



(22) Date de dépôt/Filing Date: 2007/09/28

(41) Mise à la disp. pub./Open to Public Insp.: 2007/12/06

(30) Priorités/Priorities: 2006/09/29 (US11/540,353);
2007/09/27 (USPCT/US07/79779)

(51) Cl.Int./Int.Cl. *G10L 11/00* (2006.01),
H04M 3/42 (2006.01)


(71) Demandeur/Applicant:
VERINT SYSTEMS INC., US

(72) Inventeur/Inventor:
BLAIR, CHRISTOPHER D., GB

(74) Agent: SIM & MCBURNEY

(54) Titre : SYSTEMES ET METHODES PERMETTANT D'ANALYSER DES SESSIONS DE COMMUNICATION AU
MOYEN DE FRAGMENTS

(54) Title: SYSTEMS AND METHODS FOR ANALYZING COMMUNICATION SESSIONS USING FRAGMENTS



(57) Abrégé/Abstract:

Systems and methods for analyzing communication sessions using fragments are provided. In this regard, a representative method includes: delineating fragments of an audio component of a communication session, each of the fragments being attributable to a party of the communication session and representing a contiguous period of time during which that party was speaking; and automatically assessing quality of at least some of the fragments such that a quality assessment of the communication session is determined.

ABSTRACT

Systems and methods for analyzing communication sessions using fragments are provided. In this regard, a representative method includes: delineating fragments of an audio component of a communication session, each of the fragments being attributable to a party of the communication session and representing a contiguous period of time during which that party was speaking; and automatically assessing quality of at least some of the fragments such that a quality assessment of the communication session is determined.

SYSTEMS AND METHODS FOR ANALYZING COMMUNICATION SESSIONS USING FRAGMENTS

BACKGROUND

[0001] It is desirable in many situations to record communications, such as telephone calls. This is particularly so in a contact center in which many agents may be handling hundreds of telephone calls each every day. Recording of these telephone calls can allow for quality assessment of agents, improvement of agent skills and/or dispute resolution, for example.

[0002] In this regard, assessment of call quality is time consuming and very subjective. For instance, a telephone call may last from a few seconds to a few hours and may be only one part of a customer transaction or may include several independent transactions. The demeanor of the caller is also influenced by events preceding the actual conversation – for example, the original reason for the call; the time spent waiting for the call to be answered or the number of times the customer has had to call before getting through to the right person.

[0003] Assessing the “quality” of a telephone call is therefore difficult and subject to error, even when done by an experienced supervisor or full-time quality assessor. Typically, the assessment of a call is structured according to a pre-defined set of criteria and sub-criteria. Some of these may relate to the initial greeting, the assessment of the reason for the call, the handling of the core reason for the call, confirming that the caller is satisfied with the handling of the call, and leaving the call.

[0004] Automation of the assessment process by provision of standardized forms and evaluation profiles have made such assessment more efficient, but it is still impractical to assess more than a tiny percentage of calls. Moreover, even with a structured

evaluation form, different assessors will evaluate a call differently with quite a wide variation of scores.

SUMMARY

[0005] In this regard, systems and methods for analyzing communication sessions using fragments are provided. An embodiment of a method comprises: delineating fragments of an audio component of a communication session, each of the fragments being attributable to a party of the communication session and representing a contiguous period of time during which that party was speaking; and automatically assessing quality of at least some of the fragments such that a quality assessment of the communication session is determined.

[0006] An embodiment of such a system comprises a communication analyzer operative to: delineate fragments of an audio component of a communication session, each of the fragments being attributable to a party of the communication session and representing a contiguous period of time during which that party was speaking; and automatically assess quality of at least some of the fragments such that a quality assessment of the communication session is determined.

[0007] Computer readable media also are provided that have computer programs stored thereon for performing computer executable methods. In this regard, an embodiment of such a method comprises: delineating fragments of an audio component of a communication session, each of the fragments being attributable to a party of the communication session and representing a contiguous period of time during which that party was speaking; and automatically assessing quality of at least some of the fragments such that a quality assessment of the communication session is determined.

