

(81) Designated States (*unless otherwise indicated, for every kind of national protection available*): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, IT, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW,

```

graph TD
    EC([Electronic Content  
150]) --> PAM[Profile Attribution Model  
600]
    EC --> ASR[ASR Model  
140]
    ST[Speech Tags  
152] --> PAM
    ST --> ASR
    PAM --> UP([User Profiles  
662])
    ASR --> TO([Transcript (Original)  
154A])
    ASR --> M([Metadata  
156])
    TO --> PPM[Post-Process Model  
142]
    M --> TMD([Transcript (Metadata)  
154C])
    PPM --> TPD([Transcript (Post-Processed)  
154B])
    TMD --> TPD
    TPD --> SM1[Summarization Model  
146]
    TPD --> TM[Template Model  
146]
    SM1 --> SA([Summary (Abstractive)  
158A])
    TM --> TP([Template (Populated)  
160A])
    TP --> SM2[Summarization Model  
146]
    SM2 --> SA
    SA --> SM3[Summarization Model  
146]
    SM3 --> SA
  
```

(S7) Abstract: Attributes of electronic content from a meeting are identified and evaluated to determine whether sub-portions of the electronic content should or should not be attributed to a user profile. Upon determining that the sub-portion should be attributed to a user profile, attributes of the sub-portion of electronic content are compared to attributes of stored user profiles. A probability that the sub-portion corresponds to at least one stored user profile is calculated. Based on the calculated probability, the sub-portion is attributed to a stored user profile or a guest user profile.



SA, SC, SD, SE, SG, SK, SL, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, WS, ZA, ZM, ZW.

- (84) Designated States** (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

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

Published:

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

AUTOMATED MEETING MINUTES GENERATION SERVICE

BACKGROUND

[0001] Meetings are a very common forum where people exchange ideas, make plans, and share information. With the ubiquity of automatic speech recognition systems come vast amounts of meeting transcripts. With the abundance of these automatically generated meeting transcriptions, meeting summarization, and other downstream operations, are of great importance to both the participants of the meetings and other parties interested in the meeting content. Because of the large volume of text included in the meeting transcripts, it is often difficult to sift through the content to find relevant information or get a general idea of the content of the meeting. Therefore, the need to succinctly summarize the content of a meeting naturally arises.

[0002] Several methods of generating summaries for meeting have been created. One non-computerized method includes assigning one participant of the meeting to manually keep meeting minutes notes and prepare a formalized summary from the meeting for use by the rest of the participants. Other methods of summarizing meetings include complex multi-step pipelines that make joint optimization intractable. There are also a handful of existing deep neural models for generating text summarization. However, existing systems for are not configured to handle the semantic structures and styles of meeting transcripts, which are quite different from transcriptions from articles and conversations. Existing systems are also not configured to extract content from meetings to train models for machine transcription and subsequent summarization.

[0003] There are several types of meeting summaries that can be created, including extractive summaries and abstractive summaries. Extractive summarization is the strategy of concatenating extracts taken from a corpus (i.e., meeting transcription) into a summary, while abstractive summarization involves paraphrasing the corpus using novel sentences. Abstractive summaries are often more readable and useful but present more challenges in achieving accurate summarization as compared to extractive summarization techniques.

[0004] For example, current abstractive summarization requires complicated multi-stage machine learning pipelines, such as template generation, sentence clustering, multi-sentence compression, candidate sentence generation, and ranking. As these approaches are not end-to-end optimizable, it is hard to jointly improve various parts in the pipeline to enhance the overall performance. Moreover, some components (e.g., template generation), require extensive human involvement, rendering the solution difficult to scale or transfer

between systems.

[0005] Some systems have been configured to perform document summarization, such as pointer-generator networks, reinforced summarization networks, and memory networks. These deep learning methods can effectively generate abstractive document summaries by directly optimizing defined goals. However, existing systems are not able to perform the meeting summarization task which, in contrast to general document summarization, inherently bears a number of challenges that make it more difficult for end-to-end training than document summarization.

[0006] First, for example, a meeting is carried out between multiple participants, wherein each participant has a different role or contributive function in the meeting, different semantic styles, and other variable factors that affect the meeting transcription and downstream actions in the summarization pipeline. Thus, the different semantic styles, standpoints, and roles of each participant all contribute to the heterogeneous nature of the meeting transcript. Furthermore, it is often difficult to ascertain the identity of the speaker solely based on a meeting transcription and/or audio recording for the same reasons.

[0007] Second, the transcript and summary of a single meeting are usually much longer than those of a document. For instance, in some news articles and/or news transcriptions (i.e., common datasets used to train models on summarization processes), there are on average hundreds of tokens per article. In contrast, some meeting transcriptions contain thousands of tokens per transcript and hundreds of turns (i.e., changes between speakers). This poses a great challenge to the time and space efficiency of meeting summarization methods because as models trained on non-similar datasets (e.g., news content or other non-meeting sources) are used for meeting summarization, the summary output is not optimized for the length nor the structure of the meeting transcript.

[0008] A third challenge is that the structure of the meeting transcript is very distinct from other document formats, as a meeting usually progresses according to an agenda and according to defined goals (or goals that are defined during the concourse of the meeting), while other documents tend to have a general form including an introduction, several main paragraphs, and a conclusion.

[0009] Other challenges to summarization include the problem of error propagation from the initial transcription creation through to other downstream operations. Oftentimes, the speech service (e.g., automatic speech recognition (ASR) system and/or speech-to-text system) will transcribe the audio detected from the speakers (i.e., participants) of a meeting. However, the ASR system often will not correctly transcribe every word from every

speaker, causing the transcription to include certain word errors which can be defined by a word error recognition (WER) percentage.

[0010] In some instances, the speakers themselves introduce errors into the transcription, wherein the manner in which the speakers expressed themselves produced grammatical and/or syntactic errors are included in the transcription, even if the transcription process did not introduce more errors. Both machine and human induced errors can be difficult to find and correct and therefore, can be propagated through to other operations performed on the transcript. For example, in some instances, a summary generated from a meeting transcript could include errors introduced by a speech-to-text translation and/or spoken errors. Such errors greatly decrease the ability of existing systems to generate summaries or other products based on meeting content.

[0011] It is also noted, as alluded to earlier, that existing systems are not well configured to identify participants of a meeting, particularly those that are not previously associated with a known profile, and to attribute content of the meeting to the different participants, especially the unknown participants.

[0012] Accordingly, there is an ongoing need for improved systems and methods for providing meeting transcriptions and for facilitating accurate and readable subsequent downstream operations with the transcribed meeting content, such as meeting summarizations, meeting minutes, outlines, as well as to identify meeting participants and to attribute meeting content to the different meeting participants, including unknown participants.

[0013] The subject matter claimed herein is not limited to embodiments that solve any disadvantages or that operate only in environments such as those described above. Rather, this background is only provided to illustrate one exemplary technology area where some embodiments described herein may be practiced.

BRIEF SUMMARY

[0014] Disclosed embodiments are directed towards embodiments for automatically processing electronic content and for attributing sub-portions of the electronic content to contributing entities.

[0015] Some embodiments include methods and systems for attributing a sub-portion of electronic content associated with a meeting to either a stored user profile or a guest user profile corresponding to a participant of the meeting. In these embodiments, a computing system identifies electronic content associated with a plurality of stored user profiles. The computing system then identifies a sub-portion of electronic content and identifies attributes

corresponding to the sub-portion of electronic content. The computing system accesses a database storing a plurality of stored user profiles including a plurality of known user profiles and a plurality of previously generated guest user profiles.

[0016] After identifying attributes of the sub-portion of electronic content, the computing system evaluates a first attribute of the sub-portion of electronic content to determine that the portion of the electronic content should be attributed to a user profile of the plurality of user profiles or, alternatively, should not be attributed to a user profile of a plurality of user profiles. In response to a determination that the sub-portion of electronic content should be attributed to a user profile of the plurality of user profiles, the computing system compares a second attribute of the sub-portion of electronic content to one or more attributes of the stored user profiles. The computing system then calculates a probability that the sub-portion of electronic content corresponds to one or more of the stored user profiles. Subsequently, based at least in part on the calculated probability, the computing system attributes the sub-portion of electronic content to at least one of the stored user profiles or to a guest user profile.

[0017] Some disclosed embodiments are directed towards embodiments for automatically processing electronic content and for generating corresponding output. In some instances, the disclosed embodiments are also directed towards embodiments for automatically generating synthesized data (i.e., a summary and/or meeting minutes template) from electronic content obtained from a meeting. In some instances, the meeting transcripts include metadata associated user profiles corresponding to participants associated with the meeting transcript.

[0018] This Summary is provided to introduce a selection of concepts in a simplified form that are further described below in the Detailed Description. This Summary is not intended to identify key features or essential features of the claimed subject matter, nor is it intended to be used as an aid in determining the scope of the claimed subject matter.

[0019] Additional features and advantages will be set forth in the description which follows, and in part will be obvious from the description, or may be learned by the practice of the teachings herein. Features and advantages of the invention may be realized and obtained by means of the instruments and combinations particularly pointed out in the appended claims. Features of the present invention will become more fully apparent from the following description and appended claims or may be learned by the practice of the invention as set forth hereinafter.

BRIEF DESCRIPTION OF THE DRAWINGS

[0020] In order to describe the manner in which the above-recited and other advantages and features can be obtained, a more particular description of the subject matter briefly described above will be rendered by reference to specific embodiments which are illustrated
5 in the appended drawings. Understanding that these drawings depict only typical embodiments and are not therefore to be considered to be limiting in scope, embodiments will be described and explained with additional specificity and detail through the use of the accompanying drawings in which:

[0021] Figure 1 illustrates a computing environment in which speech services are in
10 communication with a computing system that is configured to transcribe and post-process electronic content and generate an output based on the post-processed electronic content.

[0022] Figure 2 illustrates one embodiment of a process flow diagram of a computing system configured to include and/or implement aspects of the disclosed embodiments.

[0023] Figure 3 illustrates one embodiment of a flow diagram having a plurality of acts
15 associated with methods for automatically processing electronic content and generating corresponding output.

[0024] Figure 4 illustrates one embodiment of a design overview of a machine learning model configured for abstractive summarization.

[0025] Figure 5 illustrates one embodiment of a user interface displaying an example of
20 a summary generated by a machine learning model configured for abstractive summarization.

[0026] Figure 6 illustrates one embodiment of a process flow diagram of a computing system configured to include and/or implement aspects of the disclosed embodiments, including attributing electronic content to a user profile.

[0027] Figure 7 illustrates one embodiment of a flow diagram having a plurality of acts
25 associated with methods for automatically processing electronic content and attributing sub-portions of the electronic content to one or more contributing entities associated with the electronic content.

[0028] Figure 8 illustrates an example embodiment of a process flow diagram for
30 determining the attribution of electronic content to user profiles.

[0029] Figure 9 illustrates an example embodiment of a plurality of contributing entities associated with a meeting from which electronic content is obtained with a meeting assistant device.

[0030] Figure 10 illustrates an example embodiment of electronic content included in

the user profiles referenced in Figure 1.

DETAILED DESCRIPTION

[0031] Disclosed embodiments are directed towards embodiments for automatically processing electronic content and for attributing sub-portions of the electronic content to contributing entities.

[0032] In some instances, the disclosed embodiments provide technical benefits over conventional systems and methods for automatically attributing sub-portions of electronic content to contributing entities (i.e., contributors to or creators of the electronic content) associated with the electronic content by improving the accuracy and speed of identifying contributing entities by storing, in a known database, a plurality of user profiles corresponding to one or more contributing entities. Furthermore, in some instances, if a contributing entity is not recognized in the known database, the sub-portion will be attributed to a guest user profile, wherein the guest user profile can be later tagged with the correct identity of the contributing entity associated with the sub-portion of electronic content.

[0033] Technical benefits of the disclosed embodiments also include facilitating the automatic generation of improved outputs of downstream operations performed on the electronic content and user profiles, including the creation of transcriptions of the electronic content, summaries generated by abstractive summarization techniques and/or automatically populated templates, for example. In some instances, the transcription beneficially includes metadata obtained from one or more user profiles in the electronic content and/or transcribed audio speech to facilitate the generation of improved meeting outputs from the transcription. In some embodiments, the meeting output is a generated guest user profile or one or more updated stored user profiles.

[0034] Some of the disclosed embodiments are also directed towards embodiments for automatically processing electronic content and for generating corresponding output, wherein electronic content associated with a meeting is identified and audio speech detected in the electronic content is transcribed by an ASR model trained on speech-to-text data. The transcription of the electronic content is then post-processed by modifying text included in the transcription. Additionally, output based on the post-processed transcription is generated, such as, for example, a summary and/or a template for the meeting. Disclosed embodiments are also directed to a plurality of acts associated with computer-implemented methods for automatically processing electronic content and for generating corresponding output.

[0035] In some instances, the disclosed embodiments provide technical benefits over conventional systems and methods for automatically generating a meeting output (e.g., summaries, templates, etc.) by improving the readability and accuracy of a transcription of electronic content associated with a meeting. For example, during post-processing, the original transcription is modified by changing the punctuation, grammar, or formatting introduced by the ASR model during the transcription process. Notably, in some embodiments, the transcription is also modified by changing or omitting one or more words that were included in both the audio speech and the transcription (e.g., errors introduced by one or more of the participants).

[0036] Technical benefits of the disclosed embodiments also include facilitating the automatic generation of improved outputs of downstream operations performed on the transcriptions, including summaries generated by abstractive summarization techniques and/or automatically populated templates, for example. In some instances, the transcription beneficially includes speech tags identified in the electronic content and/or transcribed audio speech to facilitate the generation of improved meeting outputs.

[0037] Additional benefits and functionality of the disclosed embodiments will be described below, including the processing and post-processing of electronic content, as well as the generation of output based on the processed and/or post-processed electronic content.

[0038] Attention will now be directed to Figures 1 and 2, which illustrate corresponding components of a computing system 110 which may include and/or be used to implement aspects of the disclosed invention. As shown, the computing system includes a plurality of machine learning models and data types associated with inputs and outputs of the machine learning models.

[0039] Attention will be first directed to Figure 1, which illustrates the computing system 110 as part of a computing environment 100 that also includes speech service(s) 120A, image service(s) 120B, and video service(s) 120C in communication (via a network 150) with the computing system 110. The computing system 110 is configured to transcribe and post-process electronic content and to generate an output based on the post-processed electronic content, as described throughout this disclosure. The computing system 110 is also configured to attribute electronic content to one or more contributing entities associated with the electronic content.

[0040] The computing system 110, for example, includes one or more processor(s) 112 and a storage 114 (i.e., hardware storage device) storing computer-executable instructions (see executable components 116 and 118), wherein ellipses “...” 170A indicate that the

storage 114 is able to house any number of computer-executable instructions by which the computing system 110 is configured to implement one or more aspects of the disclosed embodiments when the computer-executable instructions are executed by the one or more processor(s) 112.

