



- (51) International Patent Classification:  
G10L 15/00 (2006.01) G10L 15/22 (2006.01)  
G10L 15/26 (2006.01)
- (21) International Application Number:  
PCT/US2012/043185
- (22) International Filing Date:  
19 June 2012 (19.06.2012)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:  
61/498,586 19 June 2011 (19.06.2011) US
- (71) Applicant (for all designated States except US): **MMOD-AL IP LLC** [US/US]; 9009 Carothers Parkway, Suite C-2, Franklin, Tennessee 37067 (US).

- (72) Inventors; and
- (75) Inventors/Applicants (for US only): **KOLL, Detlef** [DE/US]; 6954 Reynolds St., Pittsburgh, Pennsylvania 15208 (US). **FRITSCH, Juergen** [DE/US]; 509 Gettysburg Street, Pittsburgh, Pennsylvania 15206 (US). **FINKE, Michael** [DE/US]; 1172 Murray Hill Avenue, Pittsburgh, Pennsylvania 15217 (US).
- (74) Agent: **PLOTKIN, Robert**; Robert Plotkin, P.C., 15 New England Executive Office Park, Burlington, Massachusetts 01803 (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, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

[Continued on next page]

(54) Title: DOCUMENT EXTENSION IN DICTATION-BASED DOCUMENT GENERATION WORKFLOW

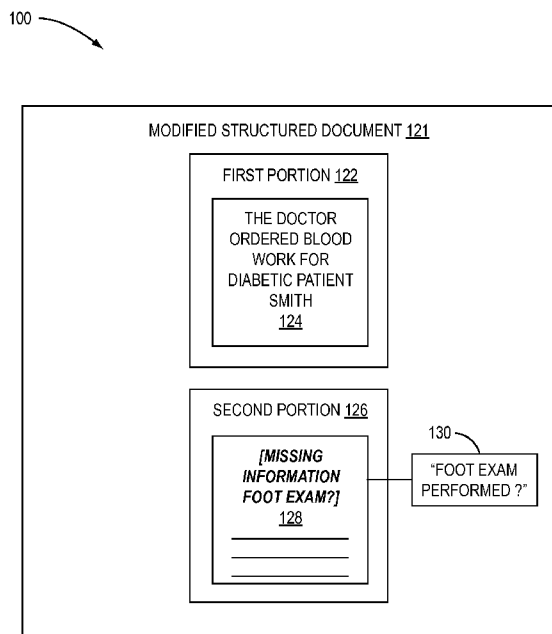


FIG. 1D

(57) Abstract: An automatic speech recognizer is used to produce a structured document representing the contents of human speech. A best practice is applied to the structured document to produce a conclusion, such as a conclusion that required information is missing from the structured document. Content is inserted into the structured document based on the conclusion, thereby producing a modified document. The inserted content may be obtained by prompting a human user for the content and receiving input representing the content from the human user.

WO 2012/177662 A1

- (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, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).
- Published:**
- *with international search report (Art. 21(3))*
  - *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))*

# Document Extension in Dictation-Based Document Generation Workflow

## BACKGROUND

[0001] A variety of automatic speech recognition (ASR) systems exist for recognizing speech and for generating text and/or commands based on such speech. ASR systems that generate text based on speech are typically referred to as “dictation” or “transcription” systems.

[0002] Some dictation systems allow users to dictate freeform text that is entered directly into documents exactly as spoken, except possibly with spelling corrections or other minor corrections applied. A common example of such systems is the dictation systems commonly used in conjunction with word processors to dictate memos, articles, and other prose documents. In contrast, some dictation systems are designed to generate and input data in a structured format based on the user’s speech. For example, such an ASR system may be used to enter data into a database form, in which data in each field is constrained to have only certain values and to be represented in certain formats. For example, a “Priority” field in such a form may be associated with a dropdown list having only three values, such as “High,” “Medium,” and “Low,” in which case the speech of a user dictating into such a field may be constrained to produce only one of these three possibilities as speech recognition results.

[0003] Whether the user provides input using speech input or other kinds of input (such as mouse or keyboard input), both freeform and structured input modalities have a variety of advantages and disadvantages. For example, one advantage of freeform input modalities is that they allow the user to provide a wide range of input and therefore to capture subtleties in the information provided. A corresponding disadvantage of structured input modalities is that by constraining the input options available to the user, they may fail to capture information that cannot accurately be represented using the available options.

[0004] One advantage of structured input modalities is that they require data to be stored in the form of discrete data elements that computers can act on easily and automatically. For example, if a patient has an allergy to allergen X and currently takes medication Y, if these two facts are input using structured input modalities (such as by using dropdown lists for selecting

allergens and medications, respectively), then the resulting data can be encoded in discrete data elements representing allergen X and medication Y, respectively, and a computer can performing processing on such information easily and automatically, such as by comparing the data to contraindications on a predetermined list to determine whether allergen X is contraindicated with medication Y. Structured input modalities, in other words, enable data to be input in a form such that the meaning of the data is available to and unambiguously processable by a computer without the need for further interpretation. In contrast, data input using freeform input modalities (such as the text “the patient is allergic to X and is taking medication Y”) must be parsed and interpreted in an attempt to discern the meaning of the text before a computer can attempt to parse the information represented by the text. Such parsing and interpretation are subject to errors which can impede the ability to perform the kind of processing that can be performed easily and automatically using data obtained using structured input modalities.

**[0005]** Another advantage of structured input modalities is that they can enable data to be input with fewer errors because the user is prevented, for example, from providing the wrong type of information (such as by inputting a name into a date field) or from providing information that is outside a permissible range of values (such as by entering -10 into a field representing the body temperature of a patient). A corresponding disadvantage of freeform input modalities is that they can allow input to contain a wide variety of errors because the user’s input is not constrained.

**[0006]** Yet another advantage of freeform input modalities is that they do not predispose the user toward providing any one particular input over another. In contrast, structured input modalities can bias the input provided by the user towards the set or range of permitted inputs. This bias can be undesirable when the goal of the system is to faithfully capture and represent information. For example, a documentation system that only offers a checkbox as a means for inputting information about the presence or absence of a particular fact—such as whether a patient has hypertension—forces the user to pigeonhole the user’s knowledge about the patient’s condition into a binary value of “yes” or “no.” If the user’s knowledge of the patient indicates, for example, that the patient currently has mild hypertension, that the patient may possibly have hypertension, or that the patient is in the process of developing hypertension, then requiring the user to provide an answer of “yes” or “no” to the question, “Does the patient have

hypertension?,” will result in a misrepresentation of the true state of the patient in relation to hypertension.

[0007] Some input systems (such as EMR systems) attempt to address this problem by adding additional input choices, such as by enabling the user to provide not only binary answers to questions about facts, but also to provide information about additional qualities related to those facts, such as degree, likelihood, conditionality, and interdependency of such facts. For example, such a system might require the user to provide a “yes” or “no” answer to the question, “Does the patient have hypertension?,” but also ask the user, in connection with that question, to provide a degree to which the user’s “yes” or “no” answer is correct, a likelihood that the user’s “yes” or “no” answer is correct, and so on.

