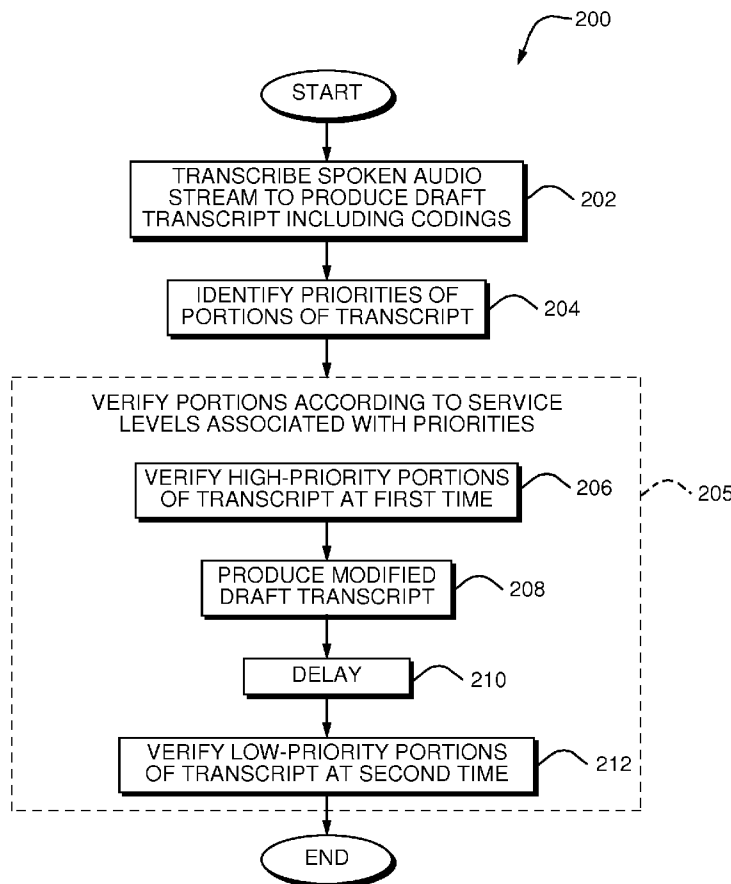




US 20140039880A1

(19) **United States**(12) **Patent Application Publication**
Koll et al.(10) **Pub. No.: US 2014/0039880 A1**(43) **Pub. Date: Feb. 6, 2014**(54) **APPLYING SERVICE LEVELS TO
TRANSCRIPTS****Publication Classification**(71) Applicant: **Multimodal Technologies, LLC**,
Pittsburgh, PA (US)(51) **Int. Cl.**
G06F 17/27 (2006.01)(72) Inventors: **Detlef Koll**, Pittsburgh, PA (US);
Michael Finke, Pittsburgh, PA (US)(52) **U.S. Cl.**
CPC **G06F 17/2785** (2013.01)
USPC **704/9**(73) Assignee: **Multimodal Technologies, LLC**,
Pittsburgh, PA (US)(57) **ABSTRACT**(21) Appl. No.: **14/052,005**(22) Filed: **Oct. 11, 2013****Related U.S. Application Data**(63) Continuation of application No. 11/766,784, filed on
Jun. 21, 2007, now Pat. No. 8,560,314.(60) Provisional application No. 60/815,689, filed on Jun.
22, 2006, provisional application No. 60/815,688,
filed on Jun. 22, 2006, provisional application No.
60/815,687, filed on Jun. 22, 2006.

Speech is transcribed to produce a draft transcript of the speech. Portions of the transcript having a high priority are identified. For example, particular sections of the transcript may be identified as high-priority sections. As another example, portions of the transcript requiring human verification may be identified as high-priority sections. High-priority portions of the transcript are verified at a first time, without verifying other portions of the transcript. Such other portions may or may not be verified at a later time. Limiting verification, either initially or entirely, to high-priority portions of the transcript limits the time required to perform such verification, thereby making it feasible to verify the most important portions of the transcript at an early stage without introducing an undue delay into the transcription process. Verifying the other portions of the transcript later ensures that early verification of the high-priority portions does not sacrifice overall verification accuracy.