5 **[0041]** The storage 114 is shown as a single storage unit. However, it will be appreciated that the storage 114 is, in some embodiments, a distributed storage that is distributed to several separate and sometimes remote systems. The system 110 can also comprise a distributed system, in some embodiments, with one or more of the system 110 components being maintained/run by different discrete systems that are remote from each other and that
10 each perform different tasks. In some instances, a plurality of distributed systems performs similar and/or shared tasks for implementing the disclosed functionality, such as in a distributed cloud environment.

[0042] The storage 114 is presently shown as storing a plurality of machine learning models and/or engines. For example, computing system 110 comprises one or more of the
15 following: ASR model 140, face recognition model 141, post-process model 142, speech tag model 144, template model 146, summarization model 148, or profile attribution model 149 which are configured to implement the different functionality described herein.

[0043] In some instances, the storage 114 includes computer-executable instructions for instantiating or executing one or more of the models shown in computing system 110. In
20 some instances, the one or more models are configured as machine learning models or machine learned models. In some instances, the one or more models are configured as deep learning models and/or algorithms. In some instances, the one or more models are configured as engines or processing systems (e.g., computing systems integrated within computing system 110), wherein each engine (i.e., model) comprises one or more processors
25 and corresponding computer-executable instructions.

[0044] For example, in some instances, the ASR model 140 is configured as a speech recognition engine configured to identify electronic content including audible speech, isolate audible speech from other noise and/or data, understand the speech, and apply accurate labeling to the speech data (e.g., generate a transcription matching the spoken
30 words recorded from or being recorded during a meeting).

[0045] Because meetings generally include a plurality of speakers having different tones, semantic styles, and roles, the ASR model 140 is trained to be speaker independent while yielding a low word error recognition (WER) percentage. In some embodiments, the ASR model 140 is trained (i.e., adapted with corresponding training data) to frequently

recorded speakers to improve the accuracy of the transcription. In some instances, this includes adapting/training the model based on content of each individual speaker's audio data, as well as the acoustic features of the audio speech pertaining to one or more particular speakers.

5 **[0046]** In some embodiments, the acoustic features are used to generate and/or update a meeting voice profile for the corresponding speaker, wherein the meeting voice profile can be matched to a saved voice profile included in a user profile corresponding to the speaker. This is beneficial when the computing system is determining to which user profile a certain audio recording should be attributed. For example, if the acoustic features of the audio
10 recording match one or more of the acoustic features of the voice profile of the user profile, the probability that the audio recording corresponds to that user profile is increased.

[0047] In some instances, one or more of the models are integrated and/or combined such that a single model is configured to implement a plurality of aspects of the disclosed embodiments. For example, in some embodiments, the ASR model 140 is also configured
15 to perform speech tagging techniques embodied by the speech tag model 144. In some embodiments, the ASR model 140 is integrated with one or more speech services 120 to detect audio speech from a meeting and transcribe the audio speech. In some embodiments, the speech tag model 144 is integrated with the speech services 120 and/or the ASR model 140, wherein the speech services 120 provide metadata to the ASR model corresponding to
20 highlighted (i.e., tagged) portions of the electronic content 150 and/or transcription (e.g., transcripts 154).

[0048] For example, in some instances, the face recognition model 141 is configured as a face recognition engine configured to identify electronic content including video and/or image data obtained during or from a meeting, isolate frames of the video and/or image data,
25 and accurately associate the data to a particular entity (i.e., a participant of the meeting) included in the electronic content. More specifically, the face recognition model 141 is configured to detect and/or identify one or more human faces from a captured image or video frame.

[0049] In some embodiments, the face recognition model 141 recognizes facial features
30 such as hair color, skin color, eye color, face shape and dimensions, facial contour, height, pre-determined multi-point reference structure, or other feature to facilitate the detection and/or identification of a face of a particular contributing entity or speaker.

[0050] Because meetings generally include a plurality of participants having different styles, appearances, characteristics, and even changing characteristics over time from

meeting to meeting, the face recognition model 141 is trained to be participant independent in some instances.

[0051] In some embodiments, the face recognition model 141 is trained (i.e., adapted with corresponding training data) to frequently recorded participants to improve the accuracy of the user profile attribution. In some instances, this includes adapting/training the model based on content of each individual participants video and/or image data, as well as the visual features pertaining to one or more particular participants.

[0052] In addition to a plurality of machine learning models, algorithms, and engines, the computing system 110, for example, the storage 114 of the computing system 110, is shown housing a plurality of data types utilized by or generated by the machine learning models, including electronic content 150 (e.g., electronic content obtained via the speech service(s) 120, speech tags 152 (e.g., speech tags identified and/or stored by the speech tag model 144), transcripts 154 (e.g., automatically created transcripts by ASR model 140 and/or post-processed transcripts generated by post-process model 142), metadata 156 associated with the electronic content 150 and/or other data types, summaries 158 (e.g., summaries authored by meeting participants and/or summaries generated by the computing system, for example by summarization model 148), and templates 160 (e.g., templates generated and/or automatically populated by template model 146).

[0053] Additionally, in some embodiments, the stored data types may also include user profiles (i.e., known, stored and/or guest user profiles corresponding to contributing entities that are associated with the electronic content).

[0054] In some instances, the electronic content 150 also includes training data for training and/or evaluating one or more of the plurality of models associated with computing system 110.

25 Post-Processing

[0055] In some embodiments, a post-process model 142 is configured to perform downstream processing (i.e., editing, modification) to the meeting transcripts 154. In some embodiments, post-processing includes actions to improve the accuracy of the meeting transcription based on the full context of the electronic content, domain knowledge, human knowledge, and other information. Additionally, post-processing includes action taken to enhance the readability of the meeting transcription.

[0056] One aspect of the post-processing system is to generate more readable, user-friendly sentences given the audio and recognition outputs from the ASR engines. Some metrics to evaluate the readability of text from different dimensions include, for example,

one or more of the following: syntactic correctness, semantic closeness, fluency, or style. In some embodiments, these metrics will be also used to guide the manual labelling for data collection and automatic evaluation of the systems. Readability is also directed to a style transfer from a spoken language (from a speaker's audio stream) to a written language (beneficial for meeting transcripts). Spoken language refers to spontaneous speech that does not naturally follow normal grammatical conventions. Additionally, readability includes proper segmentation, capitalization, fluency, and/or grammar, as well as properly formatted dates, times, names, locations, and/or other numerical entities.

[0057] In some embodiments, post-processing includes modifications based on white lists, blacklists and/or the removal or replacement of hesitation words (e.g. um, uh) and repeated phrases removal, grammar error correction. In some embodiments, a sequence to sequence (seq2seq) system is used to generate more diversified sentences, followed by the ranking system to rank the output in terms of various metrics (e.g. readability, accuracy) and/or other metrics (WER, BLEU, MaxMatch, and/or GLEU).

[0058] In some embodiments, the WER metric is configured as a readability-aware WER (RA-WER) by removing the text normalization before calculating Levenshtein distance. All mismatches due to grammatical mistakes, disfluency, as well as improper formats of capitalization, punctuation, and written numerical entity are treated as errors. If there are alternative references, the closest one to the candidate is selected. Metrics are specific to the speech recognition task, the machine translation task, and grammar error correction task, among other tasks.

[0059] In some embodiments, the post-processing systems and methods are directed to the idea of "understanding then rewriting" as a new NLP task to provide conversion from raw ASR transcripts to error-free and highly readable text. In this solution, error correction and readability improvements will be performed in a single step.

[0060] Traditional ASR post-processing methods offer improvements in readability; however, there are two important shortcomings, as follows. (1) Since the whole process is divided into several sub-processes, the mistakes in the previous steps will accumulate. For example, in the sentence, "Mary had a little lamb. It's fleece was white as snow.", if in the punctuation step, a period '.' is added after the word "had," the rule-based capitalization will capitalize the word 'a.' (2) The traditional methods tend to transcribe the speech verbatim while ignoring the readability of the output text. It cannot detect and correct disfluency in spontaneous speech transcripts. For example, in an utterance such as "I want a flight ticket to Boston, uh, I mean to Denver on Friday", the speaker means to

communicate “I want a flight ticket to Denver on Friday.” The segment “to Boston, uh, I mean” in the transcript is not useful for interpreting the intent of the sentence and hinders human readability and the performance of many downstream tasks. The inclusion of these words in the transcription increases the cognitive load to the reader, whether that reader is an automated machine reader or a human reader.

[0061] According to the disclosed embodiments, the post-processing model (e.g., post-process model 142) is configured to automatically and beneficially modify the original transcription by removing the segment “to Boston, uh, I mean,” thereby increasing the overall readability of the transcription and decreasing the computational load associated with post-processing of the transcript.

[0062] In some embodiments, a baseline setup is used as the post-processing model including the production 2-step post-processing pipe of a speech recognition system, namely n-best language model (LM) rescoring followed by inverse text normalization (ITN). In some embodiments, other models (e.g., MASS, UniLM, RoBERTa, RoBERTa-UniLM and/or other models) are adapted to the post-processing task of increasing readability for a machine and/or human reader and overall accuracy of the transcription.

[0063] Disclosed embodiments are also directed to systems and methods for building a pipeline to evaluate the performance of the post-processing system, including creating datasets and metrics. In some embodiments, the post-processing system is configured as offline and/or online processing. Due to the lack of relevant data, we constructed a dataset for the APR task using a GEC dataset as seed data. The GEC dataset is composed of pairs of grammatically incorrect sentences and corresponding sentences corrected by a human annotator. First, a text-to-speech (TTS) system converts the ungrammatical sentences to speech. Then, an ASR system transcribes the TTS output. Finally, the output of the ASR system and the original grammatical sentences are used to create the data pairs.

[0064] In some embodiments, post-processing also refers to actions taken to enrich the meeting transcription such as automatic summarization, extraction of keywords, topics, entities, as well as enriching the content with externally obtained metadata such as knowledge graphs and organizational people graphs that describe the hierarchical relationship and duties between meeting participants. Additionally, or alternatively, the seq2seq system is used to incorporate external knowledge (e.g. attendee names, project names) and full context with the electronic content, audio stream, and/or meeting transcript. In some instances, the inclusion of external knowledge is performed via bias vectors. In some embodiments, external knowledge includes attendee names, project names, project

deadlines, project locations, meeting locations, organizational structure, and/or other relevant information. In some embodiments, the transcription is embedded with links to content (e.g., meeting documents, emails, directions, user accounts, etc.) related to the meeting but external to the transcription based on the audio speech recognized by the ASR model.

[0065] From the extracted information and/or added external knowledge, the system is able to generate summarization for each individual participant and/or the whole group. Additionally, the system is configured to generate follow-up items, action items, with associated parties and deadlines based on the data enriching post-processing. In some embodiments, summarization across multiple related meetings is also generated.

[0066] Additionally, or alternatively, the post-processing also refers to actions taken to facilitate the retrieval of information from the audio speech and corresponding transcription. For example, in some instances, the meeting transcript is modified to be searchable content, searchable by keywords, attendees, entities, and/or topics. In some instances, the audio speech and/or transcription is configured for conversational question and answer (QnA) functionality. In this manner, a user is provided a functionality for searching across meetings. Conversational QnA will incorporate a search function to narrow down the reader's (machine or human) reader's focus to a few paragraphs in order to increase the speed of finding a relevant answer to the question searched.

[0067] The disclosed embodiments for the post-processing model herein are therefore directed to a novel natural language processing (NLP) task for ASR post-processing to improve readability. This post-processing transforms the noisy ASR output into readable text for humans and downstream tasks, while maintaining the semantic meaning of the speaker. Noise, in this regard, is characterized by syntactic and grammatical errors, disfluency, informal words, background noise incorrectly detected as speech, and other features that make the transcription difficult to understand.

[0068] The disclosed embodiments are also directed to methods to address the lack of task-specific data by synthesizing examples for the post-processing task, including the use of text-to-speech (TTS) data and grammatical error correction (GEC) data.

Speech Tagging

[0069] In some embodiments, the speech tag model 144 is performed to provide a way for a user (i.e., speaker, meeting participant, etc.) and/or a system to highlight important parts of the electronic content obtained from a meeting (e.g., speech tags included with recorded audio data and/or transcription of the audio). In some embodiments, the speech

tags are attributed to a particular speaker (i.e., contributing entity) via the profile attribution model 141 (see Figure 6).

[0070] As shown in Figures 1 and 2, a speech tag model 144 is provided for use in and/or for the computing system 110 and is configured to identify spoken keywords in the audio speech of the electronic content 150 obtained from a meeting. For example, given an utterance wherein a user asks (i.e., gives a speech command) to an AI assistant to take a note (i.e., highlight) what the user will say next, the computing system 110 detects keywords, understands the note type, and identifies the content. Identified keywords trigger the computing system to insert speech tags 152 as metadata 156 into the electronic content 150 and/or transcription (e.g., original transcript 154A of Figure 2).

[0071] For example, an initial or starting keyword (or keyword phrase) is pre-defined and detectable by the speech tag model 144, wherein the speech tag model 144 recognizes a starting index for a new portion of the electronic content and/or associated transcription. In some instances, an ending keyword (or keyword phrase) is also pre-defined wherein the speech tag model 144 recognizes an ending or closing index for the highlighted portion of the electronic content and/or associated transcription.

[0072] For example, a user speaks “Hi Princeton” (i.e., starting keyword phrase) to indicate a start of a speech tag, and speaks “That’s it” (i.e., ending keyword phrase) to indicate an end of the speech tag. A user then can say “Hi Princeton, take a note on a follow up action: upload the design document to Teams. That’s it.” The speech tag model 144 will detect and recognize the starting keyword phrase “Hi Princeton” and recognize the ending keyword phrase “That’s it.” Once the boundaries of the highlighted portion (e.g., speech tag 152) are defined, the note type is understood (e.g., “follow up action”) wherein the speech tag 152 is determined to correspond to a special meaning of the speech content (e.g., a follow up action, a to-do list, open issues, resolved issues, summary points). Then, the content of the speech tag is determined (e.g., “upload the design document to Teams”).

[0073] In some embodiments, the speech tag 152 includes corresponding metadata associated with attributes/features detected for the tagged content, including, for example, an assignment to a responsible party (e.g., a team member responsible for completing a task), assignment to an assignment party (i.e., the person identifying a task and/or person to whom the responsible party should report a status of the task), an assignor, content that is discovered by the system as being associated with the task and that can be linked to the task, a deadline associated with the task, and/or other relevant information.

[0074] In some embodiments, the attributes/features are used by the profile attribution

model 141 to update a user profile 162 with the tagged content. For example, a user profile 162 can include a speech tag 152 that they spoke and/or a user profile 162 can include a speech tag 152 including content that is related to them (e.g., an assignment and/or deadline).