[0008] Although this kind of solution can address some of the problems with structured input modalities, adding such additional input choices can quickly make using the system unwieldy due to the large amount and variety of inputs that the user must provide. Furthermore, the meanings of the additional input choices may not be clear to the users responsible for selecting such choices. For example, if the user is provided with choices of “Low,” “Medium,” and “High” for rating the likelihood that the patient has hypertension, it is not clear whether these choices represent equally-distributed probabilities such as 0-33.3%, 33.3-66.7%, and 66.7-100%, or some other ranges, such as 0-10%, 10-90%, and 90-100%. As a result, the user must make some decision about how to interpret the available input choices, and the result of that decision may differ both from that intended by the designer of the system and from the decisions made by other users when faced with the same set of input choices. As a result, the inputs provided by users may result in inaccurate information being entered into the system.

[0009] In light of these various advantages and disadvantages of freeform and structured input modalities, what is needed are improved techniques for capturing data to maximize accuracy and minimize errors.

### **SUMMARY**

[0010] In one aspect, a method includes applying automatic speech recognition to an audio signal to produce a structured document representing contents of the audio signal. The

method includes determining whether the structured document includes an indication of compliance with a best practice to produce a conclusion. The method includes inserting content into the structured document, based on the conclusion, to produce a modified structured document.

[0011] In another aspect, a system includes an automatic speech recognition system and a structured document analyzer. The automatic speech recognition system applies automatic speech recognition to an audio signal to produce a structured document representing contents of the audio signal. The automatic speech recognition system may include a structured document generator producing the structured document. The structured document analyzer determines whether the structured document includes an indication of compliance with a best practice to produce a conclusion and inserting content into the structured document, based on the conclusion, to produce a modified structured document

[0012] Other features and advantages of various aspects and embodiments of the present invention will become apparent from the following description and from the claims.

### **BRIEF DESCRIPTION OF THE DRAWINGS**

[0013] The foregoing and other objects, aspects, features, and advantages of the disclosure will become more apparent and better understood by referring to the following description taken in conjunction with the accompanying drawings, in which:

[0014] FIG. 1A is a block diagram depicting one embodiment of a system for document extension in dictation-based document generation workflow;

[0015] FIG. 1B is a block diagram depicting one embodiment of a structured document representing contents of an audio signal; and

[0016] FIG. 1C is a block diagram depicting one embodiment of a system for document extension in dictation-based document generation workflow;

[0017] FIG. 1D is a block diagram depicting one embodiment of a structured document modified to request additional content; and

[0018] FIG. 2 is a flow diagram depicting one embodiment of a method for document extension in dictation-based document generation workflow.

### **DETAILED DESCRIPTION**

[0019] As described above, freeform and structured input modalities have a variety of advantages and disadvantages. In particular, it was mentioned that a structured input modality can bias the user toward providing the particular inputs that the structured input modality permits, thereby potentially causing inaccurate data to be generated. The bias imposed by structured input modalities, however, is not always harmful. In some circumstances it can have beneficial effects. For example, by forcing users to commit to certain discrete input choices, such as a particular diagnosis in the case of an electronic medical record, the resulting data may be less accurate but more unambiguous than it would have been if a freeform input modality had been used. Despite the potential decrease in accuracy that results from requiring input to be selected from among a discrete set of permissible inputs, forcing the user to make such a selection may cause the physician or other user to make the crucial decision about which discrete input to select early in the process (namely, at the time of providing the input), rather than leaving such a certain to the same user, or to a different and less qualified user, at a later date. As a result, the resulting data may be more likely to be useful for making accurate decisions without the need for further review or revision in certain circumstances, such as when the input data is to be used to submit a medical report to a payer for a treatment plan, if such payer requires specific diagnoses to be provided before payment can be made.

[0020] As another example, structured input forms can be useful for providing guidance to the user about the set of information that must be provided and about the possible contents of such information. For example, the presence of a form field containing the question, "Does the patient have hypertension?," can serve as a reminder to the user that the answer to such a question, and possibly other information related to the patient in connection with hypertension, is required. As a result, using forms to create documentation can result in fewer omissions of required or desirable information than using freeform text input. Because different forms can be used in different contexts and to create different kinds of documentation, forms can be tailored to prompt the user for the information that is relevant and/or required in each context. Therefore, one significant benefit of structured input modalities is that they can reduce the frequency of errors in the form of omissions of relevant and/or required data.

[0021] Embodiments of the present invention may be used to obtain the benefit of the advantages of both freeform and structured input modalities, while reducing or eliminating the disadvantages of both such input modalities.

[0022] Referring now to FIG. 1A and 1B, in conjunction with FIG. 2, a method 200 includes applying automatic speech recognition to an audio signal to produce a structured document representing contents of the audio signal (210). In particular, embodiments of the present invention may enable users to dictate information in freeform, as if they were dictating a document to be transcribed into a freeform transcript. An automatic speech recognition system 102 may then be used to generate a structured document 120 based on the user's speech, where the structured document 120 encodes concepts from the user's speech into discrete data elements. For example, the automatic speech recognition system 102 may include a structured document generator 104 receiving an audio signal via an I/O device 103 (such as, for example, a microphone) and generating the structured document 120. As another example, the automatic speech recognition system 102 may generate a transcription 108 of the audio signal and provide the transcription 108 to the structured document generator 104 for generation of the structured document 120. As shown in FIG. 1B, the structured document 120 may include a plurality of portions (which may also be referred to as sub-templates, sub-portions or document components). A portion in the plurality of portions may or may not contain information – as shown in FIG. 1B, a first portion 122 contains information 124 (e.g., “The doctor ordered blood work for diabetic Patient Smith”) but the second portion 126 does not contain any information. Examples of techniques that may be used to create such structured documents may be found, for example, in commonly-owned U.S. Pat. No. 7,584,103, entitled, “Automated Extraction of Semantic Content and Generation of a Structured Document from Speech”; commonly-owned U.S. Pat. No. 7,716,040, entitled, “Verification of Extracted Data”; and the patents, patent applications, and other documents cited therein. These techniques, however, are merely examples of techniques that may be used to generate a structured document 120 based on speech, and do not constitute limitations of the present invention.

[0023] The method 200 includes determining whether the structured document includes an indication of compliance with a best practice to produce a conclusion (220). For example, the automatic speech recognition system 102 may include, or be in communication with, a structured

document analyzer 106 that analyzes the structured document 120 to determine whether the structured document 120 includes an indication of compliance with a best practice and produces a conclusion that represents the outcome of the determination. The structured document analyzer 106 may determine that the structured document 120 includes an indication of compliancy due to the presence of information within the structured document. Alternatively, for example, the structured document analyzer 106 may determine that the structured document 120 includes an indication of compliancy due to the absence of information within the structured document.

**[0024]** The structured document analyzer 106 may determine whether the form and/or content of the structured document 120 include the indication of compliance with the best practice (e.g., by determining whether the structured document 120 contains a particular section or piece of information); in such an embodiment, the structured document analyzer 106 may determine whether the contents of the structured document 120 comply with the best practice.

**[0025]** Additionally or alternatively, the structured document analyzer 106 may determine whether an action or other fact described by the structured document 120 complies with the best practice; in such an embodiment, the structured document analyzer 106 may determine whether a fact represented by the contents of the structured document 120 complies with the best practice. As an example, the structured document analyzer 106 may determine that to comply with a best practice, a particular structured document 120 should include an indication as to whether a foot examination was performed during a patient-physician encounter (e.g., the structured document 120 should include a statement such as “During the visit, Dr. Jones performed a foot examination for Patient Smith”); in determining that the structured document 120 includes no such indication, or that the structured document 120 includes an indication that no such examination occurred, the structured document analyzer 106 may produce a conclusion that the patient-physician encounter itself did not comply with the best practice since no foot examination was performed.

