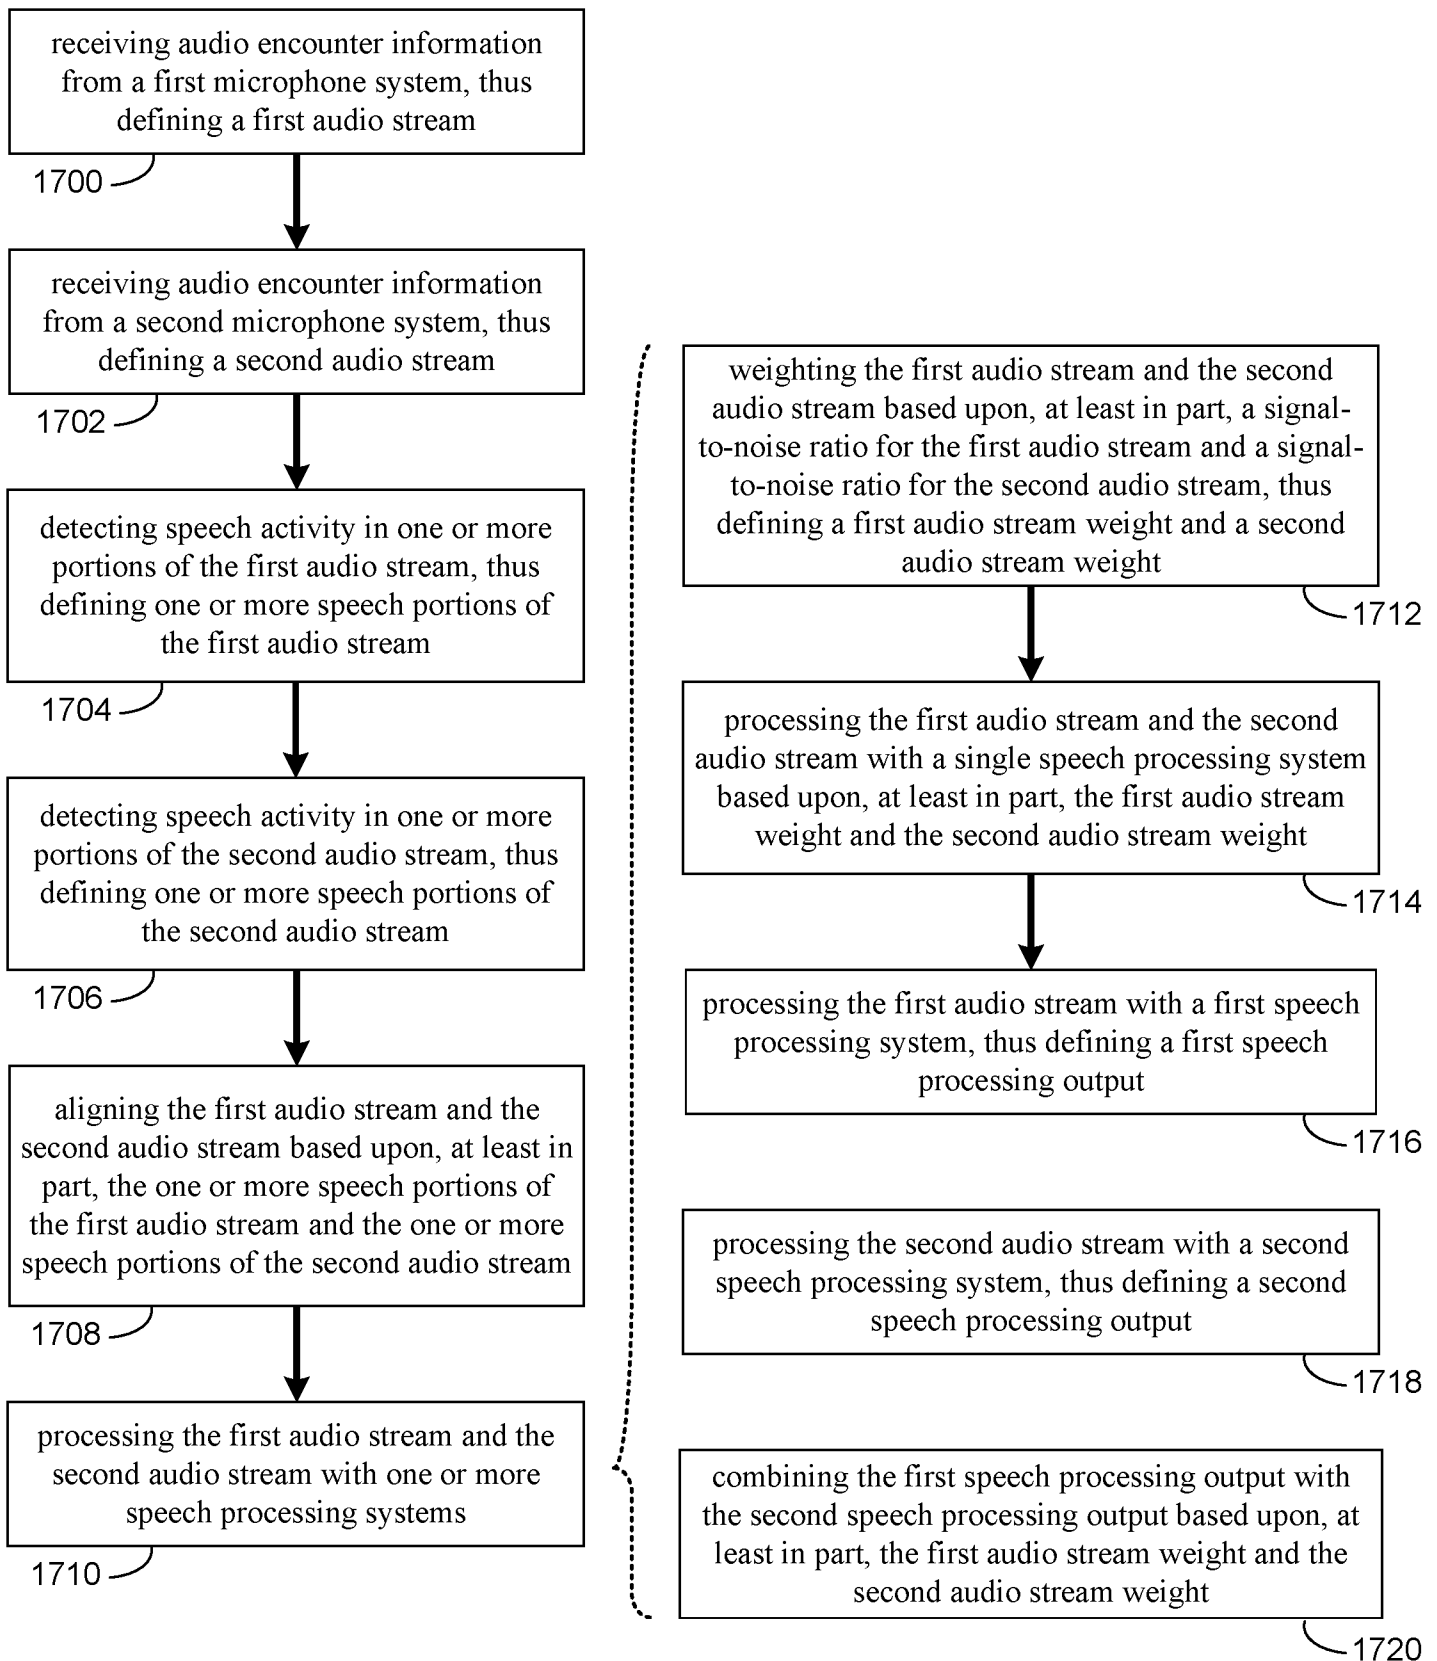


FIG. 17






FIG. 18




FIG. 19




FIG. 20A

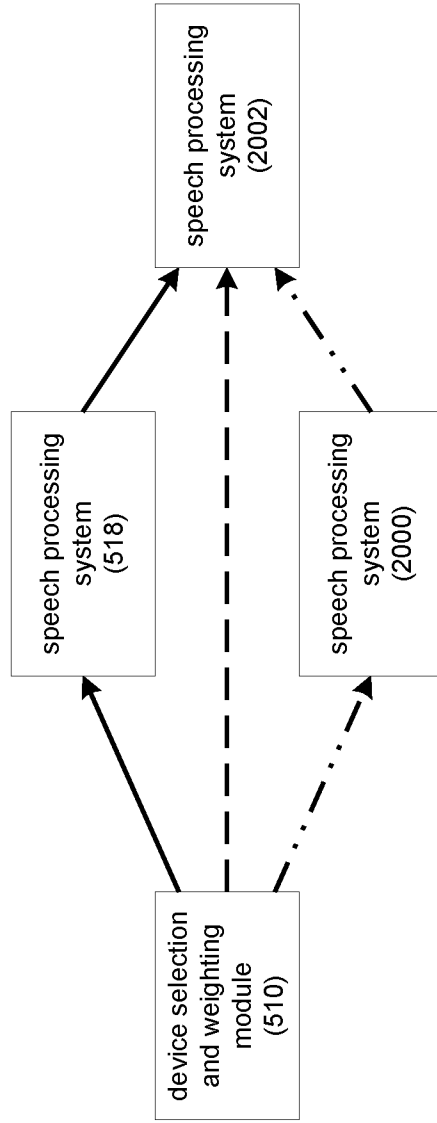


FIG. 20B

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US 21/31512

A. CLASSIFICATION OF SUBJECT MATTER
 IPC - G10L 15/20, 15/26, 21/02 (2021.01)
 CPC - G10L 15/20, 15/26, 21/02, 2021/02166

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
 See Search History document

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

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

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2019/0051375 A1 (NUANCE COMMUNICATIONS, INC.) 14 February 2019 (14.02.2019) entire document, especially para [0042], [0058]-[0060], [0066]	1-20
A	US 2017/0047079 A1 (SONY CORPORATION) 16 February 2017 (16.02.2017) entire document, especially para [0155], [0341], [0729], [0731]	1-20
A	US 2018/0176679 A1 (VERIZON PATENT AND LICENSING INC.) 21 June 2018 (21.06.2018) entire document, especially para [0014], [0017], [0018], [0021], [0023], [0045], [0046]	1-20
X,P	US 2020/0335088 A1 (TENCENT TECHNOLOGY (SHENZHEN) COMPANY LIMITED) 22 October 2020 (22.10.2020) entire document	1-20

Further documents are listed in the continuation of Box C.

See patent family annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"D" document cited by the applicant in the international application

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

21 July 2021

Date of mailing of the international search report

AUG 16 2021

Name and mailing address of the ISA/US

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents
 P.O. Box 1450, Alexandria, Virginia 22313-1450
 Facsimile No. 571-273-8300

Authorized officer

Kari Rodriguez

Telephone No. PCT Helpdesk: 571-272-4300