5 [0075] In another example, the user defines the keywords as custom phrases, wherein the starting keyword is “Hi Denmark” and the ending keyword is “Done Denmark” to vocally indicate the start and end of a portion of the transcription (e.g., transcript 154) that the user desires to highlight and/or assign a level of importance. In some instances, the speech tag model 144 is deployed on a speech-enabled device implementing a one or more
10 speech services (e.g., speech service(s) 120).

[0076] In some embodiments, a computing system (e.g., computing system 110) provides a user with a user interface, wherein a user customizes definitions for the start and end parts of the speech tag. Through the user interface, a user records a few examples of the keywords (and/or keyword phrases) for the starting tag and ending tag.

15 [0077] A speech recognition engine (e.g., ASR model 140) is used to decode the examples via a beam search to estimate the labels of the keyword (i.e., possible text transcriptions corresponding to the spoken keywords). The algorithm (i.e., speech recognition engine) maintains an N-best list of label sequence hypotheses, along with a confidence for each hypothesis, in order to minimize the error that will be incurred by
20 incorrectly estimating the labels. The labels include one or more of the following: phonemes, graphemes, or other linguistic subunits. The algorithm loops (i.e., repeats) for both the starting tag and the ending tag of the speech tag to obtain a start tag model and an end tag model for the use of speech tag triggering during inference (i.e., during streaming speech).

25 [0078] At inference time, a voice activity detector (VAD) is used to determine which frames of the electronic content (e.g., electronic content 150) include speech audio. Frames including speech audio are sent to an ASR model (e.g., ASR model 140) and/or a speech tag model (e.g., speech tag model 144). For each segment of speech, a recognition score is calculated, the score indicating the probability of a confidence score of the speech
30 recognition engine output (i.e., how likely the segment includes speech and/or the likelihood the ASR model correctly understood and applied labels corresponding to parts of the speech segment).

[0079] In some embodiments, the ASR model also extracts audio features of the frames that include audio speech to determine the identity of the speaker associated with the audio

speech. For example, the acoustic features of the audio speech can be matched to a known user profile or associated with a guest profile. In such instances, the corresponding audio speech and speech tag will be mapped/attributed to the known user profile and/or guest profile.

5 **[0080]** To calculate and obtain speech tag scores, each hypothesis in the tag models (e.g., a start tag model and/or an end tag model) is scored using a forward algorithm, wherein the hypothesis scores are aggregated to obtain a single detection score for start and end tags, respectively. If a detection score exceeds a pre-determined confidence threshold, the computing system (e.g., computing system 110) is triggered to identify and catalog
10 individual speech tags.

[0081] In some instances, the confidence threshold is manually tuned. Additionally, or alternatively, the confidence threshold is machine-learned. To initiate a triggered state of speech tagging, a final determination of whether a tagging keyword is detected is made by a rule-based and/or machine-learned model based on (1) the relative values of the speech
15 recognition confidence scores and/or speech tag confidence scores, (2) based on a rule that the start tag must occur (i.e., be detected) before the end tag, (3) based on a rule that there is a time-out or time-space between the start and end tags, and/or (4) based on the allowance and/or correction of user input errors such as repeated tags or unclosed tags.

[0082] Once the speech tags (e.g., a starting keyword and an ending keyword) are
20 detected, the transcript or content between the tags is extracted for subsequent processing (e.g., note classification and content span detection). In some embodiments, note classification and content span detection are performed by a deep learning model configured with multi-task learning. For example, the model comprises two task-specific head layers, wherein one layer is an individual sentence or sentence-pair classification head, while the
25 other layer is a token classification head.

[0083] It is noted that a sentence typically comprises a plurality of tokens. For instances in which there is an accessible database (i.e., list) of predefined note classes, the individual sentence classification is used. For cases involving custom note types, a system administrator is able to define a list of note types with corresponding descriptions and
30 examples. In such embodiments, a sentence-pair classification head is used in the neural network to predict to which custom note the transcript corresponds.

[0084] The token classification head is used to predict the start and end position of the note span in the transcript. Pretrained language models, or encoders, are used to provide meaningful generalization ability. Specifically, the task-specific head layers are configured

to be “on top” of the pretrained layers. In some embodiments, the entire neural network is fine-tuned (i.e., trained) on supervised task-specific data, wherein supervised training comprises validated ground truth training data (e.g., human labels). In the case that any of the classification output is invalid (e.g., undetectable note type and/or misclassified note type), the triggering of the speech tagging is identified as a false accept (i.e., false alarm) and the computing system ignores the previously detected speech tag.

[0085] It will be appreciated that, in some instances, speech tagging will be performed during and/or after the meeting from which the electronic content is being obtained in order to highlight important phrases and/or portions of the electronic content. The disclosed embodiments beneficially include speech tagging functionality as a novel mechanism to insert structured semantics into a speech stream (e.g., electronic content 150) and/or into metadata 156 corresponding to a transcript 154, thus facilitating an improvement in downstream processing to easily and quickly separate highlighted content from non-highlighted content.

[0086] In some instances, the highlighted portions will be used to facilitate the generation of meeting minutes (e.g., transcript-based output such as a summary and/or template) after the meeting. In such embodiments, rule-based and/or neural network models are used to analyze the content of the speech tags 152 and the corresponding highlighted portions of the transcription to generate a structured summary 158 and/or automatically populate a template 160. The rule-based and neural network models help to make the inclusion of a speech tag 152 more natural during speech (i.e., naturally embedded as part of typical sentences heard in meetings). The speech tag 152 is also more flexible for use in downstream processing (i.e., post-processing) and simultaneously keep semantics more accurate throughout the downstream processing and corresponding outputs.

[0087] In some embodiments, the system also searches local and/or remote databases for content that is related to the detected speech terms, tasks, dates and/or relevant parties, to link the related content to the tags and/or with the tags as metadata for the tags.

[0088] In some embodiments, the databases include a database of user profiles 162 including information about the user (see also Figure 10). These user profiles 162 may be specific to particular people and may include particular personal attributes associated with those particular people. In some instances, the user profiles 162 may also include one or more generic guest profiles (e.g., a guest profile having particular/predetermined guest attributes).

[0089] In some embodiments, the downstream processing is performed by a

summarization model 148 configured to generate an abstractive summary of the transcription of the meeting. The summarization model 148 is described in further detail in reference to Figures 4 and 5. Additionally, or alternatively, the downstream processing is performed by a template model 146 configured to generate a plurality of templates and/or automatically populate one or more of the generated templates as a meeting output.

Meeting Minutes Templates

[0090] As shown in Figures 1 and 2, computing system 110 also comprises a template model 146 configured to generate and/or automatically populate a template corresponding to a structured outline of input fields associated with a particular source of electronic content (e.g., a particular meeting type and/or corresponding transcript), based on content detected in the transcript and/or based on linked/related content.

[0091] In some embodiments, the template model 146 automatically generates a plurality of templates (e.g., templates 146A, 146B, 146C, 146D), and/or automatically selects a particular template from the plurality of templates, based on a plurality of meeting types and/or user preferences and further based on content identified in the meeting transcript and/or based on the tags.

[0092] In some embodiments, the user preferences are learned over time by a deep learning network associated with the template model 146. Then, from the plurality of templates, one or more templates is automatically selected and automatically populated with information extracted from the post-processed transcription of the meeting's electronic content and/or other information that is linked with the one or more tags.

[0093] In some embodiments, the output generated from the post-processed meeting transcription (e.g., by a template model) are structured meeting notes configured as a meeting minutes template so that users will quickly and easily know what was discussed during the meeting, what was accomplished, and what remains to be done. The template provides a format wherein the outcomes, and group and personal action items are easily accessible by meeting participants and/or third parties interested in the meeting content.

[0094] The template version of the meeting minutes will expose important meeting information via an extensible template (e.g., defined in an XML file). Much of the data will be intelligently auto populated using information from the transcript, meeting audio, and meeting metadata, including metadata extracted from speech tags. In later phases, the Meeting Minutes service will consume additional inputs from the Office graph and other external information. For accessibility, and to improve remote meetings, the chat logs for a given meeting are incorporated as an input.

[0095] In some embodiments, the template includes elements that can be modified by users. Ensuring that edits can flow back to the service is vital for continuous learning. It should be appreciated that the order, presentation, and subset of the template items vary depending on the industry, the organization's admin settings, the meeting administrator's settings, or the type of meeting (classified). Due to the current state of related technologies, in some embodiments, certain categories will be supervised via pre-configured meeting commands, e.g., "Hey Princeton, assign an action item to Will to schedule a meeting with Cindy before next Friday." (See speech tagging above).

[0096] Several examples of a skeleton/template of a meeting minutes are presented below. Depending on user/designer input, certain elements will be editable, purely navigational, and/or only visible to certain users (e.g. certain analytics). Certain items such as the additional media may be solely populated via user uploads. In some embodiments, one or more templates are generated based on a meeting type classification (e.g., working meetings, team/project meetings, one-on-one meetings, staff meetings, management meetings, kickoff meetings, committee meetings, stand-up meetings, townhalls, steering committees, negotiations, performance reviews, and/or status update meetings).

[0097] Working meetings are meetings where people take time to work on something together. These are characteristic of medium or large organizations and span all meeting sizes. Collaboration, disagreement, and tangents are common in these meetings. Information flow tends to be distributed amongst attendees rather than coming from a single source. Common forms of working meetings include brainstorming, problem-solving, decision making, and issue resolution. Therefore, an overview of important information included with meeting minutes beneficially includes a focus on personal action items, team action items, next steps (e.g., next meetings, follow-up discussions), and outcomes (e.g., actions taken during the meetings, decisions, results of votes, etc.). Context for the meeting minutes includes the time and data of the meeting, list of attendees, files shared at or leading up to the meeting. Secondary information is also included in some embodiments such as additional private and/or public notes, whiteboard notes, and/or highlighted sections (i.e., pinned or tagged portions).

[0098] Team/project meetings are similar to working meetings but focus more on a project's "status." A meeting minutes template therefore beneficially focuses on who is working on what assignment and the status of individual and/or collective task. Outcomes are characterized by decisions made and/or issues blocking the completion of the task.

[0099] One-on-one meetings tend to occur between a manager and their employees,

though may also occur between two or more parties. Often these meetings are centered around career development, so the focus of a template may further be directed to deadlines and priority adjustments that should be made and a sense of progress from previous meetings and/or across an individual's career.

- 5 **[00100]** Staff meetings are also similar to working meetings but generally are defined by the transmission of information from one (or a select few) participants to a large volume of participants. In some embodiments, templates generated and populated for staff meetings beneficially include custom executive summaries and agenda items based on salient information for a particular reader (e.g., different departments of a company receive
- 10 different content relating to their specific goals and tasks). In some embodiments, the customized features are based on preferences stored in association with the user profile corresponding to the particular reader, wherein the user profile includes information about their specific goals, tasks, organizational function, role, etc.

- [00101]** Here are a few non-limiting examples of meeting templates that can be
- 15 automatically generated, selected and/or populated, as described:

Example 1: Generic Meeting Minutes Template

I. Meeting Metadata

- a. Date and time
- b. Attendees (*not invitees*)
 - 20 i. (Optional) roles / affiliations of attendees if important for sector
- c. Subject
- d. Meeting ID

II. Executive Summary (abstractive)

- a. Description: A short (3-6 sentence) summary of the high-level content of
 - 25 the meeting. In most cases, this focuses on what the *team* discussed as a whole, without getting lost in the details. In some embodiments, a participant and/or group is associated with the summary, wherein the executive summary is customizable to what is more salient or informative to a particular reader.
 - b. (Extended) – Summary generated for each participant.
- 30

III. Agenda / Verbose Summary – Nested list of items:

- a. Topic
- b. Content: series of key points formatted as participants or the whole group's contributions related to the topic.

- c. Associated Attendees
- d. Transcript IDs associated with the topic.

IV. Follow-up Actions

- a. (readable / abstractive) Action content
- b. Status
- c. Point(s) of contact (responsible) (In some examples, regarding a user interface or post-meeting email includes the primary reader's individual TO-DO list items first).
- d. Due Date
- e. Transcript IDs associated with the action

V. Key Words

- a. Word
- b. Links to Transcript IDs for navigation

VI. Next Steps(s) – (Mainly taken from metadata. Some could be inferred from NLU)

- a. **Next Meetings(s)**
 - i. Date
 - ii. Location
 - iii. Subject
 - iv. Tasks to address at the meeting (optional)
- b. Any Additional information

VII. Additional Notes, Comments, and Files (Perhaps a binding with existing Teams meeting notes content, but any place for people to be taking private and/or shared notes and linking to other content)

- a. Public notes – for additional shared comments
- b. Private notes – per user

Example 2: Expanded Meeting Minutes Template

I. Meeting Metadata

- a. Date and time
- b. Location (probably hidden in most scenarios)
- c. Invitees (probably hidden in most scenarios)
- d. Attendees (*not invitees*)
 - i. (Optional) roles / affiliations of attendees if important for sector
- e. Subject

f. Meeting ID

II. Executive Summary (abstractive)

a. Description: A short (3-6 sentence) summary of the high-level content of the meeting. In most cases, this focuses on what the *team* discussed as a whole, without getting lost in the details.

b. (Extended) – Summary generated for each participant.

III. Agenda / Verbose Summary – Nested list of items:

a. Topic

b. Content: series of key points formatted as participants or the whole group's contributions related to the topic. E.g.:

i. Person A updated the group that s/he completed task X.

ii. The group discussed whether to delay the release.

iii. Person B voiced that the group needs to consider issue Y for this topic.

iv. The group voted on this resolution.

v. Etc.

c. Associated Attendees

d. Associated Entities

e. Transcript IDs associated with the topic.

f. Sub-topics

i. Sub-topic

ii. Sub-topic content

iii. Etc.

IV. Follow-up Actions

a. (readable) Action content

b. Status

c. Assigned by [nullable list]

d. Point(s) of contact (responsible), or assigned to [list]

e. Due Date

f. Transcript IDs associated with the action

g. Linked projects

h. Linked topics

i. Other Links

V. Decisions

- a. Content
- b. Linked agenda items (from above)

VI. Unresolved Items

- a. Content
- b. Linked agenda items (from above)

VII. Status updates/Progress

- a. Linked entities (project or task)
- b. Linked topics (from agenda)
- c. Status content
 - i. Evaluation of progress
- d. Linked transcript IDs
- e. Links to similar status in previous meetings in series or of same type
(visibility may be limited to relevant individuals)

VIII. Key Words

- a. Word
- b. Links to Transcript IDs for navigation

IX. Entities (could be subset of key words - Leverage Entity Linking Intelligence Service)

- a. Term
- b. Content (perhaps like search result cards on Bing)
- c. Links to Transcript IDs for navigation
- d. Links to associated topics
- e. Links to associated documents (URL to SharePoint/emails, etc.)