**[0026]** Once such a structured document 120 has been created, embodiments of the present invention may be used to automatically provide context-dependent guidance to encourage the use of best documentation practices by users of the system, such as the speaker whose speech was recognized to create the structured document 120 and/or other users who are responsible for editing, reformatting, and proofreading the document. For example, consider a

case in which a user speaks into an audio capture device (such as a microphone connected to a desktop computer, or a handheld audio device containing an embedded microphone, or other input/output device 103) and in which the user has access to a computer monitor (or other display device 105) for reviewing and signing off on the resulting transcript. While (or after) the user dictates a freeform report into the audio capture device, any of the techniques referenced above may be used to generate a structured document 120 based on the user's speech. As another example, a user of a computing device 101b (not shown) may use an input/output device 103 to provide audio signals to the computing device 101b; the computing device 101b may transmit the audio signals to the computing device 101 shown in FIG. 1A.

[0027] An embodiment of the present invention may include or have access to information about best practices for use in connection with the report dictated by the user. Such best practices may, for example, include rules, guidelines, or preferences related to any one or more of the following:

- information that must be included in the report (such as the patient's name and date of the patient's visit); and
- actions that are required to be taken before the report can be completed (such as performing a particular procedure, such as a foot exam, on the patient during the visit described in the report).

[0028] Embodiments of the present invention may use a single set of best practices for all reports, or may use different sets of best practices to different reports. For example, a distinct set of best practices may apply to each of a plurality of report types and/or content. For any particular report, the set of best practices corresponding to the report's type (or corresponding to content in the report) may be applied to the report. As another example, a generic (e.g., universal) set of best practices may be applied to all reports in addition to specific sets of best practices based on report types and/or content. These and other examples may be combined with each other in any of a variety of ways.

[0029] The best practices may be represented and stored in any of a variety of ways, as will be understood by those having ordinary skill in the art. For example, best practices may take the form of rules that operate on discrete data elements in structured reports 120. For example, a best practice rule may dictate that if the discrete data elements in a structured report

120 indicate that the patient described by the report has diabetes, then the structured report 120 must indicate whether a foot exam was performed on the patient because it is a best practice to perform regular foot exams on diabetic patients. Such rules may operate on freeform text in addition to or instead of discrete data elements.

**[0030]** Best practices need not take the form of rules. Alternatively or additional for example, they may take the form of heuristics or procedures for drawing conclusions about whether a particular report is missing required (or recommended) data, includes data that it is prohibited from including (or that is advisable not to include), or otherwise fails to comply with or falls short of any required or desired standard applicable to the report.

**[0031]** The structured document analyzer 106 may identify a best practice based on data from a variety of sources, such as any one or more of the report itself, other reports (such as other reports related to the same patient or subject), database records (such as electronic medical records (EMRs)), or any other source of data. For example, the structured document analyzer 106 may analyze a data source relating to an event to which the structured document 120 relates to identify the best practice.

**[0032]** A structured report 120, once created, may be displayed to the user who dictated the report and/or to other users in any of a variety of ways. For example, the structured report 120 may be displayed as a document containing the text of the report, visually structured into sections (such as by adding section headings before each section or by emphasizing dictated section headings), and with discrete data elements (such as dates, medications, allergies, and diagnoses) emphasized or otherwise formatted to indicate that they are discrete data elements and/or to indicate the types of such discrete data elements. The structured report 120, in other words, may be rendered in a format that generally resembles that of a plain text document rendered by a conventional word processor, but with the structured elements of the document (e.g., sections and lower-level concepts) visually emphasized or otherwise indicated.

**[0033]** As mentioned above, the structured document analyzer 106 may produce a conclusion 114, such as a conclusion 114 that certain required information is missing from the report. Embodiments of the present invention may provide guidance to the user to provide input, such as additions to, deletions from, or modifications to the structured report 120 to cause the

structured report 120 to conform to the best practices or to reduce the extent to which the structured report 120 fails to conform to the best practices.

**[0034]** Such guidance may be repeated for each of a plurality of best practices, and the kind and content of guidance may be the same or vary for each such best practice. Similarly, the same best practice may be applied multiple times to the same structured report 120, such as by applying the same best practice to multiple portions of the same structured report 120, thereby resulting in multiple instances of guidance within the same structured report 120 based on the same best practice.

**[0035]** The method 200 includes inserting content into the structured document, based on the conclusion, to produce a modified structured document (230). The automatic speech recognition system 102 may update the structured document 120 to include the user input; for example, the structured document generator 104 may include functionality for updating generated structured documents. In such an example, the structured document generator 104 may receive user input from the structured document analyzer 106 and update the structured document 120 with the user input. As another example, the structured document analyzer 106 may transmit an instruction to the structured document generator 104 to generate a prompt within the structured document 120. The structured document generator 104 may render the modified structured document for review by the user. FIG. 1D, described in greater detail below, provides an example of a modified structured document 121.

**[0036]** Although depicted in FIG. 1A as executing within the automatic speech recognition system 102, the structured document generator 104 and the structured document analyzer 106 may be provided as separate components in communication with the automatic speech recognition system 102. Referring now to FIG. 1C, and by way of example, the automatic speech recognition system 102 may transmit a transcription 108 of an audio signal to a stand-alone separate structured document generator 104, which in turn may transmit a structured document 120 to the structured document analyzer 106. In some embodiments, the structured document generator 104 and the structured document analyzer 106 may be provided as part of a report generation component (not shown); for example, the structured document generator 104 and the structured document analyzer 106 may be provided as plug-ins, add-ons or other extensions to software for generating components.

[0037] Embodiments of the present invention may provide output to the user to prompt the user to provide the input mentioned above for revising the structured report 120. Such prompting output may take any of a variety of forms. For example, a prompt (such as a window or a dialog box) that is distinct from the rendering of the report itself may be displayed to the user to prompt the user to provide input (e.g., by displaying the question “Was a foot exam performed on the patient?”). The user may provide such input in response to the distinct prompt, in response to which the system may update the structured report 120 accordingly (such as by adding a word, phrase, or sentence to the report representing the user’s input, or by checking off an appropriate box in the report). The user may provide the input in freeform or in a structured form; for example, where the prompt includes an open-ended question, the user may provide the input in freeform. As an alternative example, when the prompt includes a closed question, the user may provide the input in a structured form. The form of the input provided in response to the prompt may differ from the form originally used to generate the structured document 120; for example, the structured document 120 may have been generated based on freeform speech, while the prompt may prompt for and receive input in the form of structured data, or vice versa.

[0038] As another example, a prompt may be displayed within the rendering of the report itself. Such a prompt may, for example, take the form of text appended to or inserted into the rendering of the report at an appropriate location. For example, if application of a best practice rule leads to the conclusion that the report is missing information about whether a foot exam was performed on the patient, and if the report contains a section on procedures performed on the patient during the current visit, then a prompt such as “Foot exam performed?” may be inserted into that section. The text input cursor may then be positioned immediately after this prompt, or the user may otherwise be notified of the prompt and given an opportunity to provide input directly into the document (such as by typing text) after the prompt.