FIG. 1

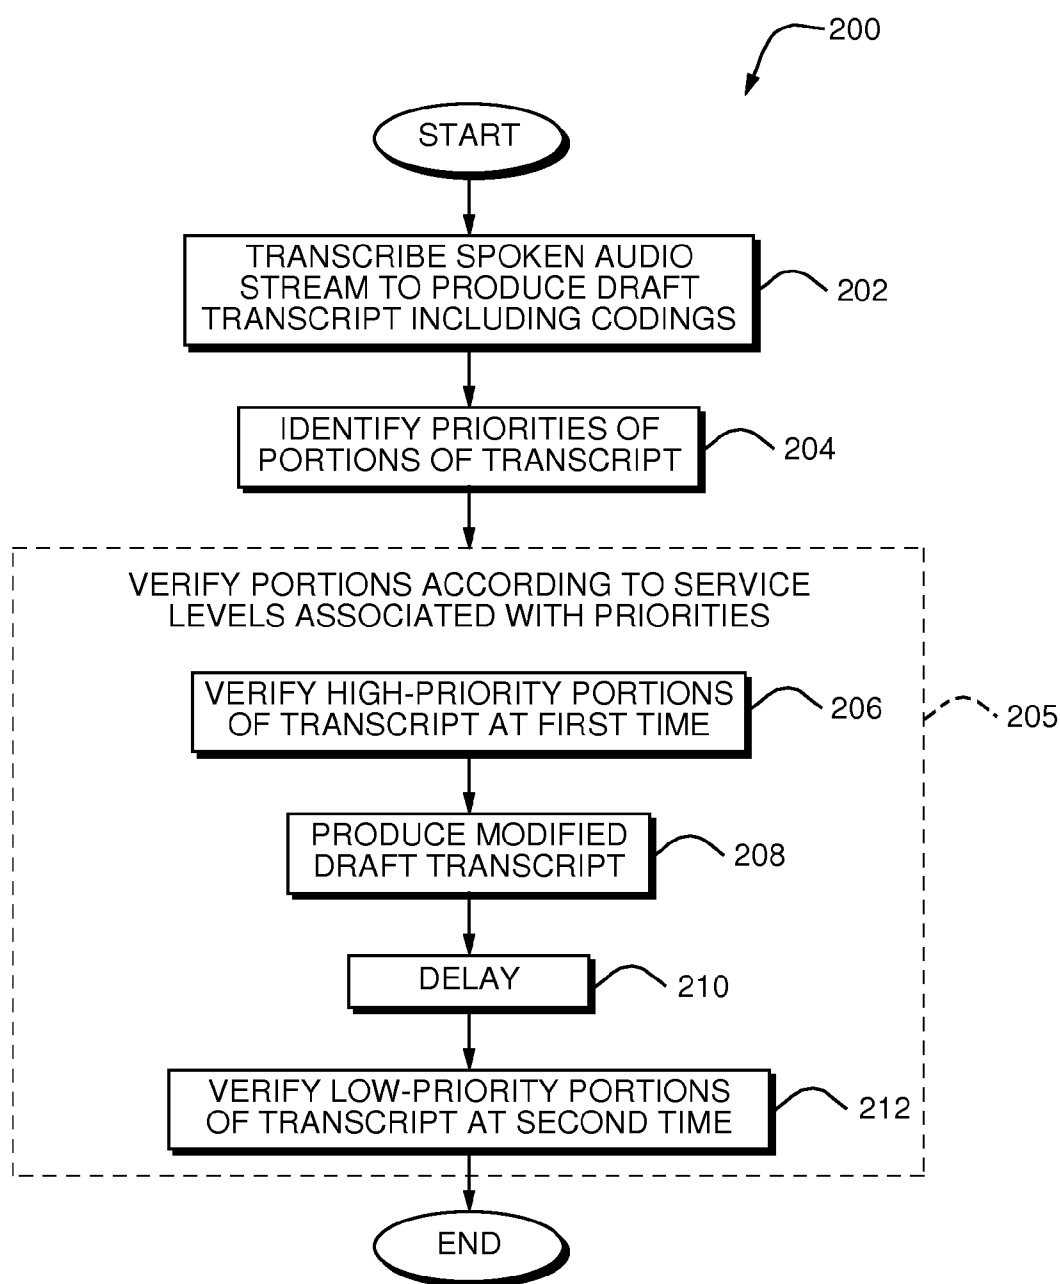


FIG. 2

APPLYING SERVICE LEVELS TO TRANSCRIPTS

CROSS REFERENCE TO RELATED APPLICATIONS

[0001] This application is related to U.S. patent application Ser. No. 10/923,517, filed on Aug. 20, 2004, entitled "Automated Extraction of Semantic Content and Generation of a Structured Document from Speech," now U.S. Pat. No. 7,584,103, which is hereby incorporated by reference herein.

BACKGROUND

[0002] It is desirable in many contexts to generate a structured textual document based on human speech. In the legal profession, for example, transcriptionists transcribe testimony given in court proceedings and in depositions to produce a written transcript of the testimony. Similarly, in the medical profession, transcripts are produced of diagnoses, prognoses, prescriptions, and other information dictated by doctors and other medical professionals.

[0003] Producing such transcripts can be time-consuming. For example, the speed with which a human transcriptionist can produce a transcript is limited by the transcriptionist's typing speed and ability to understand the speech being transcribed. Although software-based automatic speech recognizers are often used to supplement or replace the role of the human transcriptionist in producing an initial transcript, even a transcript produced by a combination of human transcriptionist and automatic speech recognizer will contain errors. Any transcript that is produced, therefore, must be considered to be a draft, to which some form of error correction is to be applied.

[0004] Producing a transcript is time-consuming for these and other reasons. For example, it may be desirable or necessary for certain kinds of transcripts (such as medical reports) to be stored and/or displayed in a particular format. Providing a transcript in an appropriate format typically requires some combination of human editing and automatic processing, which introduces an additional delay into the production of the final transcript.

[0005] Consumers of reports, such as doctors, nurses, and radiologists in the medical context, often stand to benefit from receiving reports quickly. If a diagnosis depends on the availability of a certain report, for example, then the diagnosis cannot be provided until the required report is ready. Similarly, if a doctor dictates a report of an operation into a handheld dictation device while still in the operating room, it may be desirable for a nurse to receive a transcript or other report of the dictation as soon as possible after the patient leaves the operating room. For these and other reasons it is desirable to increase the speed with which transcripts and other kinds of reports derived from speech may be produced, without sacrificing accuracy.