X. Highlights – (both automatically generated using hot-spot or other detectors and manually pinned)

- a. Attendee(s)
- b. Transcript IDs

XI. Post-processed transcript

- a. Transcript results after being re-processed to correct errors and improve readability

XII. Additional Media

- a. Uploaded images / diagrams
- b. Linked documents/files
- c. Etc.

XIII. Meeting Analytics – extensible

- a. Proportion of meeting taken up by each participant’s speech
- b. List of interrupts and overlapping speech
- c. Number of unresolved issues
- 5 d. Subjective Meeting Success (documented from participants)
- e. Time per topic
- f. Quantity of off-topic items
- g. Personal feedback on contributions (e.g. use of filler words, inclusive language, etc.)
- 10 h. Locations and proportion of meeting associated with “Hot Spots” (were participants are highly engaged or excited)
- i. Etc.

XIV. Next Steps(s) – (Mainly taken from metadata. Some could be inferred from NLU)

- 15 a. **Next Meetings(s)**
 - i. Date
 - ii. Location
 - iii. Subject
 - iv. Tasks to address at the meeting (optional)
- 20 b. Any Additional information

XV. Additional Notes and Comments (Perhaps a binding with existing Teams meeting notes content, but any place for people to be taking private and/or shared notes and linking to other content)

- a. Public notes – for additional shared comments
- 25 b. Private notes – per user

*Example 3: Working Meetings Template***I. Meeting Metadata**

- a. Date and time
- b. Attendees (*not invitees*)
 - 30 i. (Optional) roles / affiliations of attendees if important for sector
- c. Subject
- d. Meeting ID

II. Executive Summary (abstractive)

- a. Description: A short (3-6 sentence) summary of the high-level content of

the meeting. In most cases, this focuses on what the *team* discussed as a whole, without getting lost in the details.

III. Follow-up Actions – Categorized based on assigned to me and those for the whole group.

- 5 a. (readable) Action content
- b. Status
- c. Point(s) of contact (person responsible – for personalized action items lists)
- d. Due Date (could be null)
- e. Transcript IDs associated with the action

IV. Outcomes

- 10 a. **Decisions**
 - i. Content
 - ii. Linked agenda items (from agenda)
- b. **Unresolved Items**
 - 15 i. Content
 - ii. Linked agenda items (from agenda)

V. Highlights – (both automatically generated using hot-spot or other detectors and *manually pinned*)

- 20 a. Attendee(s)
- b. Transcript IDs

VI. Additional Media – (Not auto-populated)

- a. Uploaded images / diagrams
- b. Linked documents
- c. Etc.

VII. Additional Notes and Comments (Not auto populated. A binding for users to take/share notes, follow-up after the meeting, etc.)

- 25 a. Public notes – for additional shared comments
- b. Private notes – per user

VIII. Next Meeting – (Mainly taken from metadata)

- 30 a. Date
- b. Location
- c. Subject
- d. Tasks to address at the meeting (optional)

IX. Agenda / Verbose Summary – Nested list of items:

- a. Topic
 - b. Content: series of key points formatted as participants or the whole group's contributions related to the topic. E.g.:
 - i. Person A updated the group that s/he completed task X.
 - ii. The group discussed whether to delay the release.
 - iii. Person B voiced that the group needs to consider issue Y for this topic.
 - iv. The group voted on this resolution.
 - v. Etc.
 - c. Associated Attendees
 - d. Associated Entities
 - e. Transcript IDs associated with the topic.
 - f. Sub-topics
 - i. Sub-topic
 - ii. Sub-topic content
 - iii. Etc.
- X. Additional Content**
- a. **Key Words**
 - i. Word
 - ii. Links to Transcript IDs for navigation
 - b. **Entities** (could be subset of key words)
 - i. Term
 - ii. Content (perhaps like search result cards on Bing)
 - iii. Links to Transcript IDs for navigation
 - iv. Links to associated topics
 - v. Links to associated documents (URL to SharePoint/emails, etc.)
 - c. **Meeting Analytics** – extensible
 - i. Proportion of meeting taken up by each participant's speech
 - ii. List of interrupts and overlapping speech
 - iii. Number of unresolved issues
 - iv. Subjective Meeting Success (documented from participants)
 - v. Time per topic
 - vi. Quantity of off-topic items
 - vii. Personal feedback on contributions (e.g. use of filler words, inclusive

language, etc.)

viii. Etc.

Example 4: Team/Project Meeting Template

I. Meeting Metadata

- 5
- a. Date and time
 - b. Subject
 - c. Meeting ID

II. Executive Summary (abstractive)

- 10
- a. Description: A short (3-6 sentence) summary of the high-level content of the meeting. In most cases, this focuses on what the *team* discussed as a whole, without getting lost in the details.

III. Current Status – (can be hierarchical)

- 15
- a. Linked entities (project or task)
 - b. Linked topics (from agenda)
 - c. Status content
 - d. Attendees
 - e. Linked transcript IDs

IV. Follow-up Actions (team first, then personal, then others)

- 20
- a. (readable) Action content
 - b. Status
 - c. Point(s) of contact (responsible)
 - d. Due Date
 - e. Transcript IDs associated with the actions

V. Highlights – (both automatically generated using hot-spot or other detectors and manually pinned)

- 25
- a. Attendee(s)
 - b. Transcript IDs

VI. Additional Media – (Not auto-populated)

- 30
- a. Uploaded images / diagrams
 - b. Linked documents

VII. Additional Notes and Comments

- a. Public notes – for additional shared comments
- b. Private notes – per user

VIII. Agenda / Verbose Summary – Nested list of items:

- a. Topic
- b. Content: series of key points formatted as participants or the whole group's contributions related to the topic. E.g.:
 - i. Person A updated the group that s/he completed task X.
 - ii. The group discussed whether to delay the release.
 - iii. Person B voiced that the group needs to consider issue Y for this topic.
 - iv. The group voted on this resolution.
 - v. Etc.
- c. Associated Attendees
- d. Associated Entities
- e. Transcript IDs associated with the topic.
- f. Sub-topics
 - i. Sub-topic
 - ii. Sub-topic content
 - iii. Etc.

IX. Additional Content

- a. **Decisions**
 - i. Content
 - ii. Linked agenda items (from above)
- b. **Unresolved Items**
 - i. Content
 - ii. Linked agenda items (from above)
- c. **Key Words**
 - i. Word
 - ii. Links to Transcript IDs for navigation
- d. **Meeting Analytics** – extensible
 - i. Proportion of meeting taken up by each participant's speech
 - ii. List of interrupts and overlapping speech
 - iii. Number of unresolved issues
 - iv. Subjective Meeting Success (documented from participants)
 - v. Time per topic
 - vi. Quantity of off-topic items
 - vii. Personal feedback on contributions (e.g. use of filler words, inclusive

language, etc.)

viii. Etc.

e. **Entities** (could be subset of key words)

i. Term

5 ii. Content (perhaps like search result cards on Bing)

iii. Links to Transcript IDs for navigation

iv. Links to associated topics

v. Links to associated documents (URL to SharePoint/emails/uploaded files etc.)

10 f. **Next Meeting** – (Mainly taken from metadata)

i. Date

ii. Location

iii. Subject

iv. Tasks to address at the meeting (optional)

15 *Example 5: One-on-one Meeting Template*

I. Meeting Metadata

a. Date and time

b. Subject

c. Meeting ID

20 **II. Executive Summary (abstractive)**

a. Description: A short (3-6 sentence) summary of the high-level content of the meeting. In most cases, this focuses on what the *team* discussed as a whole, without getting lost in the details.

III. Follow-up Actions (personal, then other)

25 a. (readable) Action content

b. Status

c. Point(s) of contact (responsible)

d. Due Date

e. Transcript IDs associated with the actions

30 **IV. Agenda / Verbose Summary** – Nested list of items:

a. Topic

b. Content: series of key points formatted as participants or the whole group's contributions related to the topic. E.g.:

i. Person A updated the group that s/he completed task X.

- ii. We both discussed whether to delay the release.
- iii. Person B voiced that the group needs to consider issue Y for this topic.
- iv. Etc.

- 5 c. Associated Attendees
- d. Associated Entities
- e. Transcript IDs associated with the topic.
- f. Sub-topics

- i. Sub-topic
- 10 ii. Sub-topic content
- iii. Etc...

V. Outcomes

a. Decisions

- i. Content
- 15 ii. Linked agenda items (from above)

b. Unresolved Items

- i. Content
- ii. Linked agenda items (from above)

VI. Current Status/Progress – (can be hierarchal)

- 20 a. Linked entities (project or task)
- b. Linked topics (from agenda)
- c. Status content
 - i. Evaluation of progress
- d. Linked transcript IDs
- 25 e. Links to similar status in previous meetings in series or of same type
(Visibility can be restricted)

VII. Entities

- a. Term
- b. Content (perhaps like search result cards on Bing)
- 30 c. Links to Transcript IDs for navigation
- d. Links to associated topics
- e. Links to associated **documents** (URL to SharePoint/emails/uploaded files etc.)
- f. Links to associated objects in previous meetings

VIII. Highlights – (both automatically generated using hot-spot or other detectors and manually pinned)

- a. Attendee(s)
- b. Transcript IDs

5 **IX. Additional Media – (Not auto-populated)**

- a. Uploaded images / diagrams
- b. Linked documents

X. Additional Notes and Comments

- a. Public notes – for additional shared comments
- b. Private notes – per user

10

XI. Additional Content

a. **Key Words**

- i. Word
- ii. Links to Transcript IDs for navigation

15

b. **Meeting Analytics** – extensible

- i. Proportion of meeting taken up by each participant's speech
- ii. List of interrupts and overlapping speech
- iii. Number of unresolved issues
- iv. Subjective Meeting Success (documented from participants)
- v. Time per topic
- vi. Quantity of off-topic items
- vii. Personal feedback on contributions (e.g. use of filler words, inclusive language, etc.)
- viii. Etc.

20

c. **Next Meeting** – (Mainly taken from metadata)

- i. Date
- ii. Location
- iii. Subject
- iv. Tasks to address at the meeting (optional)

25

30

Profile Attribution

[00102] As shown in Figures 1 and 6, the computing system 110 also includes a profile attribution model 149 and/or profile attribution model 600 configured to attribute electronic content obtained, for example, from a meeting to one or more meeting participants by storing the electronic content in association with one or more user profiles corresponding to

the one or more meeting participants (i.e., contributing entities).

[00103] In some embodiments, with reference to Figure 1, the profile attribution model 149 receives audio and visual data obtained via one or more of the speech service(s) 120A, image service(s) 120B, or video service(s) 120C. The video service(s) 120C are configured to record video and/or obtain a video stream from an event, for example, a meeting having one or more participants (i.e., contributing entities). The video service(s) 120C also include recording equipment such as video recorders and/or cameras having video recording capabilities.

[00104] The image service(s) 120B are configured to capture and record images of an even, such as a meeting, and also include equipment such as cameras. In some instances, the image services(s) 120B record images directly from the environment of the meeting. Additionally, or alternatively, the image service(s) capture image data (e.g., a frame) from video data recorded by the video service(s) 120C. The computing system 110 obtains data from one or more of these services via the network 130 and is able to store the data in the storage 114 via the profile attribution model 149 and/or other models that are able to access and process the data (i.e., electronic content).

[00105] The profile attribution model 149 is also associated with the data type of user profiles 162 which are stored in storage 114 in some instances. The user profiles 162 are configured to as a database corresponding to a particular meeting participant, wherein the user profile stores data that is associated with or linked to the meeting participant.

[00106] For example, referring to Figure 10, the user profile 162 includes a voice profile 1002 (e.g., a voice profile comprised of acoustic features extracted from audio speech associated with the meeting participant), a face profile 1004 (e.g., a face profile comprised of visual features extracted from video and/or image data associated with the meeting participant), a role and/or position data 1006 (e.g., a role vector 460 or organizational data), one or more tags 1008 (e.g., tag for linking the user profile to a meeting participant and/or speech tags), links 1010 (e.g., links to external data, recordings, documents, etc.) electronic content 1012 (e.g., attributed sub-portions of electronic content), identifiers 1014 (e.g., formal names, user names, numerical identifiers, , and/or action items 1016 (i.e., action items that were either assigned by or assigned to the meeting participant)).

[00107] In some embodiments, the voice profile 1002 is generated and/or updated, at least in part, by the ASR model 140 and/or data processed by the ASR model 140, wherein the voice profile 1002 is an aggregation of audio speech having similar acoustic features. In some instances, the voice profile 1002 includes labeled characteristics such as native

language, spoken language, accent, dialect, etc. In some embodiments, the face profile 1004 is generated and/or updated, at least in part, by the face recognition model 141, wherein the face profile 1004 comprises an aggregation of image and/or video data having similar visual features as associated with a face and/or body of a person (e.g., contributing entity and/or meeting participant). The face profile 1004, in some instances, includes labeled characteristics such as eye color, hair color, face shape, prominent features, height, etc.

[00108] In some embodiments, the user profiles 162 include a plurality of stored user profiles (e.g., previously generated and/or updated user profiles) and/or a plurality of guest profiles. In some instances, the guest profiles are newly generated guest profiles that are generated in response to processing meeting content. Additionally, or alternatively, the guest profiles are previously generated and stored guest profiles and/or the guest profiles are stored profile templates (e.g., unpopulated, non-updated guest profiles).

[00109] In some embodiments, the guest user profile becomes a known user profile when the guest user profile is tagged with the identity of a meeting participant, thereby linking the user profile and its associated electronic content to the meeting participant. Thus, the electronic content associated with the user profile is attributed to that particular meeting participant. In some embodiments, the guest user profile remains untagged, for example, when the identity of the contributing entity associated with the electronic content is unknown and/or unidentifiable. These user profiles are accessible by the profile attribution model 149. In some embodiments, the profile attribution model 149 is configured to generate new user profiles (e.g., a guest user profile) and/or update existing user profiles (e.g., a tagged user profile and/or previously stored guest user profile).

[00110] In some embodiments, the profile attribution model 149 utilizes the electronic content 150 (audio and visual data) to identify sub-portions to attribute to a user profile. Additionally, or alternatively, the profile attribution model 149 utilizes only audio data and/or only visual data to attribute sub-portions to a user profile. In some embodiments, each of a plurality of sub-portions of the electronic data (e.g., electronic data 150) are attributed to at least one user profile. Alternatively, one or more of the plurality of sub-portions of the electronic data are attributed to one or more user profile. In some embodiments, a sub-portion is attributed to one user profile, while in some embodiments, a sub-portion is attributed to one or more user profiles. In some embodiments, one or more sub-portions are not attributed to any user profiles. In some embodiments, the profile attribution model 149 utilizes data extracted from a transcription of the electronic content (e.g., transcripts 154).