[0039] Upon receiving input from the user, the system may update the report to include the input. For example, the system may insert the user input into the report substantially at or near the location of the prompt within the structured document 120; in such an example, the system may either maintain the prompt or remove the prompt from the location.

[0040] As another example, the prompt may take the form of a form field, such as a text field, check box, or dropdown list, into which the user may directly provide input in response to

the prompt. An appropriate label (such as “Foot exam performed?”) may be displayed near the form field to guide the user in providing input into the form field. The user may then provide input directly into the form field to cause the structured document 120 to be updated in accordance with the corresponding best practice. Referring to FIG. 1D, a block diagram depicts one example of a modified structured document 121. As shown in the example depicted in FIG. 1D, a modified structured document 121 may include a plurality of portions (which may also be referred to as sub-templates, sub-portions or document components). In the example shown in FIG. 1D, the modified structured document 121 is a version of the structured document 120 and includes the first portion 122, including information 124, and the second portion 126, which does not include any information. However, in the modified structured document 121, unlike in the structured document 120, the second portion 126 includes a prompt 128 and a label 130. The prompt 128 in this example is a text field into which a user may provide missing information. The label 130, displayed near the prompt 128, draws attention to the prompt 128 and provides guidance as to a type of information requested.

[0041] More generally, embodiments of the present invention may be contained within and/or operate in conjunction with a report generation procedure. Prompts provided by embodiments of the present invention may be part of and integrated into the report generation procedure. As a particular example, prompts may be integrated into the report itself. For example, as described elsewhere herein, as a report is being generated (and therefore after part of the report has been generated but before the entire report has been generated), any of the prompts disclosed herein may be provided to the user. Such prompts may be added or otherwise integrated into the report as the report is being generated. Similarly, input provided in response to such prompts may be added to or otherwise used to generate data that is added to the report as the report is being generated.

[0042] As one example of the above, a prompt may be a report sub-template, and such a sub-template may be inserted into the report while the report is being generated. A report sub-template may, for example, include one or more of a plurality of any kind(s) of form fields. For example, if an embodiment of the present invention determines that a particular portion of the report describes or otherwise relates to a particular type of medical procedure, in response to such a determination a report sub-template corresponding to the particular type of medical

procedure may be identified and inserted into the report at the current location within the report. Such an inserted report sub-template may include content (e.g., text) that is added to the report and/or form fields (or other kinds of prompts) to obtain input from the user to insert into the report within the inserted report sub-template. The report sub-template and any data entered by users into the report sub-template may become part of the final report. As in the examples provided above regarding inserting, into the structured document 120, input provided in response to a prompt, the system may receive input from a user in response to the report sub-template and incorporate the user input into the report sub-template; and, as above, the system may determine to maintain or remove the prompt after the insertion of the user input. In examples in which the prompt is provided as a user interface element (e.g., a window displaying text, the window rendered as part of the report generation system), the system may determine whether to continue rendering the user interface element or to remove the user interface element from a display after the user input has been received.

**[0043]** As described above, the system may prompt the user to provide input representing missing information. For example, when a best practice indicates that a particular type of content (such as a particular concept, e.g., report section or fact) should be included in a particular type of report, but the structured document analyzer 106 determines that the structured document 120 does not contain the particular type of content, the content is said to be missing from the report. By way of example, if a best practice requires that a type of report include information indicating that a foot examination was performed on a patient but the structured document 120 does not include information indicating that a foot exam was performed, then information relating to a foot exam is said to be “missing” from the report.

**[0044]** However, embodiments of the present invention are not limited to prompting the user to provide input representing missing information. As another example, embodiments of the present invention may automatically update the structured report 120 with the missing information, an estimate of the missing information, or an indication that such information is missing, without receiving input from the user. For example, if the structured document analyzer 106 determines that to comply with a particular best practice, the report is required to indicate whether a foot exam was performed and the report does not indicate whether a foot exam was performed, embodiments of the present invention may take any of a variety of actions, such as

updating the report to indicate that no information is available about whether a foot exam was performed (e.g., by inserting the text “Foot exam performed? No information available”), or by retrieving information about whether a foot exam was performed from an electronic medical record associated with the same patient and inserting that information automatically into the report (e.g., by inserting the text “Foot exam performed? Yes”) into the report. As another example, the system may prompt a user to take an action. For instance; the structured document analyzer 106 may determine whether a fact represented by the structured document 120 complies with the best practice; if the described facts are determined not to comply with the best practice, then the system may prompt the user to take an action that complies with the best practice. As an example, the system may prompt a reviewer of the structured document 120 to perform a foot exam in an embodiment in which the structured document 120 is being generated during a patient visit; the system may also prompt a reviewer of the structured document 120 to schedule, or arrange to have scheduled, a follow-up appointment with a patient to take an additional action (such as to perform a foot exam). The system may also prompt a reviewer of the structured document 120 to follow up with a practitioner to determine how to address the missing information (for example, a user reviewing the structured document 120 may be prompted to contact the physician who originally generated the audio signals and determine how the physician prefers to obtain the missing information or complete an incomplete task). In some embodiments, therefore, the system prompts a human user to take an action, receives input from the human user, the input indicating a status of the action, and inserts the input as content into the structured document 120. As with the examples above, such automatically generated information could be inserted into the structured report 120 in other ways, such as by inserting it into form fields or using it to provide input to user interface controls such as checkboxes or dropdown lists.

**[0045]** Any information inserted automatically into a report may be marked as “automatically generated.” When the report is rendered, such automatically generated information may be flagged as such, such as by highlighting such information in a particular color, to call attention to the fact that such information was generated automatically and not input manually by a human. The user may even be specifically prompted to review and, if necessary, edit such information.

[0046] Any of the techniques described above may be performed while a particular report is being transcribed or after the report has been transcribed. For example, after a particular portion of a report has been transcribed, best practices may be applied to that portion of the report in any of the ways disclosed above, even as the user continues to dictate additional portions of the report and even as such portions are transcribed to generated additional text.

[0047] Furthermore, any of the techniques described above may be used in connection with a single user or multiple users across any number of computing devices. For example, a single user may dictate a report and also provide additional input in response to prompts generated based on application of best practices to the report. As another example, a first user may dictate the report and a second user may provide additional input in response to prompts generated based on application of best practices to the report.

[0048] Embodiments of the present invention have a variety of advantages. For example, embodiments of the present invention combine some of the advantages of freeform input modalities with some of the advantages of structured input modalities, while reducing or eliminating some of the disadvantages of both. In particular, embodiments of the present invention combine the ability of freeform input modalities to enable users to provide unbiased and nuanced input with the ability of structured input modalities to enforce requirements that certain kinds of data be input. One way in which embodiments of the present invention achieve this result is by enabling the user to first provide input using a freeform input modality, then determining whether the input includes an indication of compliance with one or more best practices, and then extending the resulting report (either automatically or based on additional input from the user) based on conclusions drawn from the determination. The report extension may involve obtaining input from the user using a structured input modality, such as by obtaining input from the user in a form field. By requesting and obtaining such structured input from the user only if and when necessary, embodiments of the present invention avoid predisposing the user toward providing input that is biased in the way caused by conventional structured input modalities.