SUMMARY

[0006] Speech is transcribed to produce a draft transcript of the speech. Portions of the transcript having a high priority are identified. For example, particular sections of the transcript may be identified as high-priority sections. As another example, portions of the transcript requiring human verification may be identified as high-priority sections. High-priority portions of the transcript are verified at a first time, without verifying other portions of the transcript. Such other portions

may or may not be verified at a later time. Limiting verification, either initially or entirely, to high-priority portions of the transcript limits the time required to perform such verification, thereby making it feasible to verify the most important portions of the transcript at an early stage without introducing an undue delay into the transcription process. Verifying the other portions of the transcript later ensures that early verification of the high-priority portions does not sacrifice overall verification accuracy.

[0007] For example, one embodiment of the present invention is a computer-implemented method comprising: (A) identifying a first semantic meaning of a first portion of a first transcript of speech; (B) assigning a first service level to the first portion of the transcript based on the first semantic meaning; (C) identifying a second semantic meaning of a second portion of the first transcript; (D) assigning a second service level to the second portion of the transcript based on the second semantic meaning; and (E) verifying the transcript in accordance with the first and second service levels.

[0008] Another embodiment of the present invention is an apparatus comprising: first meaning identification means for identifying a first semantic meaning of a first portion of a first transcript of speech; first service level assignment means for assigning a first service level to the first portion of the transcript based on the first semantic meaning; second meaning identification means for identifying a second semantic meaning of a second portion of the first transcript; second service level assignment means for assigning a second service level to the second portion of the transcript based on the second semantic meaning; and verification means for verifying the transcript in accordance with the first and second service levels.