[00111] The profile attribution model 149 is configured to recognize one or more contributing entities (e.g., participants in a meeting) based on visual and/or acoustic features of each contributing entity. For example, sub-portions of electronic content obtained from a particular meeting are aggregated based on similar and/or matching visual and/or acoustic features. One or more matching sub-portions of the electronic content are then attributed to a particular user profile.

[00112] In some embodiments, the profile attribution model 149 is configured to attribute, in real time, a sub-portion of electronic content to a user profile as soon as the sub-portion is detected and/or recorded (i.e., profile attribution during a meeting). In some embodiments, the profile attribution 149 is configured to attribute sub-portions of electronic content after a meeting.

[00113] During a meeting and/or a meeting transcription process, the profile attribution model 149 is configured to automatically detect one or more speakers (e.g., contributing entities) and distinguish the speakers from each other with one or more user profiles. If a speaker is not able to be matched with a previously stored user profile, the speaker is linked to a guest profile. The speakers are linked to user profiles based on audio and/or video signals. The computing system 110 is also configured to allow users to tag the user profiles, especially guest profiles) by providing identifiers like a username, an email, and other identifying information that helps to identify the speaker to which the user profile has been linked and/or any electronic content attributed to the user profile. The computing system 110 is also configured to allow users to optionally store the user profiles (tagged and/or guest user profiles) for future usage, wherein the user profiles become stored user profiles.

[00114] The profile attribution model 149 is configured to identify electronic content (i.e., data retrieved from one or more of the service(s) 120). Once the electronic content is identified, a sub-portion of the electronic content is selected. Based on one or more characteristics of the sub-portion of the electronic content, the profile attribution model 149 determines whether or not the sub-portion of the electronic content should be attributed to a user profile. The attribution process is described in further detail in reference to Figures 6-10 below.

[00115] In some embodiments, the profile attribution model 149 is configured to receive data output from an ASR model (in some examples, ASR model 140) such as acoustic features extracted by an acoustic model of the ASR model and/or labels applied via a language model of the ASR model. Additionally, or alternatively, the profile attribution model 141 is configured to receive data output (e.g., visual features extracted from an image

or video) from a face recognition model (in some examples, face recognition model 141). In some embodiments, the aforementioned ASR model is incorporated into the speech service(s) 120A and/or the aforementioned face recognition model is incorporated into the video and/or image services 120B, 120C.

5 **[00116]** Attention will now be directed to Figure 2 which illustrates one embodiment of a process flow diagram, and possible corresponding process flow paths, of a computing system (such as computing system 110) that is configured to include and/or implement aspects of the disclosed embodiments.

10 **[00117]** As shown, electronic content 150 is obtained from a meeting (e.g., audio speech from a recorded/ recording audio stream, meeting documents, emails, memos, etc.). The electronic content 150, including audio speech, is sent to a speech tag model 144, wherein the speech tag model generates a plurality of speech tags 152 which highlight certain portions of the electronic content for use in transcription generation (e.g., input to the ASR model 140), post-processing (e.g., input to post-process model 142), template generation and population (e.g., input to the template model 146) and/or summarization (e.g., input to the summarization model 148). These speech tags 152 facilitate the inclusion of specific user-identified portions of the transcript and aid the model(s) in identifying salient information for post-processing, metadata enriching, summarization and/or for meeting minutes templates. In some embodiments, the note type identified in the speech tag 152
20 corresponds to input fields of one or more templates of the template model 146.

25 **[00118]** The electronic content 150 is also sent to the ASR model (in parallel or in series with the speech tag model 144), wherein the ASR model recognizes speech and applies speech labels to generate an original transcription (e.g., transcript 154A). In some embodiments, the ASR model 140 also extracts metadata 156 from the electronic content 150, the metadata 156 being associated with the transcript 154A. The original transcript 154A is then sent to the post-process model 142 to under post-processing as described herein (e.g., error correction and/or improvements to readability).

30 **[00119]** Additionally, or alternatively, the original transcript 154A is sent to the template model 146 having access to a plurality of templates (e.g., templates 146A, 146B, 146C, 146D) to provide a basis from which to generate a new template, modify an existing template, and/or populated a template (e.g., populated template 160A). In some embodiments, the original transcript 154A is sent to the summarization model 148 to generate an abstractive summary 158A.

35 **[00120]** Once the original transcript 154A is post-processed, and the post-process model

142 generates a post-processed transcript 154B, the post-processed transcript 154B is input to the summarization model 148, wherein an abstractive summary 158A is generated based from the post-processed transcript 154B. Additionally, or alternatively, the post-processed transcript 154B is input to the template model 146. In some embodiments, the post-processed transcript 154B is further enriched with metadata 156 to create a Metadata-Enriched Transcript 154C. This transcript 154C is then input to the summarization model 148 and/or the template model 146.

[00121] In some embodiments, the input fields (i.e., framework) for the generated templates are extracted from the original transcript 154A, the post-processed transcript 154B, and/or the summary 158A. In some embodiments, the content that is used for auto-population of a template is extracted from similar sources.

Methods

[00122] The following discussion now refers to a number of methods and method acts that may be performed. Although the method acts may be discussed in a certain order or illustrated in a flow chart as occurring in a particular order, no particular ordering is required unless specifically stated, or required because an act is dependent on another act being completed prior to the act being performed.

[00123] Attention will now be directed to Figure 3 which illustrates a flow diagram 300 that includes various acts associated with exemplary methods that can be implemented by computing systems, such as computing system 130 described above in reference to Figure 1. As shown in Figure 3, the flow diagram 300 includes a plurality of acts (act 310, act 320, act 330, and act 340) which are associated with various methods for automatically processing electronic content and generating output based on the processed (and/or post-processed) electronic content. Examples of claimed components claimed in the acts are described in reference to features of a computing system, for example, computing system 110 of Figure 1 and/or associated process flow diagram paths as described in reference to Figure 2.

[00124] As shown in Figure 3, the flow diagram 300 and corresponding methods include an act of a computing system (e.g., computing system 110) identifying electronic content (e.g., electronic content 150) associated with a meeting, the electronic content including audio speech (act 310). After identifying electronic content, the computing system creates a transcription (e.g., transcript 154A) of audio speech detected in the electronic content with an ASR model (e.g., ASR model 140) trained on speech-to-text training data, wherein the transcription is a text-based transcription (act 320).

[00125] The computing system then performs post-processing on the original transcription to generate a post-processed transcription (e.g., transcript 154B) by using a machine learning model (e.g., post-process model 142) trained on post-processing training data for modifying text included in the original transcription (act 330). In some
5 embodiments, the post-processing includes both (1) modifying at least one of a punctuation, grammar, or formatting of the transcription that was introduced by the ASR model and (2) changing or omitting one or more words in the transcription which were included in both the audio speech and the transcription.

[00126] Once the transcription is post-processed, the computing system generates output
10 based from the post-processed transcription, the output comprising at least one of a meeting summary (e.g., summaries 158 and/or summary 158A) generated by a machine learning summarization model (e.g., summarization model 148) that summarizes content of the post-processed transcription or a template (e.g., templates 160, generated templates 146A, 146B, 146C, 146D, and/or populated template 160A), wherein the template is generated (and/or
15 populated) at least in part from the post-processed transcription (act 340).

Abstractive Summarization

[00127] Typically, the input to the summarization generation engine (e.g., summarization model 148) is a meeting transcript consisting of multiple turns from multiple speakers, coming from the ASR (e.g., ASR model 140) results. Each turn corresponds to a speaker
20 and associated utterance. There are several challenges to generating summaries for transcriptions, especially for transcriptions of meetings. First, because the meeting is carried about between multiple participants, the transcript includes utterances from multiple speakers wherein each speaker is characterized by different semantic styles, standpoints, contributing roles, and even different acoustic features cause the ASR results to be variable
25 between speakers. The heterogeneous nature of the transcript makes it difficult for a singularly trained model to adapt to so many dependent variables.

[00128] In some embodiments, the input to the summarization model 148 also includes data extracted from one or more user profiles 162 generated by the profile attribution model 149, wherein each turn is associated to a user profile 162.

[00129] Conventional summarization models are also typically trained and well adapted
30 for document summarization. However, meeting transcriptions are, on average, are between 5 and 10 times longer than standard documents. This poses a great challenge to the time and space efficiency of summarization models. Additionally, the structure and flow of content in a meeting transcript is often highly variable as compared to written documents.

[00130] Disclosed embodiments are provided herein to overcome the aforementioned challenges and present improved solutions to meeting transcription summarization. First, as meeting transcripts are usually lengthy (e.g., over 4,000 tokens), a direct application of a canonical transformer (i.e., conventional models) is not feasible. For instance, conducting a multi-head self-attention mechanism on a transcript with thousands of tokens is very time consuming and will cause memory overflow problems for the computing system. In contrast, leveraging a hierarchal structure facilitates a reduction in the computational burden. Since a meeting consists of utterances from different participants, it naturally forms a multi-turn hierarchy. Thus, a hierarchal structure of the summarization model follows the multi-turn hierarchy of the meeting transcription.

[00131] Attention will now be directed to Figure 4 which illustrates one example of a neural network 400 configured as an abstractive summarization model (e.g., summarization model 148). Thus, the summarization model carries out both a token-level understanding within each turn (e.g., word-level transformer 410) and a turn-level understanding (e.g., turn-level transformer 420) across the whole of the meeting transcription. During summary generation, the neural network 400 (e.g., Hierarchal Meeting summarization Network 'HMNet') applies attention (e.g., self-attention 414, 424) to both levels of understanding to ensure that each party of the summary stems from different portions of the transcript with varying granularity.

[00132] The overall training process for the neural network is end-to-end, optimizing the cross entropy of the generated summary. Consequently, the neural network 400 makes it very convenient to jointly fine-tune each component to enhance summarization performance. Each transformer in the network 400 is used to incorporate contextual information. For example, a transformer block consists of a multi-head attention layer (e.g., self-attention layers 414, 424) and a feed-forward layer (e.g., feed forward layers 418, 428). These layers are each followed by a layer-norm with residuals (e.g., add & norm layers 416A, 416B, 426A, 426B).

[00133] The attention is based on an inner product of query and key:

$$Attention(Q, K, V) = softmax(\frac{QK^T}{\sqrt{d}})V$$

[00134] Multi-head attention employs the attention h times, each time projecting Q, K, V to a d-dim space:

$$MultiHead(Q, K, V) = Concat(A_1, A_2, \dots, A_h)W^O$$

$$where A_i = \mathbf{Attention}(QW_i^Q, KW_i^K, VW_i^V)$$

[00135] The feed-forward network utilizes two linear transformations:

$$FFN(x) = \max(0, xW_1 + b_1) W_2 + b_2$$

[00136] In addition, as attention mechanism juxtaposes words from all positions together, the position information is lost. To solve this problem, transformer has positional encoding
 5 470 (e.g., positional encoding 470A, 470B, 470C for the word level transformer, turn-level transformer, and decoder, respectively).

[00137] PE is added to input vectors:

$$PE_{(i,2j)} = \sin(i/10000^{\frac{2j}{d}})$$

$$PE_{(i,2j+1)} = \cos(i/10000^{\frac{2j}{d}})$$

10 Where $PE_{(i,j)}$ stands for the j-th dimension of positional embedding for the i-th word in input. Thus, the transformer block on a sequence of n tokens can generate n output embeddings of the same dimension as input. Thus, multiple transformer blocks can be sequentially stacked to form a transformer network, denoted by:

$$Transformer(\{x_1, \dots, x_n\}) = \{y_1, \dots, y_n\}$$

15 [00138] In some embodiments, the problem of meeting summarization is formalized as follows. The input consists of meeting transcripts X and meeting participants P. Suppose there are 's' meetings in total. The transcripts are $X = \{X_1, \dots, X_s\}$. Each meeting transcript consists of multiple turns, where each turn is the utterance of a participant. Thus, $X_i = \{(p_1, u_1), (p_2, u_2), \dots, (p_{L_i}, u_{L_i})\}$, where $p_j \in P, 1 \leq j \leq L_i$, is a participant, and $u_j = (w_1, \dots, w_{l_j})$
 20 is the tokenized utterance from p_j . The human-labelled summary for meeting X_i , denoted by Y_i , is also a sequence of tokens. So, the goal of the system is to generate meeting summary $Y = (y_1, \dots, y_n)$ given the transcripts $X = \{(p_1, u_1), (p_2, u_2), \dots, (p_m, u_m)\}$. In other words, the goal is to maximize the conditional probability of the meeting summary Y given transcript X and network parameters.

25 [00139] The word-level transformer 410 processes the token sequence (e.g., 412A, 412B, 412C) of one turn in the meeting. Each token is encoded in one turn using a trainable embedding matrix D. Thus, the j-th token in the i-th turn, $w_{i,j}$, is associated with a uniform length vector $D(w_{i,j}) = g_{i,j}$. To incorporate syntactic and semantic information, two embedding matrices are trained to represent the part-of-speech (POS) and entity (ENT) tags.
 30 Therefore, the token $w_{i,j}$ is represented by the vector $x_{i,j} = [g_{i,j}; POS_{i,j}; ENT_{i,j}]$. Note a special token $w_{i,0}=[BOS]$ is added before the sequence to represent the beginning of a turn (e.g., <u> 412A). Then, the output of the word-level transformer is denoted as follows:

$$Word_Transformer(\{x_{i,0}, \dots, x_{i,L_i}\}) = \{x_{i,0}^w, \dots, x_{i,L_i}^w\}$$

[00140] The turn-level transformer 420 processes the information of all m turns in a meeting. To represent the i -th turn, the output embedding of the special token [BOS] is employed from the word-level transformer, i.e. $x_{i,0}^w$. Furthermore, it is concatenated with
 5 the role vector of the speaker for this turn, p_i . It follows that the output of the turn-level transformer is:

$$Turn_Transformer(\{[x_{1,0}^w; p_1], \dots, [x_{m,0}^w; p_m]\}) = \{x_1^t, \dots, x_m^t\}$$

[00141] Furthermore, to accommodate the multi-speaker scenario, the neural network 400 incorporates the role of each speaker to encode different semantic styles and standpoints
 10 between participants. For example, a project manager usually emphasizes the progress of the project, whereas a user interface designer tends to focus on user experience and design features. Thus, the neural network 400 employs one or more role vectors 460 (e.g., role vector 460A, role vector 460B) for each meeting participant to represent the speaker's information during encoding. In some embodiments, roles are pre-defined (e.g., program
 15 manager, industrial designer, marketing expert, etc.). Additionally, or alternatively, the role of the speaker is inferred and defined by the model during analysis of the transcription. For each role (of any number of roles), a role vector 460 is trained to represent the role of the speaker as a fixed-length vector. This role vector 460 is appended to the embedding of the speaker's turn (e.g., 422A, 422B).