[0049] Another benefit of embodiments of the present invention is that they may increase the accuracy of the resulting report, either by guiding the user toward providing more complete information that is compliant with best practices, or by obtaining and inserting such information

into the report automatically. Reduced errors and higher degrees of completeness can be particularly beneficial in contexts, such as medical reporting, in which the accuracy and completeness of reports can have an impact on human health and safety, and in which accurate and complete reports are necessary to comply with regulations and to obtain reimbursement for the procedures that resulted in the reports.

**[0050]** Yet another advantage of embodiments of the present invention is that by reminding a physician and/or other user to take a recommended or required action (such as performing a procedure on a patient), the user can be reminded to take such action and thereby be more likely to take the action. Even if the user is reminded of the action to take at a time when it is no longer possible to take the action in connection with the event that led to the creation of the report (such as a patient visit), the user may be reminded of the need to take the action in similar circumstances in the future. In this way, embodiments of the present invention may serve to educate users and to increase the likelihood that they will take recommended and/or required actions in the future.

**[0051]** Another advantage of embodiments of the present invention is that they may be used to encourage or require users to provide input that is useful and/or required for subsequent processing, such as information that is required to generate a bill for reimbursement of healthcare services. For example, if a report is input solely in the form of freeform text using conventional methods, it might be difficult to extract information needed for billing from such text. For example, if a freeform text report describes a surgery performed on a patient and, in the course of describing the surgery, describes not only stents that actually were inserted into the patient but also stents that were attempted to be inserted into the patient without success, it might be difficult or impossible to automatically extract from such a report the number of stents that actually were inserted into the patient. Yet it may be necessary to provide such a number to generate a bill that may be used to obtain reimbursement for the surgery. If instead, using embodiments of the present invention, the physician or other user is required to provide the number of stents actually inserted as a discrete data element using a structured input modality, then the bill may be generated automatically because the system has enforced the requirement that information necessary to generate the bill be provided as input. The same logic applies to

any information that is contained within a document but is difficult to extract discretely and automatically.

**[0052]** Yet another benefit of embodiments of the present invention is that they may be used to reconcile existing stored data (such as discrete data, e.g., data stored in an EMR or other database record) with data input into a report. For example, if the dictating physician indicates that a patient has uncontrolled diabetes and existing data (e.g., data for the patient stored in an EMR) already indicates that the patient has diabetes but does not indicate whether the diabetes is uncontrolled, embodiments of the present invention may be used to revise the existing data (e.g., EMR) about the patient to specify that the patient's diabetes is uncontrolled.

**[0053]** It is to be understood that although the invention has been described above in terms of particular embodiments, the foregoing embodiments are provided as illustrative only, and do not limit or define the scope of the invention. Various other embodiments, including but not limited to the following, are also within the scope of the claims. For example, elements and components described herein may be further divided into additional components or joined together to form fewer components for performing the same functions.

**[0054]** The techniques described above may be implemented, for example, in hardware, software tangibly stored on a computer-readable medium, firmware, or any combination thereof. The techniques described above may be implemented in one or more computer programs executing on a programmable computer including a processor, a storage medium readable by the processor (including, for example, volatile and non-volatile memory and/or storage elements), at least one input device, and at least one output device. Program code may be applied to input entered using the input device to perform the functions described and to generate output. The output may be provided to one or more output devices.

**[0055]** Each computer program within the scope of the claims below may be implemented in any programming language, such as assembly language, machine language, a high-level procedural programming language, or an object-oriented programming language. The programming language may, for example, be a compiled or interpreted programming language.

**[0056]** Each such computer program may be implemented in a computer program product tangibly embodied in a machine-readable storage device for execution by a computer processor. Method steps of the invention may be performed by a computer processor executing a

program tangibly embodied on a computer-readable medium to perform functions of the invention by operating on input and generating output. Suitable processors include, by way of example, both general and special purpose microprocessors. Generally, the processor receives instructions and data from a read-only memory and/or a random access memory. Storage devices suitable for tangibly embodying computer program instructions include, for example, all forms of non-volatile memory, such as semiconductor memory devices, including EPROM, EEPROM, and flash memory devices; magnetic disks such as internal hard disks and removable disks; magneto-optical disks; and CD-ROMs. Any of the foregoing may be supplemented by, or incorporated in, specially-designed ASICs (application-specific integrated circuits) or FPGAs (Field-Programmable Gate Arrays). A computer can generally also receive programs and data from a storage medium such as an internal disk (not shown) or a removable disk. These elements will also be found in a conventional desktop or workstation computer as well as other computers suitable for executing computer programs implementing the methods described herein, which may be used in conjunction with any digital print engine or marking engine, display monitor, or other raster output device capable of producing color or gray scale pixels on paper, film, display screen, or other output medium.

**[0057]** What is claimed is:

## CLAIMS

1. A computer-implemented method performed by at least one computer processor, the method comprising:

- 5 (A) applying automatic speech recognition to an audio signal to produce a structured document representing contents of the audio signal;
- (B) determining whether the structured document includes an indication of compliance with a best practice to produce a conclusion;
- 10 (C) inserting content into the structured document, based on the conclusion, to produce a modified structured document.

2. The method of claim 1, wherein (C) comprises:

- (C)(1) prompting a human user to provide the content;
- (C)(2) receiving input from the human user representing the content; and
- 15 (C)(3) inserting the content into the structured document.

3. The method of claim 2:

- wherein (C)(1) comprises displaying a structured input element; and
- wherein (C)(2) comprises receiving input from the human user in the
- 20 structured input element.

4. The method of claim 3, wherein the structured input element comprises a form field.

5. The method of claim 1, wherein (C) comprises:

- (C)(1) obtaining the content from a data source without human input; and
- (C)(2) inserting the content into the structured document.

6. The method of claim 1, wherein (B) further comprises:

- 30 identifying a type of the structured document;

identifying a best practice associated with the identified type; and  
determining whether the structured document includes an indication of  
compliance with the identified best practice to produce the conclusion.

5           7. The method of claim 6, further comprising determining whether the structured  
document includes an indication of compliance with a universal best practice.

8. The method of claim 1, wherein (B) further comprises determining whether the  
structured document includes an indication of compliance for each of a plurality of best  
10 practices to produce the conclusion.

9. The method of claim 8, further comprising:

generating a first indication that a user should provide additional input  
of a first type to conform the structured document to a first best practice in the  
15 plurality of best practices; and

generating a second indication that the user should provide additional  
input of a second type to conform the structured document to a second best  
practice in the plurality of best practices.

20           10. The method of claim 8, further comprising:

generating a first indication that a user should provide additional input  
to conform a first portion of the structured document to a first best practice in  
the plurality of best practices; and

25           generating a second indication that the user should provide additional  
input to conform a second portion the structured document to the first best  
practice in the plurality of best practices.

11. The method of claim 1, further comprising producing a conclusion that  
information is missing from the structured document.

30

12. The method of claim 1, further comprising producing a conclusion indicating a user should provide input to the structured document.

5 13. The method of claim 1, further comprising generating an indication that a user should provide additional input to conform the structured document to the best practice.

14. The method of claim 1, further comprising generating a prompt for additional input, based on the conclusion, the prompt rendered separately from a rendering of the structured document.

10

15. The method of claim 14, further comprising:  
receiving input representing the content; and  
removing the rendered prompt.

15 16. The method of claim 1, further comprising generating a prompt for additional input, based on the conclusion, the prompt included in a rendering of the structured document.

17. The method of claim 1, wherein (C) comprises:

20 (C) inserting, into the structured document, an indication that the content is missing.