[0008] Other systems, methods, features and/or advantages of this disclosure will be or become apparent to one with skill in the art upon examination of the following drawings and detailed description. It is intended that all such additional systems, methods, features, and advantages be included within this description and be within the scope of the present disclosure.

BRIEF DESCRIPTION

[0009] Many aspects of the disclosure can be better understood with reference to the following drawings. The components in the drawings are not necessarily to scale, emphasis instead being placed upon clearly illustrating the principles of the present disclosure. Moreover, in the drawings, like reference numerals designate corresponding parts throughout the several views. While several embodiments are described in connection with these drawings, there is no intent to limit the disclosure to the embodiments disclosed herein.

[0010] FIG. 1 is a schematic diagram illustrating an embodiment of a system for analyzing communication sessions using fragments.

[0011] FIG. 2 is a flowchart illustrating functionality (or methods steps) that can be preformed by the embodiment of the system for analyzing communication sessions using fragments of FIG. 1.

[0012] FIG. 3 is a schematic representation of an exemplary communication session and corresponding call fragments.

[0013] FIG. 4 is a flowchart illustrating functionality (or methods steps) that can be preformed by another embodiment of a system for analyzing communication sessions using fragments.

[0014] FIG. 5 is a diagram illustrating an embodiment of a system for analyzing communication sessions using fragments that is implemented by a computer.

DETAILED DESCRIPTION

[0015] Systems and methods for analyzing communication sessions using fragments are provided. In this regard, several exemplary embodiments will be described in which a recording of a telephone call is divided into more manageable fragments. By way of example, each of the fragments can be configured as contiguous speech of a party of the call. Specific behaviors can, therefore, be identified automatically as each fragment can be assessed more easily and unambiguously than if the behaviors were attempted to be identified from within an undivided call. By automating the assessment of call quality, a higher proportion of calls can be analyzed and hence a higher proportion of problem behaviors, processes and issues identified and addressed with less effort and cost than alternative manual strategies.

[0016] In this regard, FIG. 1 is a schematic diagram illustrating an embodiment of a system for analyzing communication sessions using fragments. As shown in FIG. 1, system 100 incorporates a communication analyzer 110 that is configured to analyze audio components of communications. In FIG. 1, the audio component (not shown) is associated with a communication session that is occurring between a caller 112 and an agent 114 via a communication network 116. In this embodiment, the agent is associated with a contact center that comprises numerous agents for interacting with customers, e.g., caller 112.

[0017] One should note that network 116 can include one or more different networks and/or types of networks. As a non-limiting, example, communications network 116 can include a Wide Area Network (WAN), the Internet, and/or a Local Area Network

(LAN). Additionally, the communication analyzer can receive information corresponding to the communication session directly or from one or more various components that are not illustrated in FIG. 1. By way of example, the information can be provided from a long term storage device that stores recordings of the communication session, with the recordings being provided to the storage device by a recorder. Additionally or alternatively, the recordings could be provided directly from such a recorder.

[0018] In operation, the analyzer of FIG. 1 performs various functions (or method steps) as depicted in the flowchart of FIG. 2. As shown in FIG. 2, the functions include (as depicted in block 210) delineating an audio component of a communication session into fragments. In particular, in this embodiment, each of the fragments is attributable to a party of the communication session and represents a contiguous period of time during which that party was speaking. By way of example, one such fragment could involve a recording (e.g., 4 seconds in duration) of the speech of agent 114 during a communication session with customer 112, in which the agent greeted the customer. As shown in block 212, the analyzer also automatically assesses quality of at least some of the fragments such that a quality assessment of the communication session is determined.

[0019] In some embodiments, the parties to a communication session are recorded separately. In other embodiments, a session can be recorded in stereo, with one channel for the customer and one for the agent.

[0020] A vox detection analyzer of a communication analyzer can be used to determine when each party is talking. Such an analyzer typically detects an audio level above a pre-determined threshold for a sustained period (the “vox turn-on time”). Absence of audio is then determined by the audio level being below a pre-

determined level (which may be different from the first level) for a pre-determined time (which may be different from the previous “turn-on” time). By identifying audio presence on each of the two channels of recording of a call results in a time series through the call that identifies who, if anyone, is talking at any given time in the series.