[00142] The vectorized representation of speaker roles plays an important part in boosting the performance of summarization. The performance can be further maximized if richer data is available in practice. For example, in some instances, if an organization chart of participants is available, the chart is added in representations of the relationship between participants, e.g., manager and developers, into the network. If there is a pool of registered
 25 participants, each participant is assigned a personal vector which acts as a user portrait and evolves as more data about this user is collected.

[00143] The decoder 430 is a transformer to generate the summary tokens. The input to the decoder transformer contains the $k - 1$ previously generated summary tokens $\hat{y}^1, \dots, \hat{y}^{k-1}$. Each token (e.g., 432A, 432B, 432C) is represented by a vector using the same
 30 embedding matrix D as the encoder, $D(\hat{y}^i) = g_i$. The decoder transformer 430 is different from its counterpart in the encoder in two ways. First, as the tokens are generated from left to right, the input to the self-attention layer 434 has to be masked with a lower triangular matrix 480 to avoid peeking at future information.

[00144] Second, the decoder transformer block includes two additional cross-attention layers (e.g., layers 440A, 440B). After self-attention 434, the embeddings first attend with token-level outputs $\{xW\}_m$, L_i , and then with turn- i, j $i=1, j=1$ level outputs $\{xT_i\}_{m_i=1}$, each followed by layer-norm (e.g., layers 436A, 436B) with residuals. This makes the model
 5 attend to different parts of the inputs with varying scales at each inference step.

[00145] The output of the decoder transformer is denoted as:

$$Decoder_Transformer(\{g_1, \dots, g_{k-1}\}) = \{x_1^d, \dots, x_{k-1}^d\}$$

To predict the next token y_k , the network reuses the weight of embedding matrix D to decode v_{k-1} into a probability distribution over the vocabulary:

$$10 \quad P(w_k | w_{<k}) = softmax(x_{k-1}^d D^T)$$

For example, a linear layer 450 and SoftMax layer 452 are applied as final layers to the output of the decoder 430.

[00146] During the training of the model, techniques are employed to minimize the cross entropy. In some embodiments, a teacher-forcing is used during decoder training (i.e., the
 15 decoder takes the ground-truth summary tokens as input). During inference, a beam search is used to select the best candidate. The search starts with the special token <BEGIN> (e.g., token 432A). A common trigram blocking is also used during beam search, if a candidate work creates a trigram that already exists in the previously generated sequence of the beam, wherein the model forcibly sets the word's probability to zero. Finally, the summary with
 20 the highest average log-likelihood per token is selected as the final output of the summarization model.

[00147] The disclosed embodiments are also directed to methods for training a summarization model and generation of training data to be used in training the model. For example, the amount of annotation for meeting summarization is usually inadequate. To
 25 solve the data sparsity issue, several data augmentation approaches are used to enlarge the size of the training data. A first approach splits a meeting summary by speakers. In the labeling phase, the annotator writes a summary for each speaker in the meeting. During the training phase, the model takes only the turns from the specific speaker as the input and the corresponding summary as the output.

30 [00148] In some embodiments, the meeting summary is split by user profiles, wherein some user profiles are tagged with an identity corresponding to a known speaker and some user profiles are guest profiles wherein the speaker identity is unknown.

[00149] In the second approach, a pseudo meeting summarization annotation is created

from large news summarization corpus. For each news article, each paragraph is treated as a turn in the meeting and a placeholder role is assigned to all turns. The summary written by the editor is treated as the ground truth. These approaches have been shown to significantly improve the summarization quality of the summarization model where the amount of meeting annotation is limited.

[00150] The output of the neural network 400 (also see summarization model 148) is innovative by using words that are not directly extracted from the transcript to generate the summary of the transcript. The level of abstractness of the summary is measured via the ratio of novel words or phrases in the summary to directly extracted words. A higher ration, then, in some instances, indicates a more abstractive summary. Thus, the system tends to summarize meetings a high level (e.g., topics and/or decisions) without presenting large amounts of details regarding the high-level concepts.

[00151] Attention will now be directed to Figure 5, which illustrates one example of a user interface 510 displaying a portion of a meeting transcription 512 and corresponding portion of a summary 514 generated by a summarization model (e.g., summarization model 148 and/or neural network 400). As shown in Figure 5, the meeting transcript 512 includes a plurality of turns (e.g., 162 turns) corresponding to a plurality of speakers associated with a plurality of roles (e.g., marketing expert ‘ME’, program manager ‘PM’, industrial designer ‘ID’). The summary 514 includes 21 abstractive sentences, including keywords identified in the meeting transcript 512 (e.g., LCD screen, corporate image, plastic). For example, “Then, a surprisingly great deal of people w indicated that an LCD screen in the remote control would be preferred.” Is summarized by “He discussed the user requirements to include a remote, and discussed how to include in an LCD screen.”

[00152] In some embodiments, wherein the transcription includes errors introduced by the ASR model, post-processing of the transcript is performed to identify and correct those errors. For example, a post-processed version of the meeting transcript 512 omits the character “w” as it appears to be an extraneous text (i.e., mis-labeled audio speech). In some embodiments, the transcription is highly accurate in recognizing and labeling speaker words. However, in some instances, the speaker may introduce errors in their speech that decrease the readability of the transcript. For example, as shown in the user interface 510, program manager (PM) said, “We have to be very *attent* in putting the corporate image in our product.”

[00153] The post-process model (e.g., post-process model 142) is also configured to identify ungrammatical sentences, phrases, and/or words created by the user and/or by the

ASR model that persist in the transcription and either omit and/or replace the label with a more grammatical version. For example, a post-processed version of the PM's transcription is "We have to be very *intent* in putting the corporate image in our product" or "We have to be very *attentive* in putting the corporate image in our product." In these examples, "attent" which is not a canonically recognized word is replaced with "intent" and/or "attentive" which are relevant to the speaker's original meaning and improves the grammatical and semantic syntax of the sentence. These examples are also representative of modifications made during post-processing of the original transcript (e.g., transcript 154A of Figure 2) to a post-processed transcript (e.g., transcript 154B of Figure 2).

[00154] In some embodiments, the post-processing includes adding words and/or phrases to make the transcription more readable and grammatically correct. For example, a post-processed version of the transcription corresponding to the industrial designer (ID) is "... we should use default materials, *such as* simple plastics." The inclusion of the phrase "such as" improves the readability (i.e., a reader's understanding) of the following words "simple plastics" which the speaker intended to name as examples of "default materials." In this manner, the system actually modifies the content of the original audio speech that persists in the transcript by adding text to the audio speech/transcript to make it more clear and easier to summarize.

[00155] Additionally, post-processing modifications described above and/or any of the other modifications described herein (e.g., replacing/deleting/modifying terms) in the transcript that comprise errors and/or unclear terms/phrases that are either introduced by an ASR model and/or introduced by the speaker's original verbal communication, greatly improves any downstream processing of the transcription, such as, for example, in summarizing the transcription.

[00156] In this manner, trained models such as the neural network 400 which often follow some grammatical rule-based data sets, are better able to analyze and synthesize information from the transcription into a summary (e.g., summary 514). Furthermore, errors in the transcriptions are prevented from being propagated into the downstream outputs (e.g., summaries and/or templates).

[00157] In some embodiments, the resulting summary 514 results from post-processing to remove and fix any errors introduced from the transcription model and/or by the summarization model. For example, the extraneous comma in the first sentence is removed and the word "in" is removed so that the summary reads, "He discussed user requirements to include a remote and discussed how to include an LCD screen."

[00158] In some embodiments, the summary 514 is generated and/or post-processed to read, “He discussed user requirements to include an LCD screen in the remote.” In some embodiments, the summary 514 is enriched with further descriptions, such as “*The marketing expert* discussed user requirements...” and/or “*The program manager indicated that* the corporate image must be used.” In some instances, the last illustrated sentence of the summary 514 is post-processed and/or enriched to read, “The case *for the remote* will have to use plastic” wherein “*for the remote*” helps the reader understand more context for the case and it’s requirement to be plastic.

[00159] Summaries (e.g., summaries 158, abstractive summary 158A, summary 514) are able to be evaluated based on several criteria or metrics, wherein a post-processing is applied to the transcription and/or the summary output until a certain level or threshold of each criteria is reached. In some embodiments, the summaries are evaluated for readability and relevance. Readability is a measurement of how fluent the summary language is, including word and grammatical error rate and ratio of coherent to incoherent phrases. Relevance measures how well the summary sums up the main ideas of the meeting. Ultimately, the fewer errors that exist in the original transcription, the more readable and relevant the summary will be.

[00160] In some embodiments, the length of the summary is determined based on a length of the original and/or post-processed transcription of the meeting. In some embodiments, the length of the summary is determined based on summary training data ground truth. In some embodiments, one or more summaries is generated based on an individual reader’s preference and/or a collective group’s preference. In some embodiments, a plurality of summaries is generated wherein each summary is adapted to a particular reader as identified in a hierarchal organizational structure of relevant parties.

[00161] In some embodiments, these preferences and/or a placement associated with the hierarchal organizational structure of the relevant parties is stored as part of a user profile corresponding to the particular reader.

[00162] In some embodiments, certain role vectors are weighted more heavily than other role vectors, wherein the summarization is based on a great volume of content from one role than another role. In some embodiments, summaries are generated across a plurality of meetings, wherein the meetings are chronological and/or related. In some embodiments, summaries are generated based on a focus to a specific theme, topic, task, and/or project, wherein the summaries include information from one or more meeting transcriptions discussing the particular theme, topic, task, and/or project.

[00163] Furthermore, in addition to summarization, other applications such as automatic subtitle generation and meeting minutes generation require automatic speech transcription that is highly readable for humans, while machine translation, dialogue systems, voice search, voice question answering, and many other applications also require highly readable transcriptions to generate the best machine response. The existence of defects in the speech transcriptions will significantly harm the experience of the application for users if the system cannot mitigate the effects of the defects.

[00164] Attention will now be directed to Figure 6 which illustrates an example embodiment of a process flow diagram 600, and possible corresponding process flow paths, of a computing system (such as computing system 110) that is configured to include and/or implement aspects of the disclosed embodiments. It should be appreciated that process flow diagram 600 is related to the process flow diagram 200 shown in Figure 2, wherein the flow diagram 600 of Figure 6 further includes a profile attribution model 600 and user profiles 662 integrated into the process flow paths of the flow diagram 200 of Figure 2.

[00165] In some embodiments, the profile attribution model 600 obtains electronic content 150 for example, as associated with a meeting. Sub-portions of the electronic content 150 are attributed to one or more user profiles 662 which are also generated and updated by the profile attribution model 600.

[00166] In some embodiments, the profile attribution model 600 is configured to process data from the speech tag model 144, including the use of speech tags 152 generated by the speech tag model 144. Additionally, or alternatively, the speech tag model 144 is configured to process data from the profile attribution model 600, including accessing data associated with one or more user profiles. In some embodiments, the speech tags 152 are included in the user profiles 662.

[00167] In some embodiments, the sub-portion of electronic content used by the profile attribution model includes an identified speech tag 152, wherein the attribution of the sub-portion of the electronic content to a particular user profile facilitates an improvement in accuracy of linking one or more speech tags 152 to a particular entity (the entity being associated with the user profile 662 via an identifying tag). For example, data, for example, an action item, deadline, assigning party, and/or responsible party, associated with speech tag 152 is included in the data stored and/or linked to the user profile 662.

[00168] Thus, in some embodiments, the computing system uses a machine learning speech tagging model to generate the speech tags, wherein the machine learning speech tagging model generates at least one tag in response to identifying a spoken starting keyword

and a spoken ending keyword detected in the sub-portion of electronic content. For example, the generating of the at least one tag includes links the at the least one tag to the user profile to which the sub-portion of electronic content is attributed.

5 [00169] In some embodiments, metadata 156 output by the ASR model 140 and/or data extracted from the metadata enriched transcript 154C is processed by the profile attribution model 600 to facilitate an improvement in accurately attributing a sub-portion to a particular user profile and/or accurately tagging a particular user profile to the corresponding contributing entity.

10 [00170] In some embodiments, the user profiles 662 and corresponding sub-portions from electronic content 150 attributed to the user profiles 662 are utilized by the ASR model 140, summarization model 148 and/or template model 146. In some embodiments, data extracted from one or more user profiles 662 are linked to a transcript 154A and/or included in the transcript 154A generated by the ASR model 140. For example, if a sub-portion of the electronic content 150 is previously attributed to a particular user profile, the ASR model
15 140 can automatically assign the labeled portion of audio speech included in the sub-portion to a particular speaker.

[00171] For example, in some instances, the computing system creates a transcription of audio speech included in the electronic content with an automatic speech recognition (ASR) model, wherein the transcription is divided based on one or more identified sub-
20 portions of electronic content. In such instances, each sub-portion of electronic content is labeled with an identifier associated with a user profile to which the sub-portion of electronic content is attributed.

[00172] In some embodiments, wherein the transcription (e.g., transcript 154B, and/or 154C) includes a plurality of links that point to external data related to the electronic content.
25 The external data related to the electronic content is external to the electronic content and corresponds to a particular sub-portion of electronic content that is stored in the user profile to which the particular sub-portion of electronic content is attributed.

[00173] In some embodiments, the transcription includes text-based descriptions of visual performance data of a meeting participant extracted from the sub-portion of electronic
30 content attributed to the user profile corresponding the meeting participant. It is anticipated that the visual performance data will be collected by image and/or video services (e.g., image service(s) 120B and/or video service(s) 120C of Figure 1).

[00174] In some embodiments, the computing system generates output based from the transcription, the output comprising at least one of a meeting summary (e.g., summary

158A) generated by a machine learning summarization model (e.g., summarization model 148) that summarizes content of the transcription or a template (e.g., template 160A) that is generated at least in part from the transcription.

[00175] In some embodiments, the output comprises the meeting summary which is automatically generated based on abstractive summarization of the transcription by a summarization model. In such instances, the summarization model 148 is configured to break the transcription into a plurality of turns comprising a plurality of words, the turns corresponding to a plurality of sub-portions of electronic content. Then, the summarization model 148 analyzes each turn based on a role vector of a speaker corresponding to the turn, wherein the role vector that is applied is based on the user profile to which each sub-portion of electronic content corresponding to each turn is attributed. The summarization model 148 then selectively applies one or more different rules during the analyzing, each of the one or more different rules being based on one or more correspondingly different roles from which the role vector is determined.

[00176] In some embodiments, if a particular sub-portion of electronic content is included in the summary 158A, the summary is also able to be linked to the same user profile 662. In some instances, where the summary 158A includes one or more identifiers that identify the contributing entity associated with a part of the summary, the identity and/or identifier for the contributing entity is generated based on data extracted from the user profile linked to that particular contributing entity.