[0009] Another embodiment of the present invention is a computer-implemented method comprising: (A) identifying a first semantic meaning of a first transcript of first speech; (B) assigning a first service level to the first transcript based on the first semantic meaning; (C) identifying a second semantic meaning of a second transcript of second speech; (D) assigning a second service level to the second transcript based on the second semantic meaning; and (E) verifying the first and second transcripts in accordance with the first and second service levels, respectively.

[0010] Another embodiment of the present invention is an apparatus comprising: means for identifying a first semantic meaning of a first transcript of first speech; means for assigning a first service level to the first transcript based on the first semantic meaning; means for identifying a second semantic meaning of a second transcript of second speech; means for assigning a second service level to the second transcript based on the second semantic meaning; and means for verifying the first and second transcripts in accordance with the first and second service levels, respectively.

[0011] 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

[0012] FIG. 1 is a dataflow diagram of a system for verifying a transcript of speech according to one embodiment of the present invention; and

[0013] FIG. 2 is a flowchart of a method performed by the system of FIG. 1 according to one embodiment of the present invention.

DETAILED DESCRIPTION

[0014] Embodiments of the invention are directed to verifying a transcript of speech using multiple service levels. Service levels may, for example, correspond to required turnaround times for verification. A “high priority” service level may, for example, correspond to a relatively short turnaround time, while a “low priority” service level may correspond to another service level. As another example, the high priority service level may indicate that human verification of the corresponding text is required, while the low priority service level may indicate that human verification is not required.

[0015] The priorities of different portions of a transcript are identified. For example, the “impressions” section of a medical transcript may be identified as a high-priority section of the report, while other sections may be identified as low-priority sections. Verification of each portion of the document is performed according to that portion’s service level. For example, high-priority sections (such as the “impressions” section) may be verified immediately, while low-priority sections may be verified at a later time. As another example, high-priority sections may be verified by a human, while low-priority sections may be verified by software or not at all.

[0016] More specifically, referring to FIG. 1, a dataflow diagram is shown of a system 100 for selectively verifying portions of a draft transcript 106 according to one embodiment of the present invention. Referring to FIG. 2, a flowchart is shown of a method 200 performed by the system 100 of FIG. 1 according to one embodiment of the present invention.

[0017] A transcription system 104 transcribes a spoken audio stream 102 to produce a draft transcript 106 (step 202). The spoken audio stream 102 may, for example, be dictation by a doctor describing a patient visit. The spoken audio stream 102 may take any form. For example, it may be a live audio stream received directly or indirectly (such as over a telephone or IP connection), or an audio stream recorded on any medium and in any format.

[0018] The transcription system 104 may produce the draft transcript 106 using, for example, an automated speech recognizer or a combination of an automated speech recognizer and human transcriptionist. The transcription system 104 may, for example, produce the draft transcript 106 using any of the techniques disclosed in the above-referenced patent application entitled “Automated Extraction of Semantic Content and Generation of a Structured Document from Speech.” As described therein, the draft transcript 106 may include text 116 that is either a literal (verbatim) transcript or a non-literal transcript of the spoken audio stream 102. As further described therein, although the draft transcript 106 may be a plain text document, the draft transcript 106 may also, for example, in whole or in part be a structured document, such as an XML document which delineates document sections and other kinds of document structure. Various standards exist for encoding structured documents, and for annotating parts of the structured text with discrete facts (data) that are in some way related to the structured text. Examples of existing techniques for encoding medical documents include the HL7 CDA v2 XML standard (ANSI-approved since May 2005), SNOMED CT, LOINC, CPT, ICD-9 and ICD-10, and UMLS.

[0019] As shown in FIG. 1, the draft transcript 106 includes one or more codings 108, each of which encodes a “concept” extracted from the spoken audio stream 102. The term “concept” is used herein as defined in the above-referenced patent application entitled “Automated Extraction of Semantic content and Generation of a Structured Document from Speech.”