[0021] Once audio presence is determined, the call can be broken into “fragments” representing the period in which each party talks on the call. In this regard, a fragment can be delimited by one or more of the following:

- i) the start or end of the call;
- ii) the other party starting to speak and the previous party stopping speaking;
- iii) a “significant” pause – a period greater than a typical interval between one party finishing speaking and the other party beginning speaking. This interval may be pre-determined or determined by examining the actual intervals between the parties speaking on this call. If the call involves more than a few alternations of which party is speaking, these alternations can typically be grouped. For instance, one group could be “normal turns of dialog” in which the intervals are on the order of a fraction of a second to one or two seconds and another group could be “delays” in which the dialog is hesitant or significantly delayed for some reason; and
- iv) a “significant interruption” – a period during which both parties are speaking and which is longer than typical confirmatory feedback (e.g., “uh-huh”) that is heard every few seconds in a normal interaction.

[0022] A schematic representation of an exemplary communication session and corresponding call fragments is depicted in FIG. 3. As shown in FIG. 3, the communication session is a sequence of audio components (depicted as blocks) of an interaction between an agent and a customer that takes place over a 30 second time period. In particular, the agent speaks for the first 4 seconds, followed by a 1 second

pause. The customer then speaks for 7 seconds followed by a 1 second pause.

Thereafter, the agent speaks for 7 seconds, the last 2 seconds of which the customer begins speaking, with the customer continuing to speak for another 2 seconds. After another 1 second pause, the agent speaks for 5 seconds after which, without pause, the customer speaks for 2 seconds and the communication session ends. Notably, although not shown in this example, the reason for delimiting the fragment can be correlated with the fragment itself (e.g., alongside the fragment) resulting in a sequence of records.

[0023] Having broken a call into fragments, the system can analyze the sequence and duration of the fragments. By way of example, for each fragment, some embodiments can determine one or more of the following:

- i) which party is speaking (customer or agent);
- ii) which party spoke in the previous fragment;
- iii) which party speaks in the next fragment;
- iv) the delay between the previous fragment and this one;
- v) the delay between this fragment and the next;
- vi) a link to the previous fragment;
- vii) a link to the next fragment;
- viii) a transcript of the words and/or phonemes contained within the fragment – determined by phonetic analysis using a phonetic analyzer and/or speech recognition analysis using a speech recognition engine;
- ix) a time sequence of the amplitude of the audio of the speaking party throughout the fragment;
- x) an estimate of periods of loud speech or shouting. This may be determined by the fact that the audio level clipped as well as or instead of exceeded a

specified level or relative level compared to the call as a whole or the level of audio from the other party;

- xi) the time from the start of the call to the start of this fragment;
- xii) the duration of this fragment; and
- xiii) the time from this fragment to the end of the call.

In some embodiments, statistics of the call can be deduced from the individual call fragment data. These may include one or more of:

- i) number of call fragments;
- ii) number of times the speaker changed;
- iii) average duration of customer speaking;
- iv) average duration of agent speaking;
- v) percentage of total talk time that agent spoke;
- vi) percentage of total talk time that customer spoke;
- vii) percentage of total call time during which neither party spoke;
- viii) percentage of time that both parties spoke;
- ix) maximum duration of “interruptions” – defined for example, as periods of greater than 1 second during which both parties talked; and
- x) emotion indication – for example, pitch values and/or trends throughout the call.

[0024] As mentioned above, a communication analyzer can automatically assess quality of a communication session by assessing quality of at least some of its fragments. In order to accomplish quality assessment, various techniques can be used. By way of example, fragment training can be used, in which manual scoring is applied to one or more fragments and then the system applies comparable scoring to fragments that are evaluated to be similar.

[0025] In this regard, in some embodiments, individual fragments or sequences of two or more successive fragments are presented to the user of the system, typically with a clear indication of which party is speaking and the delay between the two fragments. The user listens to some or all of the fragments and then indicates, such as via a form on a screen provided by a scoring analyzer, whether the fragments relate to a good, bad or “indifferent” interaction, for example. In many cases, the isolated fragments will not indicate a particularly good or bad experience but in a small percentage of cases such fragments can indicate a particularly good or bad experience. By way of example, a long delay between two successive fragments can be considered “bad” but in other cases, the words uttered, the tone or volume of the utterance may indicate a good or bad experience. This manual (human) assessment of the quality of the fragment sequence can be stored and used to drive machine learning algorithms.