[00177] In some embodiments, a meeting output comprises a template which is a meeting template that is automatically selected from a plurality of different templates. The template is selected based on a meeting type that is determined from analyzing the transcription and which is automatically populated with content from the transcription, including data corresponding to one or more user profiles to which one or more sub-portions of the transcriptions are attributed.

[00178] Thus, in some embodiments, for example, where a populated template 160A is output by the template model 146, the populated template is generated and/or populated based on data extracted from one or more user profiles 662. For example, based on preferences included with the user profile 662, the template is customized to the user profile 662, thereby customizing the template 160A to the contributing entity tagged with the user profile 662. In some embodiments, wherein the user profile is not tagged (i.e., a guest user profile), the template is still able to be customized to the user profile without having to know the explicit identity of the contributing entity (e.g., meeting participant).

[00179] Attention will now be directed to Figure 7 which illustrates a flow diagram 700 that includes various acts associated with exemplary methods that can be implemented by computing systems, such as computing system 130 described above in reference to Figure 1. As shown in Figure 7, the flow diagram 700 includes a plurality of acts (act 710, act 720, act 730, act 740, act 750, act 760, act 770, and act 780) which are associated with various methods for automatically processing electronic content and attributing sub-portions of the electronic content to contributing entities. Examples of claimed components claimed in the acts are described in reference to features of a computing system, for example, computing system 110 of Figure 1 and/or associated process flow diagram paths as described in reference to Figure 6.

[00180] As shown in Figure 7, the flow diagram 700 and corresponding methods include an act of a computing system (e.g., computing system 110) identifying electronic content (e.g., electronic content 150) associated with a plurality of stored user profiles (act 710). After identifying electronic content, the computing system identifies a sub-portion of the electronic content (act 730). Once the sub-portion is identified, the computing system identifies attributes of the sub-portion (act 740) and evaluates at least a first attribute of the sub-portion of electronic content to determine that the sub-portion of electronic content should be attributed to a user profile (e.g., user profiles 162, 662) of a plurality of user profiles (act 750). The computing system also accesses a database storing a plurality of stored user profiles including a plurality of known user profiles and a plurality of previously generated guest user profiles (act 720).

[00181] In response to a determination that the sub-portion of electronic content should be attributed to the user profile of the plurality of user profiles, the computing system compares a second attribute of the sub-portion of electronic content to one or more attributes of the stored user profiles (act 760). Thereafter, the computing system calculates a probability that the sub-portion of electronic content corresponds to one or more of the stored user profiles (act 770). Based at least in part on the calculated probability, attribute the sub-portion of electronic content to at least one of the stored user profiles or to a guest user profile (act 780).

[00182] In some embodiments, once the sub-portion(s) of the electronic content are attributed and/or the user profile are tagged, the computing system generates output based at least in part from the user profiles and attributed content, the output comprising at least one of a meeting summary (e.g., summaries 158 and/or summary 158A) generated by a machine learning summarization model (e.g., summarization model 148) that summarizes

content of the post-processed transcription and one or more user profiles or a template (e.g., templates 160, generated templates 146A, 146B, 146C, 146D, and/or populated template 160A), wherein the template is generated (and/or populated) at least in part from the post-processed transcription (act 340) and one or more user profiles.

5 **[00183]** Attention will now be directed to Figure 8, an example embodiment of a process flow diagram illustrating possible paths for attributing electronic content to a user profile and tagging the user profile. For example, in some embodiments, one or more contributing entities (i.e., speakers) 802 are contributors to electronic content 804 (e.g., electronic content 150 of Figure 1) recorded and/or aggregated by, for example, speech service(s) 120A, image
10 service(s) 120B, video service(s) 120C, ASR model 140, and/or face recognition model 141. In some embodiments, the electronic content 804 is associated with audio and visual data collected from a meeting in which the contributing entities 802 are meeting participants. A sub-portion of electronic content is identified (e.g., identify portion 806) and selected for determining whether or not the sub-portion should be attributed to a user profile (e.g.,
15 attribute? 810).

[00184] In some embodiments, the determination that a sub-portion of electronic content should or should not be attributed to a user profile is based, at least in part, on an evaluation of one or more identified attributes corresponding to the sub-portion. For example, in some
20 embodiments, an attribute of the sub-portion of electronic content corresponds to a predetermined set of relevance criteria of the sub-portion of electronic content relative to an entire context of the electronic content, wherein the determination to attribute the sub-portion of the electronic content to at least one user profile is at least in part based on a value of the level of relevance or a value of the level of importance exceeding a value of a pre-determined attribution level threshold.

25 **[00185]** In some embodiments, the predetermined set of relevance criteria includes one or more of the following: a time duration, a word count, one or more keywords pre-determined to indicate a level of importance or level of relevance, or one or more unrecognizable words or intervals of the sub-portion of electronic content.

[00186] In some embodiments, the level of relevance and/or level of importance relates
30 to the usefulness of including the content of the sub-portion in downstream operations such as transcription, post-processing, summarization, and/or template generation. Additionally, or alternatively, the level of relevance and/or level of importance relates to the usefulness of including the content of the sub-portion in a user profile, for example, to facilitate an increased accuracy of attributing future electronic content to the correct contributing entity.

[00187] In some embodiments, an upper and/or lower threshold for the time duration and/or word count of a sub-portion is pre-determined, wherein the specified range is known to correspond to a higher likelihood that the sub-portion is relevant and/or important and should be attributed to at least one user profile. For example, in some instances, a sub-portion that is not long enough, such as less than 2 seconds and/or less than 3 words, will not include highly relevant or important information that would be beneficial to attribute to a user profile. In some instances, the short time duration corresponds to the sub-portion containing background noise and/or unintended noise generated from a speaker. In some embodiments, the profile attribution model 141 learns the best upper and/or lower thresholds to achieve an optimal attribution of relevant and/or important electronic content.

[00188] In some embodiments, the one or more keywords includes words used as “filler” words such as “uhm”, “yeah”, “so”, “uh”, etc. which do not contribute to an increased understanding of the semantic meaning of the audio speech and therefore would not be relevant and/or important to include in a user profile and/or other downstream operations.

In some instances, the sub-portion includes one or more words that are unrecognizable, and therefore, not able to be labeled by an ASR model (e.g., ASR model 140). In some embodiments, the lack of detected labels indicates a level of relevance and/or level of importance that does not exceed a profile attribution threshold. For example, in some embodiments, it is not useful to attribute unrecognizable speech to a user profile because it could not be used in downstream processing such as speech transcription, summarization, and/or template generation.

[00189] If the system determines that the sub-portion should not be attributed to a user profile, the computing system ignores that sub-portion of electronic content (e.g., ignore portion of content 814) and identifies a next sub-portion of electronic content (e.g., identify next portion 816), wherein the next sub-portion is evaluated to determine if it should be or should not be attributed to a user profile (e.g., attribute 810).

[00190] In some embodiments, the determination to attribute a sub-portion to at least one user profile (e.g., a stored user profile and/or a new guest user profile) is based on applying a labeling to audio speech included the identified sub-portion of the electronic content with an automatic speech recognition (ASR) model, calculating a confidence score associated with a probability of accuracy of the labeling characterizing the audio speech, comparing the confidence score to a predetermined confidence score threshold, and determining to attribute the sub-portion of electronic content to at least one of a stored or guest user profile when the confidence score meets or exceeds the predetermined confidence score threshold

or to, alternatively, refrain from attributing the sub-portion of electronic content to the at least one of stored or guest user profile when the confidence score fails to meet or exceed the predetermined confidence score threshold.

[00191] If the system determines that the sub-portion should be attributed to a user profile, one or more attributes (e.g., a second attribute) of the sub-portion are compared to one or more attributes of one or more stored user profiles (e.g., compare to stored profiles 812) to determine if one or more attributes of the sub-portion match one or more attributes of a stored user profile (e.g., match? 820). In some embodiments, the second attribute of the sub-portion of electronic content corresponds to one or more attributes of a stored user profile that is associated with a contributing entity.

[00192] In some embodiments, the second attribute of the sub-portion of electronic content includes a voice profile compiled from acoustic data extracted from the sub-portion of electronic. If the voice profile matches a user voice profile (e.g., voice profile 1002 of Figure 10) that is part of the stored user profile, the probability that the sub-portion should be attributed to the stored user profile is increased. In some embodiments, the second attribute of the sub-portion of electronic content includes a face profile compiled from image data extracted from a sub-portion of visual content included in the electronic content. If the face profile matches a user face profile that is part of the stored user profile, then the probability that the sub-portion should be attributed to the stored user profile is increased.

[00193] Additionally, or alternatively, attributes of the sub-portion could be determined to match to attributes of a stored user profile based on words known to be spoken at a certain frequency by a contributing entity 802 or acoustic features such as tone, pitch, rate of talking, accent, etc. Furthermore, visual attributes could match, such as gestures, facial expressions, etc.

[00194] Based on the comparison of one or more sub-portion attributes to one or more attributes of a stored user profile, a probability that the sub-portion corresponds to the stored user profile is calculated. Based on the calculated probability, the sub-portion is attributed to one of the stored user profiles that corresponds to a highest probability of correctly matching. For example, if attributes match, the sub-portion is attributed to a known user profile (e.g., attribute to known profile 822) included in the plurality of stored user profiles or is attributed to a previously generated and stored guest user profile (e.g., attribute to stored guest profile 824) included in the plurality of stored user profiles.

[00195] If the sub-portion is attributed to a known user profile, the computing system, in some instances, prompts a user to verify (e.g., verify 850) that the known user profile is the

correct user profile to which the sub-portion should be attributed. If the verification is positive, the sub-portion is attributed to the matched known user profile (e.g., attribute to matched profile 852), wherein the profile data associated with the match known user profile is updated and/or the sub-portion is stored (i.e., linked to) the known user profile (e.g., store/update profile data 854).

[00196] If the verification is negative, in some embodiments, the computing system automatically reassigns the sub-portion to the next most likely stored user profile (i.e., stored user profile having the next highest probability) (e.g., reassign? 860). If the stored user profile is a known user profile, the sub-portion is attributed to the new matched known profile (e.g., attribute to match profile 862), wherein the profile data is stored and/or updated (e.g., store/update profile data 864). If the attribution of the sub-portion is reassigned to a stored user profile that is a guest profile, the computing system prompts a user to tag (e.g., tag 830) the guest profile to link the guest profile to the identity of a particular contributing entity.

[00197] In some embodiments, the computing system prompts a user to determine if the sub-portion should be reassigned (e.g., reassign 860). Additionally, or alternatively, the reassignment step is bypassed, either automatically or manually, wherein the sub-portion is attributed to the matched profile and the profile data is stored and/or updated.

[00198] In some embodiments, if the attributes of the sub-portion match to attributes of a stored guest profile, the sub-portion is attributed to the previously generated guest user profile included in the plurality of stored user profiles (e.g., attribute to stored guest profile 824). In some embodiments, the attribution is based at least in part on determining whether a calculated probability meets or exceeds a user profile confidence threshold, and then attributing the sub-portion of electronic content to the at least one of the stored user profiles when the calculated probability meets or exceeds the user profile confidence threshold, or alternatively, attribute the sub-portion of electronic content to the guest user profile when the calculated probability meets or exceeds the user profile confidence threshold.

[00199] In some instances, the attributes of sub-portion do not match attributes of a stored user profile (i.e., the calculated probability that the sub-portion corresponds to a known user profile or a previously generated guest profile does not exceed a pre-determined threshold, such as a user profile confidence threshold). In such instances, a new guest profile is generated (e.g., generate guest profile 822). In some embodiments, the generate guest profile 822 also is directed to selecting an unpopulated guest profile included in the stored user profiles, the unpopulated guest profile including the framework and/or allocated memory

storage for the guest profile but has not been updated with any electronic content or identifiers linking to a contributing entity.

[00200] If the sub-portion of electronic is to be attributed to a guest profile (either vis 822 or 824), the computing system prompts a user to tag (e.g., tag? 830) the guest user profile with the identity of the contributing entity to which the guest user profile should correspond (i.e., the contributing entity responsible for contributing the sub-portion to the electronic content). In some embodiments, the computing system automatically tags the user profile. Additionally, or alternatively, the computing system generates a request for the guest user profile to be tagged to correspond to a contributing entity associated with the sub-portion of electronic content for linking the contributing entity to the guest user profile. For example, in response to user input that identifies the contributing entity, tag the guest user profile to correspond with the identified contributing entity, thereby linking the contributing entity to the guest user profile (e.g., tag profile 832). The profile data associated with the newly attributed sub-portion of electronic content and/or tag is stored and/or updated (e.g., store/update profile data 834).

[00201] In some instances, wherein the computing system generates a request for the guest user profile to be tagged to correspond to a contributing entity associated with the sub-portion of electronic content and in response to user input denying the request, the computing system stores the guest user profile in the database of stored user profiles (i.e., the newly generated guest user profile becomes now a previously generated guest user profile) (e.g., store/update profile data 842) or discards the guest user profile (e.g., discard 844).

[00202] In some embodiments, the updating of profile data (e.g., store/update profile data 854, 864, 834, 842) relates to updating the stored user profile based on identified attributes of the sub-portion of electronic content. For example, the identified attributes including one or more of the following: acoustic features extracted from a sub-portion of audio speech included in the sub-portion of electronic content or visual features extracted from a sub-portion of visual content included in the electronic content. In some embodiments, the updating of profile data (e.g., store/update profile data 854, 864, 834, 842) relates to storing the identified attributes corresponding to the sub-portion of electronic content as cached data associated with one or more stored user profiles and refraining from updating the one or more stored user profiles with the cached data. For example, the cached storage of attributes is used in determining which user profile of a plurality of user profiles to which a sub-portion will be attributed, but the attributes are not included in the assigned user profile.

In some instances, the cached storage is cached for a pre-designated amount of time. In some instances, the cached storage is discarded after the sub-portion is attributed to a user profile.

[00203] Attention will now be directed to Figure 9, which illustrates an environment in which a meeting assistant device (i.e., services 120 of Figure 1) is utilized to record electronic content associated with a meeting having a plurality of meeting participants (i.e., speakers). For example, Figure 9 shows a meeting assistant device 910 including at least a microphone 912 configured to record audio data and a camera 914 having image and/or video recording capabilities which is configured to record visual data. A plurality of meeting participants (i.e., speakers and/or contributing entities) (e.g., participant 930A, 930B, 930C, 930D) are shown surrounding the meeting assistant device, wherein each participant is positioned in a specific direction in relation to the meeting assistant device 910. In some instances, the direction also relates to a distance from the device 910 to the meeting participant. For example, meeting participant 930A is positioned in direction 920A, meeting participant 930B is positioned in direction 920B, and so on.