Reference numeral 108 is used herein to refer generally to all of the codings within the draft transcript 106. Although in FIG. 1 only two codings, designated 108a and 108b, are shown, the draft transcript 106 may include any number of codings.

[0020] In the context of a medical report, each of the codings 108 may, for example, encode an allergy, prescription, diagnosis, or prognosis. In general, each of the codings 108 includes a code and corresponding data, which are not shown in FIG. 1 for ease of illustration. Each of the codings 108 may be linked to text in the transcript 106. For example, coding 108a may be linked to text 118a and coding 108b may be linked to text 118b. Assume for purposes of example that linked text 118a is in a first section of the draft transcript 106 (such as an “impressions” section), and that linked text 118b is in a second section of the draft transcript 106 (such as a “patient history” section). Further details of embodiments of the codings 108 may be found in the above-referenced patent application entitled, “Verification of Data Extracted from Speech.”

[0021] A priority identifier 120 identifies verification priorities of one or more portions of the draft transcript 106 to produce portion priorities 122 (step 204). A “portion” of the draft transcript 106 may, for example, be a single coding (such as the coding 108a or coding 108b), sentence, paragraph, or section.

[0022] In general, the priority identifier 120 identifies the priority of a portion of the draft transcript 106 by identifying a semantic meaning of the portion, and identifying the priority of the portion based on the identified semantic meaning. For example, the priority identifier 120 may be configured to assign a high priority to any “impressions” sections in the draft transcript 106 and a low priority to all other sections in the draft transcript 106.

[0023] As another example, certain portions of the draft transcript 106 may require human (rather than software-based or other automated) verification. Typically, all portions of the draft transcript 106 require human verification. Some portions, however, such as certain codings representing data which are used solely for automatic computer processing, may not be rendered to the reviewer (e.g., physician) for review and not require human verification. The priority identifier 120 may be configured to assign a high priority to any portions of the draft transcript 106 which require human verification, and to assign a low priority to all other portions of the draft transcript 106. The priority identifier 120 may identify portions of the draft transcript 106 requiring human verification in any of a variety of ways. For example, the priority identifier 120 may be configured to identify certain codings, or certain types of codings, as requiring human verification, in which case the priority identifier 120 may assign a high priority to such codings.

[0024] In general, the portions of the draft transcript 106 are verified according to service levels associated with their priorities (step 205). For example, a “high-priority” service level may be applied to portions of the transcript 106 having a high priority. Similarly, a “low-priority” service level may be applied to portions of the transcript 106 having a low priority. Since the priorities 122 of the document portions are identified based on the semantic meanings of the portions, the service levels that are applied to the portions are based on the semantic meanings of the portions.

[0025] One example of applying such service levels is illustrated by steps 206-212 in FIG. 2. More specifically, an initial

transcript verifier **124** verifies, at a first time, only those portions of the draft transcript **106** which the priority identifier **120** has identified as having a high verification priority (step **206**). The portion priorities **122** may, for example, include links to the corresponding portions of the draft transcript **106**. The initial transcript verifier **124** may use such links to identify which portions of the draft transcript **106** have high verification priorities and to verify only those portions of the draft transcript **106**. The initial transcript verifier **124** produces initial transcript verification status identifiers **128** which indicate the verification status of each of the high-priority portions of the draft transcript **106**. A verification status may, for example, indicate whether the corresponding portion of the draft transcript **106** is correct or incorrect.

[0026] Any of a variety of techniques may be used to verify the high-priority portions of the draft transcript **106**, examples of which may be found in the above-referenced patent application entitled, "Verification of Data Extracted from Speech." For example, the high-priority portions of the draft transcript **106** may be presented to a human reviewer, who may be the same person who dictated the spoken audio stream. In fact, the speaker of the spoken audio stream **102** may still be dictating the remainder of the spoken audio stream **102** while portions of the draft transcript **106** representing previous parts of the spoken audio stream **102** are presented to the reviewer for verification.