[0026] In some embodiments, in contrast to a scoring of good, bad or indifferent, a continuous scale (e.g., 0-10 rating) can be used. Additionally, multiple criteria may be presented, each of which the user can choose to provide feedback on, such as “Customer empathy” and “Persuasiveness” for example. In many cases, any particular fragment or fragment pair will not be particularly good or bad but as long as those cases that are at one extreme or the other are identified, the system will receive valuable input.

[0027] In many cases, however, the fragments presented to the user may not show anything significant but may indicate that the previous or next fragments may provide more valuable input. Because of this, the user may be presented with controls that allow the user to play the previous and/or next fragment. Thus, the user can provide feedback on those fragments and/or move on to the next or previous fragment.

[0028] Where users assess whole calls, the overall quality assessment of the call and the individual criteria/sub-criteria may be noted. These are then applied to either all fragments or, where specific criteria are explicitly linked to particular regions of the call (e.g. “Quality of Greeting”, “Confirmation of resolution”), to the fragments of the call according to a weighting function. In those embodiments that use weighting, a different weighting can be applied to each fragment according to the distance of that fragment from the start of the call, the end of the call, or from some other known point within the call. It should be noted that point from which the fragment is measured for weighting purposes can be identified by an event that occurred during the call. The fragment can be subsequently stored with a timestamp linking the fragment to that point, e.g., event, in the call.

[0029] As mentioned before, manual quality assessments can then be used by the system for enabling automated scoring of other fragments that have not been manually scored. Additionally or alternatively, some embodiments can be provided with a number of heuristics, such as predefined rules, that the system can use during automated analysis by a scoring analyzer. In this regard, such rules can involve one or more of the following:

- i) calls in which the customer to agent speech ratio is $> 80/20$ or less than $20/80$ are scored as “bad”;
- ii) interruptions of >1 second are “bad”;
- iii) delays between fragments of >2 seconds are “bad”; and
- iv) audio volumes above X are “bad”.

[0030] The human input, e.g., predefined rules and/or examples of manually assessed calls/fragments, can be used as input for a variety of machine learning techniques such as neural nets and Bayesian filters expert systems, for example. By identifying the

characteristics of the call fragments that lead to the assessments given, a system employing such a technique can learn to identify the relevant characteristics that differentiate “good” from “bad” calls.

[0031] An example of this approach is a Bayesian probability assessment of the content of a call fragment. In such an approach, a transcript of a call may be processed and the frequency of the occurrence of each word within the customer’s speech is stored. The proportion of “good” fragments in which each word occurs and the proportion of “bad” fragments in which each word occurs is then noted. These probabilities can then be used to assess whether other fragments are likely to be “good” or “bad” based on the words within those and the likelihood of each of the words to be found in a “good” or “bad” fragment. From the many words within a given fragment, those that provide the strongest discrimination of good versus bad fragment can be used and the remainder discarded. Of the N strongest indicators, an overall assessment can be made of good versus bad.

[0032] Typically, the other attributes of a fragment, such as those described above, can be used as potential indicators of the good/bad decision. These inputs may be provided to train a neural network or other machine learning system.

[0033] In some embodiments, feedback can be used to further enhance analysis. Specifically, since a high proportion of fragment sequences do not indicate particularly good (or bad) experiences, it can be beneficial if a system presents to a user those fragments that is has identified as good or bad. By presenting these fragments and showing the assessment (good or bad) that the system has determined, the user can be enabled to confirm or correct the assessment. This input can then be fed back into the training algorithm either reinforcing the correct assessment or helping to avoid repetition of the mistake made.