18. The method of claim 2:

25 wherein (C)(1) comprises displaying a structured input element;  
wherein (C)(2) comprises receiving input from the human user in the structured input element.

19. A system comprising:

30 an automatic speech recognition system applying automatic speech recognition to an audio signal to produce a structured document representing contents of the audio signal; and

a structured document analyzer determining whether the structured document includes an indication of compliance with a best practice to produce a conclusion and inserting content into the structured document, based on the conclusion, to produce a modified structured document.

5

20. A non-transitory computer readable medium having instructions stored thereon, wherein the instructions are executable by at least one computer processor to perform a method, the method comprising:

- (D) applying automatic speech recognition to an audio signal to produce a structured document representing contents of the audio signal;
- (E) determining whether the structured document includes an indication of compliance with a best practice to produce a conclusion;
- (F) inserting content into the structured document, based on the conclusion, to produce a modified structured document.

10  
15

21. The computer readable medium of claim 20, wherein (C) comprises:

- (C)(1) prompting a human user to provide the content;
- (C)(2) receiving input from the human user representing the content; and
- (C)(3) inserting the content into the structured document.

20

22. The computer readable medium of claim 21:

wherein (C)(1) comprises displaying a structured input element; and  
wherein (C)(2) comprises receiving input from the human user in the structured input element.

25

23. The computer readable medium of claim 22, wherein the structured input element comprises a form field.

24. The computer readable medium of claim 20, wherein (C) comprises:

- (C)(1) obtaining the content from a data source without human input; and
- (C)(2) inserting the content into the structured document.

30

25. The computer readable medium of claim 20, wherein (B) further comprises:  
identifying a type of the structured document;  
identifying a best practice associated with the identified type; and  
5 determining whether the structured document includes an indication of  
compliance with the identified best practice to produce the conclusion.

26. The computer readable medium of claim 25, further comprising determining  
whether the structured document includes an indication of compliance with a universal  
10 best practice.

27. The computer readable medium of claim 20, wherein (B) further comprises  
determining whether the structured document includes an indication of compliance for  
each of a plurality of best practices to produce the conclusion.

28. The computer readable medium of claim 27, wherein the method further  
comprises:

generating a first indication that a user should provide additional input  
of a first type to conform the structured document to a first best practice in the  
20 plurality of best practices; and

generating a second indication that the user should provide additional  
input of a second type to conform the structured document to a second best  
practice in the plurality of best practices.

29. The computer readable medium of claim 27, wherein the method further  
comprises:

generating a first indication that a user should provide additional input  
to conform a first portion of the structured document to a first best practice in  
the plurality of best practices; and

generating a second indication that the user should provide additional input to conform a second portion the structured document to the first best practice in the plurality of best practices.

5           30. The computer readable medium of claim 20, further comprising producing a conclusion that information is missing from the structured document.

10           31. The computer readable medium of claim 20, wherein the method further comprises producing a conclusion indicating a user should provide input to the structured document.

15           32. The computer readable medium of claim 20, wherein the method further comprises generating an indication that a user should provide additional input to conform the structured document to the best practice.

            33. The computer readable medium of claim 20, wherein the method further comprises generating a prompt for additional input, based on the conclusion, the prompt rendered separately from a rendering of the structured document.

20           34. The computer readable medium of claim 33, wherein the method further comprises:

                    receiving input representing the content; and  
                    removing the rendered prompt.

25           35. The computer readable medium of claim 20, wherein the method further comprises generating a prompt for additional input, based on the conclusion, the prompt included in a rendering of the structured document.

30           36. The computer readable medium of claim 20, wherein (C) comprises:  
(C)           inserting, into the structured document, an indication that the content is missing.

37. The computer readable medium of claim 21:

wherein (C)(1) comprises displaying a structured input element;

wherein (C)(2) comprises receiving input from the human user in the

5 structured input element.