[0027] The verification process performed by the initial transcript verifier **124** may include correcting any portions of the draft transcript **106** that are found to be incorrect. For example, the reviewer may provide input to correct one or more portions of the draft transcript **106** presented for review. The initial transcript verifier **124** may therefore produce a modified draft transcript **126**, which includes any corrections to the codings **108** or other modifications made by the initial transcript verifier **124** (step **208**).

[0028] Note that the initial transcript verification status identifiers **128** and the modified draft transcript **126** may be combined. For example, portions of the modified draft transcript **126** may be tagged with their own verification statuses.

[0029] There is some delay **130** after the initial verification (step **210**). The delay **130** may occur, for example, while the modified draft transcript **126** is provided to a nurse or physician for use in providing medical care.

[0030] After the delay **130**, a subsequent transcript verifier **132** verifies portions of the modified draft transcript **126** other than those which were verified by the initial transcript verifier **124**, thereby producing a final transcript **134** (step **212**). For example, the subsequent transcript verifier **132** may verify all of the portions which were not verified by the initial transcript verifier **124**. The subsequent transcript verifier **132** may use the same verification techniques as the initial transcript verifier **124**. Note that although the initial transcript verifier **124** and subsequent transcript verifier **132** are illustrated as two distinct components in FIG. 1, they may be integrated into a single component (e.g., a single software program).

[0031] The techniques just described may be viewed as implementing two tiers of "service level," in which each tier corresponds to a different turnaround time requirement. The first (high-priority) tier of service level corresponds to a relatively short turnaround time, and the second (low-priority) tier of service level corresponds to a relatively long turnaround time. The priorities **122** assigned to portions of the

draft transcript **106** may, however, be used to implement service levels corresponding to characteristics other than turnaround times.

[0032] For example, a first tier of service level may require the corresponding portion of the draft transcript **106** to be verified by a human, while a second tier of service level may not require human verification. For example, the second tier may require software verification or not require any verification. In this example, the method of FIG. 2 would include only steps **202-208**. More specifically, step **202** would include transcribing the spoken audio stream to produce the draft transcript **106**, as described above. Step **204** would include identifying priorities **122** of portions of the draft transcript **106**. In this example, a high priority would indicate that human verification is required, while a low priority would indicate that human verification is not required.

[0033] Step **206** would include verifying only those portions of the draft transcript **106** having high priorities (i.e., only those portions of the draft transcript **106** requiring human verification). Such verification may, for example, be performed by providing the draft transcript **106** to a human reviewer (such as the dictating physician) for review.

[0034] The method **200** would then terminate, without verifying the low-priority portions of the draft transcript **106** because such portions have been identified as not requiring human verification. In yet another embodiment, a high service level would require verification by the dictating physician, a medium service level would require verification by a data entry clerk, and a low service level would not require any human verification.

[0035] Embodiments of the present invention have a variety of advantages. For example, transcripts of physician-dictated reports typically are not provided to nurses and others for use in patient care soon after a draft of the transcript has been produced because it has not been possible to verify the accuracy of the transcript quickly enough. As a result, draft transcripts typically are withheld until they have been verified, which can introduce a significant delay before such transcripts may be used.

[0036] One benefit of the techniques disclosed herein is that they make it possible to provide transcripts to consumers of such transcripts both quickly and with increased accuracy by limiting the amount of the transcript that is verified, at least at the outset. In particular, by limiting the verification, either initially or entirely, to high-priority portions of the draft transcript **106**, the techniques illustrated in FIGS. 1 and 2 increase the accuracy of those portions of the draft transcript **106** for which accuracy is most critical, without introducing the significant delay that would otherwise be incurred by verifying the entire transcript.

[0037] Furthermore, the techniques disclosed herein do not create significant additional work for the physician or other dictator/reviewer of the transcript, and may thereby reduce the total amount of time they need to devote to performing their tasks. If, for example, the initial transcript verifier **124** presents high-priority portions of the draft transcript **106** to a physician for review immediately after the physician dictates such portions, or even while the physician is dictating such portions, the physician may verify and correct any errors in such portions more efficiently than if those portions of the transcript were provided to the physician for review an hour, day, or week later because of the switching costs involved. Furthermore, overall accuracy may be increased because the dictator/reviewer of the draft transcript **106** may have a

clearer memory of the intended content of the draft transcript **106** is presented to him or her immediately after dictating it, rather than at a later time.