[0034] In this regard, FIG. 4 is a flowchart depicting functionality of an embodiment of a system that incorporates the use of feedback. As shown in FIG. 4, the functionality (or method) may be construed as beginning at block 410, in which a communication session is recorded. In block 412, an audio component of the communication session is delineated as a sequence of fragments. In block 414, inputs (such as manual scoring of a subset of the fragments and/or heuristics) are received for enabling automated scoring of at least some of the fragments. In block 416, the inputs are used in analyzing the fragments such that scores for at least some of the fragments that were not manually scored are produced. It should be noted that in some embodiments, the fragments that are manually evaluated may not be associated with the communication session that is being automatically scored.

[0035] In block 418, scores produced during automated analysis are presented to a user for review. By way of example, the scores can be presented to the user via a graphical user interface displayed on a display device. Then, in block 420, inputs from the user either confirming or correcting the scores are provided, with these inputs being used to update the analysis algorithm of the communication analyzer.

[0036] FIG. 5 is a schematic diagram illustrating an embodiment of a communication analyzer that is implemented by a computer. Generally, in terms of hardware architecture, voice analyzer 500 includes a processor 502, memory 504, and one or more input and/or output (I/O) devices interface(s) 506 that are communicatively coupled via a local interface 508. The local interface 508 can include, for example but not limited to, one or more buses or other wired or wireless connections. The local interface may have additional elements, which are omitted for simplicity, such as controllers, buffers (caches), drivers, repeaters, and receivers to enable communications.

- [0037] Further, the local interface may include address, control, and/or data connections to enable appropriate communications among the aforementioned components. The processor may be a hardware device for executing software, particularly software stored in memory.
- [0038] The memory can include any one or combination of volatile memory elements (*e.g.*, random access memory (RAM, such as DRAM, SRAM, SDRAM, *etc.*)) and nonvolatile memory elements (*e.g.*, ROM, hard drive, tape, CDROM, *etc.*). Moreover, the memory may incorporate electronic, magnetic, optical, and/or other types of storage media. Note that the memory can have a distributed architecture, where various components are situated remote from one another, but can be accessed by the processor. Additionally, the memory includes an operating system 510, as well as instructions associated with a speech recognition engine 512, a phonetic analyzer 514, a vox detection analyzer 516 and a scoring analyzer 518. Exemplary embodiments of each of which are described above.
- [0039] It should be noted that embodiments of one or more of the systems described herein could be used to perform an aspect of speech analytics (*i.e.*, the analysis of recorded speech or real-time speech), which can be used to perform a variety of functions, such as automated call evaluation, call scoring, quality monitoring, quality assessment and compliance/adherence. By way of example, speech analytics can be used to compare a recorded interaction to a script (*e.g.*, a script that the agent was to use during the interaction). In other words, speech analytics can be used to measure how well agents adhere to scripts, identify which agents are “good” sales people and which ones need additional training. As such, speech analytics can be used to find agents who do not adhere to scripts. Yet in another example, speech analytics can measure script effectiveness, identify which scripts are effective and which are not,

and find, for example, the section of a script that displeases or upsets customers (e.g., based on emotion detection). As another example, compliance with various policies can be determined. Such may be in the case of, for example, the collections industry where it is a highly regulated business and agents must abide by many rules. The speech analytics of the present disclosure may identify when agents are not adhering to their scripts and guidelines. This can potentially improve collection effectiveness and reduce corporate liability and risk.

[0040] In this regard, various types of recording components can be used to facilitate speech analytics. Specifically, such recording components can perform one or more various functions such as receiving, capturing, intercepting and tapping of data. This can involve the use of active and/or passive recording techniques, as well as the recording of voice and/or screen data.

[0041] It should be noted that speech analytics can be used in conjunction with such screen data (e.g., screen data captured from an agent's workstation/PC) for evaluation, scoring, analysis, adherence and compliance purposes, for example. Such integrated functionalities improve the effectiveness and efficiency of, for example, quality assurance programs. For example, the integrated function can help companies to locate appropriate calls (and related screen interactions) for quality monitoring and evaluation. This type of "precision" monitoring improves the effectiveness and productivity of quality assurance programs.