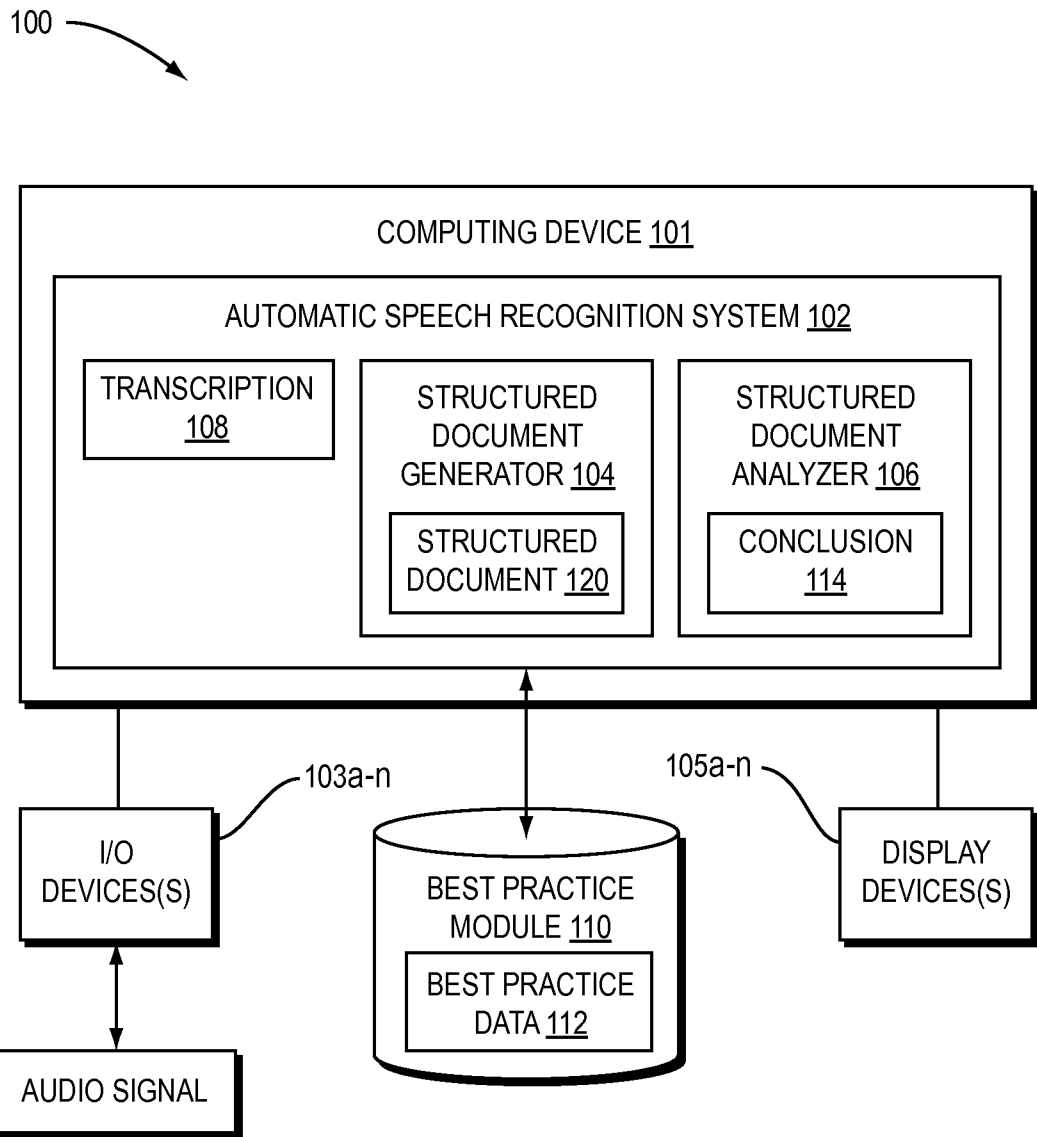


FIG. 1A

100 

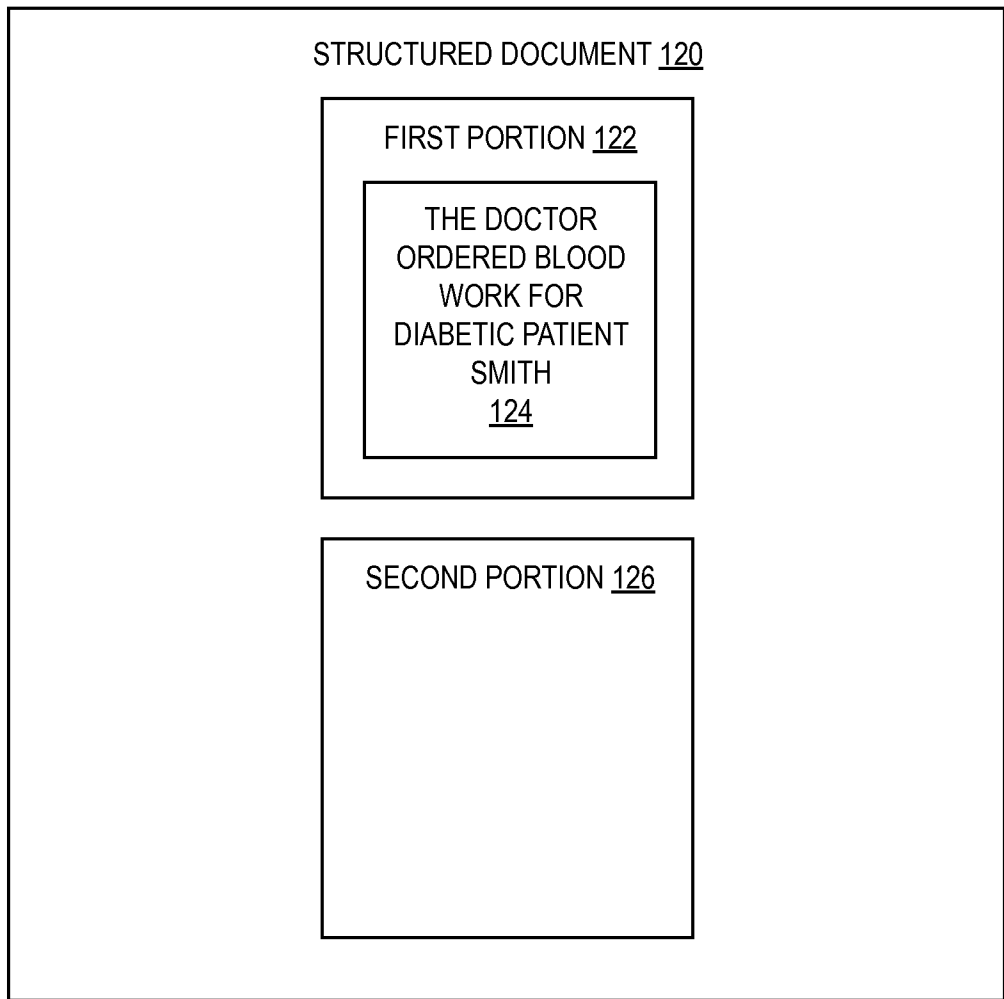


FIG. 1B

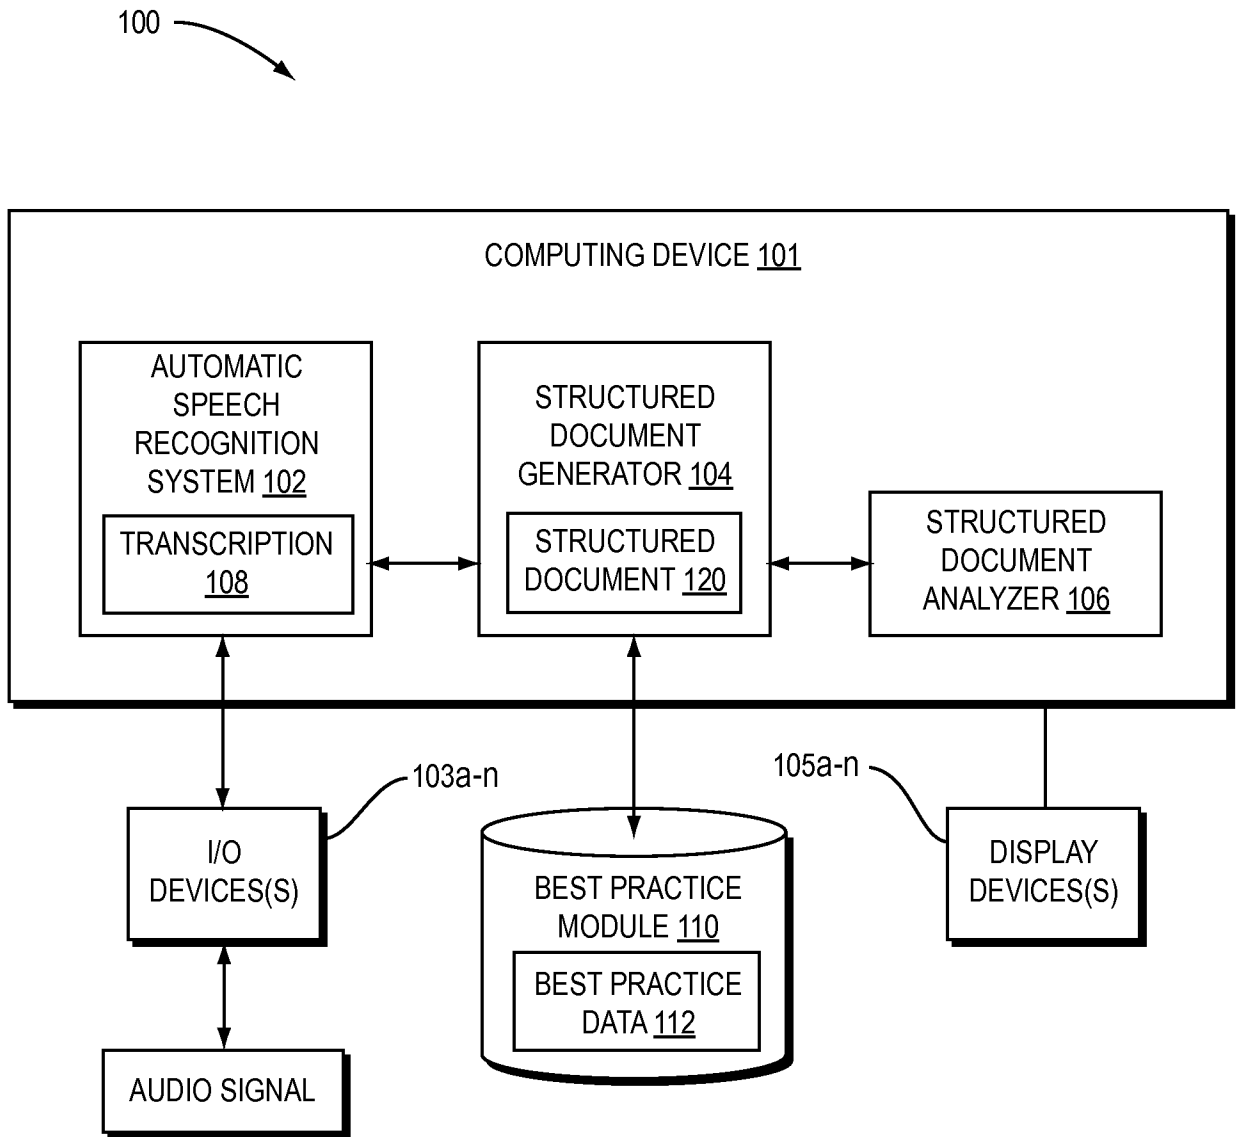


FIG. 1C

100 




FIG. 1D

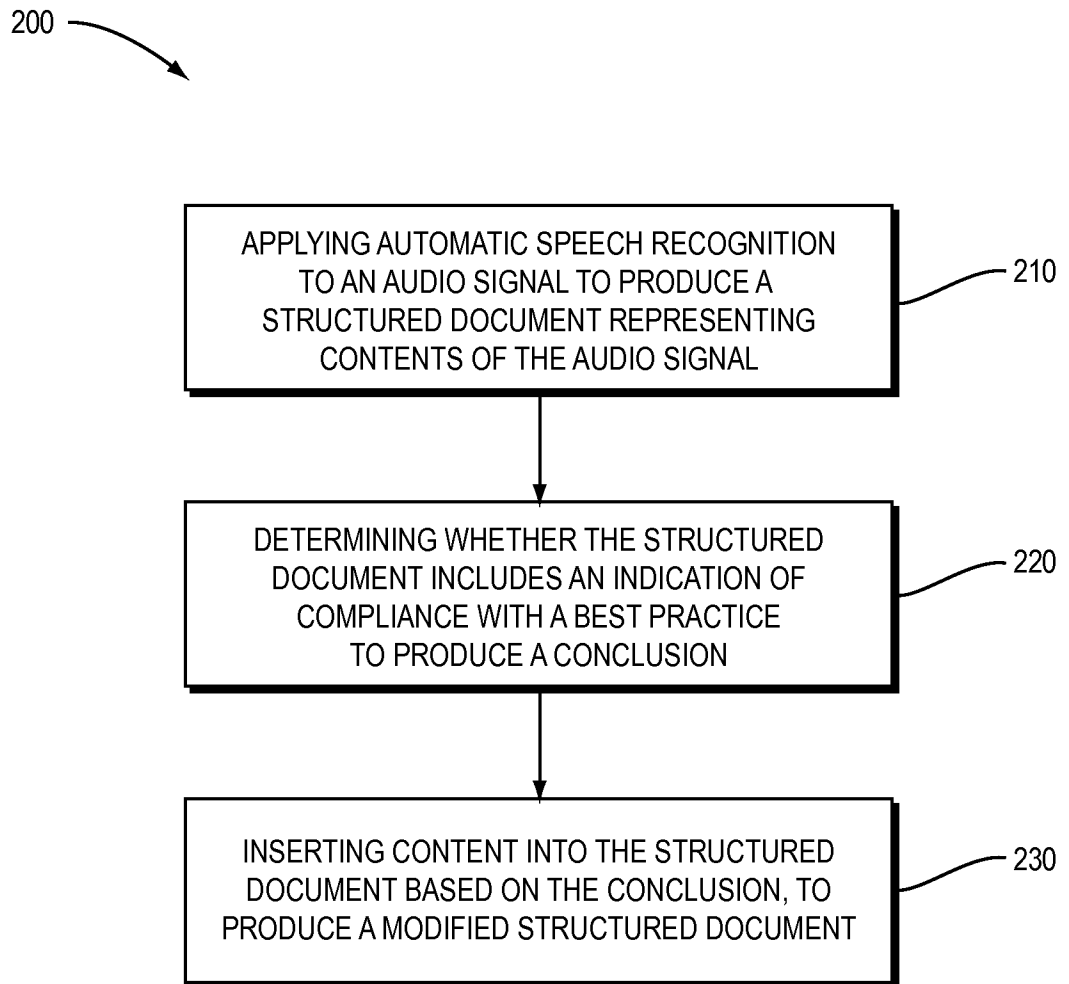


FIG. 2

## INTERNATIONAL SEARCH REPORT

International application No.  
**PCT/US2012/043185****A. CLASSIFICATION OF SUBJECT MATTER***G10L 15/00(2006.01)i, G10L 15/26(2006.01)i, G10L 15/22(2006.01)i*

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)

G10L 15/00; G09B 7/00; G06F 17/28; G10L 15/18

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

Korean utility models and applications for utility models  
Japanese utility models and applications for utility models

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

eKOMPASS(KIPO internal) &amp; Keywords: speech recognition, structured document, and best practice

**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2007-0276649 A1 (KJELL SCHUBERT) 29 November 2007 See paragraphs 7, and 67-72; claim 1; and figure 3.	1-37
A	US 2010-0291528 A1 (HUERTA JUAN MANUEL) 18 November 2010 See paragraphs 33 and 56; and the whole documents.	1-37
A	EP 0954856 B1 (MOTOROLA, INC.) 09 April 2003 See the whole documents.	1-37

 Further documents are listed in the continuation of Box C. See patent family annex.

\* Special categories of cited documents:

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

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