[0038] These techniques are particularly useful in contexts in which a transcript needs to be made available immediately to someone other than the dictator of the transcript. For example, the “impressions” section of a radiology report needs to be made available immediately for emergency room radiology studies, while the rest of the document may be completed at a later time. In order to guarantee immediate turnaround, the “impressions” section of the draft transcript **106** may be provided to the dictating physician immediately for self-correction, while leaving the remainder of the document to be corrected at a later time by professional transcriptionists.

[0039] As another example, the medication and allergy sections of a progress note may be interpreted and populated immediately in an electronic medical records (EMR) system to be available for order entry (i.e., prescription medications) and decision support at point of care. Other data elements (e.g., history of present illness) may be verified at a later time.

[0040] Another benefit of techniques disclosed herein is that they may be combined with the techniques disclosed in the above-referenced patent application entitled “Automatic Decision Support” for applying automated clinical decision support to transcripts. For example, the modified draft transcript **126** may be provided to an automated clinical decision support system, as described in the above-referenced patent application, to provide clinical decision support quickly. Such techniques may, for example, be used to provide a quick indication of whether the draft transcript **126** indicates a dangerous drug-drug allergy which requires attention by the dictating physician or other care provider. The techniques disclosed herein facilitate the provision of such rapid clinical decision support by enabling verification of high-priority portions of the draft transcript **106** to be performed soon after the draft transcript **106** is produced. If verification of such high-priority portions had to wait until all portions of the draft transcript **106** were verified, then application of automated clinical decision support to the high-priority portions would have to wait as well.

[0041] In this regard, portions of the draft transcript **106** which are necessary to provide as input a clinical decision support system may be assigned a high priority. As a result, the initial transcript verifier **124** will verify those portions of the draft transcript **106** which are necessary to provide as input to the clinical decision support system. This ensures that any desired clinical decision support can be applied to the modified draft transcript **126** produced by the initial transcript verifier **124**, without incurring additional delay.

[0042] 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.

[0043] Although certain examples provided herein involve documents generated by a speech recognizer, this is not a requirement of the present invention. Rather, the techniques disclosed herein may be applied to any kind of document,

regardless of how it was generated. Such techniques may, for example, be used in conjunction with documents typed using conventional text editors.

[0044] The spoken audio stream **102** may be any audio stream, such as a live audio stream received directly or indirectly (such as over a telephone or IP connection), or an audio stream recorded on any medium and in any format. In distributed speech recognition (DSR), a client performs preprocessing on an audio stream to produce a processed audio stream that is transmitted to a server, which performs speech recognition on the processed audio stream. The audio stream may, for example, be a processed audio stream produced by a DSR client.

[0045] The invention is not limited to any of the described domains (such as the medical and legal fields), but generally applies to any kind of documents in any domain. For example, although the reviewer **138** may be described herein as a physician, this is not a limitation of the present invention. Rather, the reviewer **138** may be any person. Furthermore, documents used in conjunction with embodiments of the present invention may be represented in any machine-readable form. Such forms include plain text documents and structured documents represented in markup languages such as XML. Such documents may be stored in any computer-readable medium and transmitted using any kind of communications channel and protocol.

[0046] Although certain examples described herein include only two priorities—high and low—these are merely examples of priorities and do not constitute limitations of the present invention. Rather, the techniques disclosed herein may be applied to any number of priorities defined and/or labeled in any manner. For example, there may be high, medium, and low priorities. As another examples, priorities may specify maximum turnaround times, such as one minute, one hour, and one day.

[0047] Furthermore, although in certain examples herein the high priority service level is described as providing “immediate” verification, this is not a requirement of the present invention. Rather, service levels may be defined and applied in any manner. For example, even when service levels define turnaround requirements, the highest-priority service level need not require immediate turnaround, but rather may require or allow any amount of delay.