[0042] Another aspect that can be accomplished involves fraud detection. In this regard, various manners can be used to determine the identity of a particular speaker. In some embodiments, speech analytics can be used independently and/or in combination with other techniques for performing fraud detection. Specifically, some embodiments can involve identification of a speaker (e.g., a customer) and correlating this identification

with other information to determine whether a fraudulent claim for example is being made. If such potential fraud is identified, some embodiments can provide an alert. For example, the speech analytics of the present disclosure may identify the emotions of callers. The identified emotions can be used in conjunction with identifying specific concepts to help companies spot either agents or callers/customers who are involved in fraudulent activities. Referring back to the collections example outlined above, by using emotion and concept detection, companies can identify which customers are attempting to mislead collectors into believing that they are going to pay. The earlier the company is aware of a problem account, the more recourse options they will have. Thus, the speech analytics of the present disclosure can function as an early warning system to reduce losses.

[0043] Additionally, included in this disclosure are embodiments of integrated workforce optimization platforms, as discussed in U.S. Application No. 11/359,356, filed on February 22, 2006, entitled "Systems and Methods for Workforce Optimization." which is hereby incorporated by reference in its entirety. At least one embodiment of an integrated workforce optimization platform integrates: (1) Quality Monitoring/Call Recording - voice of the customer; the complete customer experience across multimedia touch points; (2) Workforce Management - strategic forecasting and scheduling that drives efficiency and adherence, aids in planning, and helps facilitate optimum staffing and service levels; (3) Performance Management - key performance indicators (KPIs) and scorecards that analyze and help identify synergies, opportunities and improvement areas; (4) e-Learning - training, new information and protocol disseminated to staff, leveraging best practice customer interactions and delivering learning to support development; and/or (5) Analytics - deliver insights from customer interactions to drive business performance. By way of example, the

integrated workforce optimization process and system can include planning and establishing goals – from both an enterprise and center perspective – to ensure alignment and objectives that complement and support one another. Such planning may be complemented with forecasting and scheduling of the workforce to ensure optimum service levels. Recording and measuring performance may also be utilized, leveraging quality monitoring/call recording to assess service quality and the customer experience.

[0044] One should note that the flowcharts included herein show the architecture, functionality, and/or operation of a possible implementation of software. In this regard, each block can be interpreted to represent a module, segment, or portion of code, which comprises one or more executable instructions for implementing the specified logical function(s). It should also be noted that in some alternative implementations, the functions noted in the blocks may occur out of the order. For example, two blocks shown in succession may in fact be executed substantially concurrently or the blocks may sometimes be executed in the reverse order, depending upon the functionality involved.

[0045] One should note that any of the programs listed herein, which can include an ordered listing of executable instructions for implementing logical functions (such as depicted in the flowcharts), can be embodied in any computer-readable medium for use by or in connection with an instruction execution system, apparatus, or device, such as a computer-based system, processor-containing system, or other system that can fetch the instructions from the instruction execution system, apparatus, or device and execute the instructions. In the context of this document, a “computer-readable medium” can be any means that can contain, store, communicate, propagate, or transport the program for use by or in connection with the instruction execution

system, apparatus, or device. The computer readable medium can be, for example but not limited to, an electronic, magnetic, optical, electromagnetic, infrared, or semiconductor system, apparatus, or device. More specific examples (a nonexhaustive list) of the computer-readable medium could include an electrical connection (electronic) having one or more wires, a portable computer diskette (magnetic), a random access memory (RAM) (electronic), a read-only memory (ROM) (electronic), an erasable programmable read-only memory (EPROM or Flash memory) (electronic), an optical fiber (optical), and a portable compact disc read-only memory (CDROM) (optical). In addition, the scope of the certain embodiments of this disclosure can include embodying the functionality described in logic embodied in hardware or software-configured mediums.

[0046] It should be emphasized that the above-described embodiments are merely possible examples of implementations. Many variations and modifications may be made to the above-described embodiments. All such modifications and variations are intended to be included herein within the scope of this disclosure.