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of 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

"&amp;" document member of the same patent family

Date of the actual completion of the international search

28 NOVEMBER 2012 (28.11.2012)

Date of mailing of the international search report

**29 NOVEMBER 2012 (29.11.2012)**

Name and mailing address of the ISA/KR

Korean Intellectual Property Office  
189 Cheongsu-ro, Seo-gu, Daejeon Metropolitan  
City, 302-701, Republic of Korea

Facsimile No. 82-42-472-7140

Authorized officer

YANG, Seong Ji

Telephone No. 82-42-481-8117



## INTERNATIONAL SEARCH REPORT

International application No.

**PCT/US2012/043185**

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2007-0276649 A1	29. 11. 2007	CA 2653973 A1 EP 2024876 A2 JP 2009-538444 A US 2011-0131486 A1 US 7831423 B2 WO 2007-139624 A2 WO 2007-139624 A3 WO 2007-139624 A3	06. 12. 2007 18. 02. 2009 05. 11. 2009 02. 06. 2011 09. 11. 2010 06. 12. 2007 03. 07. 2008 06. 12. 2007
US 2010-0291528 A1	18. 11. 2010	WO 2010-130733 A1	18. 11. 2010
EP 0954856 B1	09. 04. 2003	AT 237176 T AU 1999-14650 A1 DE 69813180 D1 DE 69813180 T2 EP 0954856 A1 EP 0954856 A1 FR 2773413 A1 FR 2773413 B1 GB 2331826 A GB 2331826 B GB 9826231 D0 TW 462037 B US 6182038 B1 WO 99-28899 A1	15. 04. 2003 16. 06. 1999 15. 05. 2003 23. 10. 2003 10. 11. 1999 14. 08. 2002 09. 07. 1999 19. 05. 2000 02. 06. 1999 19. 12. 2001 20. 01. 1999 01. 11. 2001 30. 01. 2001 10. 06. 1999