[0048] Although certain examples disclosed herein apply different service levels to different portions of a single transcript, the same or similar techniques may be used to apply different service levels to different documents. In other words, a first (e.g., high-priority) service level may be applied to a first entire document and a second (e.g., low-priority) service level may be applied to a second entire document. The first and second documents may then be verified in accordance with their respective service levels.

[0049] For example, the transcription system **104** may transcribe a first spoken audio stream to produce a first draft transcript and apply automatic decision support to the first draft transcript using the techniques disclosed in the above-referenced patent application entitled, “Automatic Decision Support.” A first service level may be associated with the first draft transcript based on the results of the automatic decision support. For example, if automatic decision support determines that the draft transcript does not include any critical errors, then a low-priority service level may be associated with the first draft transcript.

[0050] The transcription system **104** may also transcribe a second spoken audio stream to produce a second draft transcript and apply automatic decision support to the second draft transcript using the above-described techniques. A second service level may be associated with the second draft transcript based on the results of the automatic decision support. For example, if automatic decision support determines that the second draft transcript contains a critical error (such as by describing a prescription that would result in a drug-drug allergy), then a high-priority service level may be associated with the second draft transcript.

[0051] The first and second draft transcripts may then be verified in accordance with their respective service levels. For example, the second draft transcript may be presented for verification before the first draft transcript, even though the first draft transcript was transcribed before the second draft transcript, because the second draft transcript is associated with a higher service level than the first draft transcript. As one example of processing the two draft transcripts in accordance with their respective service levels, a phone call might be made immediately to the dictating physician of the second draft transcript to alert him or her to the identified error. As another example, the second draft transcript might be placed at the top of the queue of documents to which full verification is to be applied. In either case, the (low-priority) first draft transcript may be verified at a later time. Such application of different service levels to different documents shares benefits of the techniques disclosed herein for applying different service levels to different portions of a single document.

[0052] Although the verification statuses **130** may include statuses such as “correct” and “incorrect,” the invention is not so limited. Rather, the verification statuses **130** may take any of a variety of forms. Further examples of the verification statuses and techniques for producing them may be found in the above-referenced patent application entitled, “Verification of Extracted Data.”

[0053] Although certain individual types of service levels are described herein, such service levels are merely examples and do not constitute limitations of the present invention. For example, the above-referenced patent application entitled, “Automatic Decision Support,” discloses techniques for performing real-time automatic decision support, in which decision support is applied to a transcript-in-progress while the physician (or other speaker) is still dictating the report. If the decision support system identifies a critical error or other problem requiring the immediate attention of the speaker-reviewer, the system notifies the speaker-reviewer of the problem while the report is still being dictated. Such immediate, real-time presentation of a partial transcript for review may be viewed as a service level, in comparison, for example, to service levels in which lower-priority errors are not deferred until dictation of the report is complete.

[0054] Furthermore, service levels not disclosed herein may be used in conjunction with the techniques disclosed herein. Furthermore, service levels may be combined with each other in various ways. For example, a “high-priority” service level may be applied to any portions of the draft transcript **106** which:

- (1) are within predetermined sections of the transcript **106**; or
- (2) which require human verification.

[0055] The verification performed by the initial transcript verifier **124** and the subsequent transcript verifier **132** may be performed in whole or in part by a human, such as the dictator of the draft transcript **106**. To facilitate such verification, the high-priority portions of the draft transcript **106** may be called out to the reviewer, such as by highlighting them or displaying them in a different color within a display of the

draft transcript **106**. Such a display may, for example, be displayed to the reviewer while the reviewer is dictating the contents of the transcript **106**.

[0056] The techniques described above may be implemented, for example, in hardware, software, 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.

[0057] 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.

[0058] 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.

What is claimed is:

1. A computer-implemented method comprising:
 - (A) identifying a first semantic meaning of a first portion of a first transcript of speech;
 - (B) assigning a first service level to the first portion of the transcript based on the first semantic meaning;
 - (C) identifying a second semantic meaning of a second portion of the first transcript;
 - (D) assigning a second service level to the second portion of the transcript based on the second semantic meaning; and
 - (E) at a processor, verifying the transcript in accordance with the first and second service levels.

* * * * *