CLAIMS

Therefore, at least the following is claimed:

- 1 1. A method for analyzing communication sessions using fragments comprising:
2 delineating fragments of an audio component of a communication session,
3 each of the fragments being attributable to a party of the communication session and
4 representing a contiguous period of time during which that party was speaking; and
5 automatically assessing quality of at least some of the fragments such that a
6 quality assessment of the communication session is determined.
- 1 2. The method of claim 1, wherein automatically assessing comprises analyzing a
2 sequence of the fragments to determine which party was speaking and for how long.
- 1 3. The method of claim 1, wherein automatically assessing comprises manually
2 assessing the quality of at least some of the fragments and using the quality
3 assessments obtained manually as inputs for automatically assessing quality of other
4 fragments.
- 1 4. The method of claim 1, wherein automatically assessing comprises defining
2 rules and analyzing the fragments for characteristics embodied by the rules.
- 1 5. The method of claim 4, wherein the rules indicate that a quality assessment is
2 to be lowered based on a determination that a party to a communication session is a
3 contact center agent, and that the agent interrupted, by speaking, another party of the
4 communication session that was speaking.

1 6. The method of claim 4, wherein the rules indicate that a quality assessment is
2 to be lowered based on a determination that a party to a communication session is a
3 contact center agent, and that the agent spoke for a duration exceeding a
4 predetermined time limit without another party to the communication session
5 speaking.

1 7. The method of claim 4, wherein the rules indicate that a quality assessment is
2 to be lowered based on a determination that a party to a communication session is a
3 contact center agent, and that the agent spoke at a volume level that at least one of: not
4 less than a high volume threshold and not higher than a low volume threshold.

1 8. The method of claim 1, wherein automatically assessing comprises, with
2 respect to the fragments analyzed, weighting scoring associated with the fragments
3 based, at least in part, on a time that the respective fragments occurred during the
4 communication session.

1 9. The method of claim 1, wherein automatically assessing comprises performing
2 script adherence analysis.

1 10. The method of claim 1, wherein automatically assessing comprises evaluating
2 the communication session for fraud.

1 11. The method of claim 1, wherein at least a portion of the communication
2 session is conducted using Internet Protocol packets.

1 12. The method of claim 1, further comprising recording the communication
2 session.

1 13. The method of claim 1, wherein:
2 one of the parties to the communication session is a contact center agent; and
3 the method further comprises altering a work schedule of the agent based, at
4 least in part, on the quality assessment of the communication session.

1 14. The method of claim 1, further comprising appending information to the
2 fragments.

1 15. A system for analyzing communications using fragments comprising:
2 a communication analyzer operative to:
3 delineate fragments of an audio component of a communication
4 session, each of the fragments being attributable to a party of the
5 communication session and representing a contiguous period of time during
6 which that party was speaking; and
7 automatically assess quality of at least some of the fragments such that a
8 quality assessment of the communication session is determined.

1 16. The system of claim 15, wherein the system comprises a speech recognition
2 engine operative to generate a transcript of at least a portion of the communication
3 session.

1 17. The system of claim 15, wherein the system comprises a phonetic analyzer
2 operative to generate a phoneme sequence of at least a portion of the communication
3 session.



1 18. The system of claim 15, wherein the system comprises a vox detection
2 analyzer operative to provide amplitude information corresponding to volume levels
3 that the audio component exhibited during the communication session, the volume
4 levels being used by the communication analyzer to determine locations at for
5 defining fragments.

1 19. A computer readable medium having a computer program stored thereon for
2 performing the computer executable method of:
3 delineating fragments of an audio component of a communication session,
4 each of the fragments being attributable to a party of the communication session and
5 representing a contiguous period of time during which that party was speaking; and
6 automatically assessing quality of at least some of the fragments such that a
7 quality assessment of the communication session is determined.

1 20. The computer readable medium of claim 19, wherein automatically assessing
2 comprises analyzing a sequence of the fragments to determine which party was
3 speaking and for how long.

1 21. The computer readable medium of claim 19, wherein automatically assessing
2 comprises manually assessing the quality of at least some of the fragments and using
3 the quality assessments obtained manually as inputs for automatically assessing
4 quality of other fragments.

1 22. The computer readable medium of claim 19, wherein automatically assessing
2 comprises defining rules and analyzing the fragments for characteristics embodied by
3 the rules.

**FIG. 1****FIG. 2**





FIG. 3

**FIG. 4**




FIG. 5

100