[00204] The microphone 912 is configured to record audio speech (e.g., audio speech 940A, 940B, 940C, 940D) from one or more of the meeting participants. It is anticipated that the audio speech produced by each of the meeting participants is able to be recorded simultaneously and/or in series. For example, in some meetings participant 930A speaks, then meeting participant 930B speaks. In such an embodiment, the recorded audio captured by the microphone includes audio speech 940A followed by audio speech 940B.

[00205] In some instances, participants 930C and 930D are speaking at the same time as wherein audio 940C and audio 940D are overlapping in the recording captured by the microphone. Thus, in such instances, an ASR model (e.g., ASR model 140 of Figure 1) is configured to identify each audio stream/ audio recording and separate them into at least two recordings, one corresponding to audio speech 940C and one recording corresponding to audio speech 940D. The audio speech captured by the microphone 912 is then able to be used in downstream data processing such as speech-to-text transcription, transcript post-processing, summarization, template generation and/ population, speech tagging, and/or user profile attribution.

[00206] It should also be appreciated that the microphone 912 is, in some instances, configured as a stand-alone speech service(s) device, such as a speech-enabled smart device. From the audio speech recorded, a voice profile (a meeting voice profile) is able to be generated for each meeting participant, wherein that voice profile is then associated with or linked to the user's base profile to supplement the user profile and/or to update a

corresponding voice profile for that user profile and which enables subsequent linking of new content. For instance, subsequent audio speech from a particular meeting participant, either later in the same meeting or in a different meeting, can be attributed to the correct meeting participant based on the saved/updated user profile of that participant based on the previously linked/updated voice profile. With regard to the foregoing, the user profile can be a guest profile.

[00207] In some embodiments, the camera 914 is configured as a plurality of cameras configured to record each of the meeting participants. It will also be appreciated that in some embodiments, the camera 914 only records a meeting participant if they are speaking and/or moving and/or speaking and moving at a same time.

[00208] The system uses the video and/ or image data captured by the camera 914 to create a face profile (a meeting face profile) for each of the meeting participants. The meeting assistant 910, via the camera 914, is also able to detect one or more faces (e.g., faces 950A, 950B, 950C, 950D) of each of the meeting participants that have saved user profiles. In particular a detected face (e.g., meeting face profile) can be matched to a saved face profile from one or more user profiles (or a guest profile) and can be used to supplement and/or update the saved user profiles (specifically the face profile portion of the user profile). Then, that updated user profile and/or the meeting face profile corresponding to the meeting participant can be used for subsequently linking new meeting content to the corresponding user/participant, similar to the voice profile matching described previously.

[00209] In some embodiments, the meeting assistant 910 is able to identify the meeting participant based on audio and/or visual features. For example, in some instances, the detected face 950A correspond to a known profile (i.e., a tagged user profile included in the plurality of stored user profiles). In some embodiments, the meeting assistant 910 detects changes or discrepancies between the detected face and the face profile corresponding to one or more meeting participants. These changes include longer or shorter hair, longer or shorter beard, change of hair color, aging, bruises, etc. In some instances, the detected changes are used to update the user profile, while in other instances, the changes are used in calculating the user profile confidence score (i.e., matching probability) for attribution but are not used to update the user profile.

[00210] In view of the foregoing, it will be appreciated that the disclosed embodiments provide many technical benefits over conventional systems and methods for automatically generating a meeting output (e.g., summaries, templates, user profiles, etc.), by improving the readability and accuracy of a transcription of electronic content associated with a

meeting by performing post-processing that corrects punctuation, grammar and formatting introduced by the ASR model during the transcription process, as well as by replacing, modifying, omitting and/or adding to one or more words that were included in both the audio speech and the transcription (e.g., errors introduced by one or more of the participants).

5 **[00211]** In view of the foregoing, it will be appreciated that the technical benefits of the disclosed embodiments include In some instances, the disclosed embodiments provide technical benefits over conventional systems and methods for automatically attributing sub-portions of electronic content to contributing entities (i.e., contributors to or creators of the electronic content) associated with the electronic content by improving the accuracy and
10 speed of identifying contributing entities by storing, in a known database, a plurality of user profiles corresponding to one or more contributing entities. Furthermore, in some instances, if a contributing entity is not recognized in the known database, the sub-portion will be attributed to a guest user profile, wherein the guest user profile can be later tagged with the correct identity of the contributing entity associated with the sub-portion of electronic
15 content.

[00212] Technical benefits of the disclosed embodiments also include facilitating the automatic generation of improved outputs of downstream operations performed on the electronic content and user profiles, including the creation of transcriptions of the electronic content, summaries generated by abstractive summarization techniques and/or
20 automatically populated templates, for example. In some instances, the transcription beneficially includes metadata obtained from one or more user profiles in the electronic content and/or transcribed audio speech to facilitate the generation of improved meeting outputs from the transcription. In some embodiments, the meeting output is a generated guest user profile or one or more updated stored user profiles.

25 **[00213]** Technical benefits of the disclosed embodiments also include facilitating the automatic generation of improved outputs of downstream operations performed on the transcriptions, including summaries generated by abstractive summarization techniques and/or automatically populated templates, for example. In some instances, the transcriptions also beneficially include speech tags identified in the electronic content and/or transcribed
30 audio speech to facilitate the generation of improved meeting outputs.

[00214] Embodiments of the present invention may comprise or utilize a special purpose or general-purpose computer (e.g., computing system 110) including computer hardware, as discussed in greater detail below. Embodiments within the scope of the present invention also include physical and other computer-readable media for carrying or storing computer-

executable instructions and/or data structures. Such computer-readable media can be any available media that can be accessed by a general purpose or special purpose computer system. Computer-readable media (e.g., storage 114 of Figure 1) that store computer-executable instructions (e.g., components 116, 118 of Figure 1) are physical storage media.

5 Computer-readable media that carry computer-executable instructions are transmission media. Thus, by way of example, and not limitation, embodiments of the invention can comprise at least two distinctly different kinds of computer-readable media: physical computer-readable storage media and transmission computer-readable media.

[00215] Physical computer-readable storage media includes RAM, ROM, EEPROM,
10 CD-ROM or other optical disk storage (such as CDs, DVDs, etc.), magnetic disk storage or other magnetic storage devices, or any other medium which can be used to store desired program code means in the form of computer-executable instructions or data structures and which can be accessed by a general purpose or special purpose computer.

[00216] A “network” (e.g., network 130 of Figure 1) is defined as one or more data links
15 that enable the transport of electronic data between computer systems and/or modules and/or other electronic devices. When information is transferred or provided over a network or another communications connection (either hardwired, wireless, or a combination of hardwired or wireless) to a computer, the computer properly views the connection as a transmission medium. Transmissions media can include a network and/or data links which
20 can be used to carry, or desired program code means in the form of computer-executable instructions or data structures and which can be accessed by a general purpose or special purpose computer. Combinations of the above are also included within the scope of computer-readable media.

[00217] Further, upon reaching various computer system components, program code
25 means in the form of computer-executable instructions or data structures can be transferred automatically from transmission computer-readable media to physical computer-readable storage media (or vice versa). For example, computer-executable instructions or data structures received over a network or data link can be buffered in RAM within a network interface module (e.g., a “NIC”), and then eventually transferred to computer system RAM
30 and/or to less volatile computer-readable physical storage media at a computer system. Thus, computer-readable physical storage media can be included in computer system components that also (or even primarily) utilize transmission media.

[00218] Computer-executable instructions comprise, for example, instructions and data which cause a general-purpose computer, special purpose computer, or special purpose

processing device to perform a certain function or group of functions. The computer-executable instructions may be, for example, binaries, intermediate format instructions such as assembly language, or even source code. Although the subject matter has been described in language specific to structural features and/or methodological acts, it is to be understood that the subject matter defined in the appended claims is not necessarily limited to the described features or acts described above. Rather, the described features and acts are disclosed as example forms of implementing the claims.

[00219] Those skilled in the art will appreciate that the invention may be practiced in network computing environments with many types of computer system configurations, including, personal computers, desktop computers, laptop computers, message processors, hand-held devices, multi-processor systems, microprocessor-based or programmable consumer electronics, network PCs, minicomputers, mainframe computers, mobile telephones, PDAs, pagers, routers, switches, and the like. The invention may also be practiced in distributed system environments where local and remote computer systems, which are linked (either by hardwired data links, wireless data links, or by a combination of hardwired and wireless data links) through a network, both perform tasks. In a distributed system environment, program modules may be located in both local and remote memory storage devices.

[00220] Alternatively, or in addition, the functionality described herein can be performed, at least in part, by one or more hardware logic components. For example, and without limitation, illustrative types of hardware logic components that can be used include Field-programmable Gate Arrays (FPGAs), Program-specific Integrated Circuits (ASICs), Program-specific Standard Products (ASSPs), System-on-a-chip systems (SOCs), Complex Programmable Logic Devices (CPLDs), etc.

[00221] The present invention may be embodied in other specific forms without departing from its spirit or characteristics. The described embodiments are to be considered in all respects only as illustrative and not restrictive. The scope of the invention is, therefore, indicated by the appended claims rather than by the foregoing description. All changes which come within the meaning and range of equivalency of the claims are to be embraced within their scope.

CLAIMS

1. A computing system for automatically processing electronic content and for attributing sub-portions of the electronic content to contributing entities, the computing system comprising:

one or more processors; and

one or more computer readable hardware storage devices having stored computer-executable instructions that are executable by the one or more processors to cause the computing system to at least:

identify electronic content associated with a plurality of stored user profiles;

identify a sub-portion of the electronic content;

identify attributes corresponding to the sub-portion of electronic content;

evaluate a first attribute of the sub-portion of electronic content to determine that the sub-portion of electronic content should be attributed to a user profile of a plurality of user profiles or, alternatively, should not be attributed to the user profile of the plurality of user profiles;

access a database storing a plurality of stored user profiles including a plurality of known user profiles and a plurality of previously generated guest user profiles; and

in response to a determination that the sub-portion of electronic content should be attributed to the user profile of the plurality of user profiles, perform the following:

compare a second attribute of the sub-portion of electronic content to one or more attributes of the stored user profiles;

calculate a probability that the sub-portion of electronic content corresponds to one or more of the stored user profiles; and

based at least in part on the calculated probability, attribute the sub-portion of electronic content to at least one of the stored user profiles or to a guest user profile.

2. The computing system of claim 1, wherein the computer-executable instructions are further executable by the one or more processors to cause the computing system to attribute of the sub-portion of electronic content to at least one of the stored user profiles or to a guest user profile by at least causing the computing system to:

determine whether the calculated probability meets or exceeds a user profile confidence threshold; and

attribute the sub-portion of electronic content to the at least one of the stored user profiles when the calculated probability meets or exceeds the user profile confidence threshold, or alternatively, attribute the sub-portion of electronic content to the guest user

profile when the calculated probability meets or exceeds the user profile confidence threshold.

3. The computing system of claim 2, wherein the sub-portion of electronic content is attributed to the guest user profile, the computing system is further caused to at least:

generate a request for the guest user profile to be tagged to correspond to a contributing entity associated with the sub-portion of electronic content for linking the contributing entity to the guest user profile; and

in response to user input that identifies the contributing entity, tag the guest user profile to correspond with the identified contributing entity, thereby linking the contributing entity to the guest user profile.

4. The computing system of claim 2, wherein the sub-portion of the electronic content is attributed to the guest user profile, the computing system is further caused to at least:

generate a request for the guest user profile to be tagged to correspond to a contributing entity associated with the sub-portion of electronic content;

in response to user input denying the request, store the guest user profile in the database of stored user profiles or discard the guest user profile.

5. The computing system of claim 1, wherein the first attribute of the sub-portion of electronic content corresponds to a predetermined set of relevance criteria of the sub-portion of electronic content relative to an entire context of the electronic content, wherein the determination to attribute the sub-portion of the electronic content to at least one user profile is at least in part based on a value of the level of relevance or a value of the level of importance exceeding a value of a pre-determined attribution level threshold.

6. The computing system of claim 5, wherein the predetermined set of relevance criteria includes one or more of the following: a time duration, a word count, one or more keywords pre-determined to indicate a level of importance or level of relevance, or one or more unrecognizable words or intervals of the sub-portion of electronic content.

7. The computing system of claim 1, wherein the computer-executable instructions are further executable by the one or more processors to cause the computing system to at least:

apply a labeling to audio speech included the identified sub-portion of the electronic content with an automatic speech recognition (ASR) model;

calculate a confidence score associated with a probability of accuracy of the labeling

characterizing the audio speech;

compare the confidence score to a predetermined confidence score threshold;

determine to attribute the sub-portion of electronic content to at least one of a stored or guest user profile when the confidence score meets or exceeds the predetermined confidence score threshold or to, alternatively, refrain from attributing the sub-portion of electronic content to the at least one of stored or guest user profile when the confidence score fails to meet or exceed the predetermined confidence score threshold.

8. The computing system of claim 1, wherein the second attribute of the sub-portion of electronic content corresponds to one or more attributes of a stored user profile that is associated with a contributing entity.

9. The computing system of claim 8, wherein the second attribute of the sub-portion of electronic content includes a voice profile compiled from acoustic data extracted from the sub-portion of electronic, wherein the voice profile matches a user voice profile that is part of the stored user profile.

10. The computing system of claim 8, wherein the second attribute of the sub-portion of electronic content includes a face profile compiled from image data extracted from a sub-portion of visual content included in the electronic content, wherein the face profile matches a user face profile that is part of the stored user profile.

11. The computing system of claim 1, wherein the sub-portion of electronic content is attributed to a stored user profile, wherein the computing system is further caused to at least:

update the stored user profile based on identified attributes of the sub-portion of electronic content, the identified attributes including one or more of the following: acoustic features extracted from a sub-portion of audio speech included in the sub-portion of electronic content or visual features extracted from a sub-portion of visual content included in the electronic content.

12. The computing system of claim 1, wherein the computer-executable instructions are further executable by the one or more processors to cause the computing system to:

store the identified attributes corresponding to the sub-portion of electronic content as cached data associated with one or more stored user profiles; and

refrain from updating the one or more stored user profiles with the cached data.

13. The computing system of claim 1, wherein the one or more stored user profiles corresponds to a plurality of meeting participants associated with a particular

meeting from which the electronic content is extracted.

14. The computing system of claim 1, wherein the computer-executable instructions are further executable by the one or more processors to cause the computing system to at least create a transcription of audio speech included in the electronic content with an automatic speech recognition (ASR) model, wherein the transcription is divided based on one or more identified sub-portions of electronic content, and wherein each sub-portion of electronic content is labeled with an identifier associated with a user profile to which the sub-portion of electronic content is attributed.

15. The computing system of claim 14, wherein the transcription includes a plurality of links that point to external data related to the electronic content, wherein the external data related to the electronic content is external to the electronic content, and wherein the external data corresponding to a particular sub-portion of electronic content is stored in the user profile to which the particular sub-portion of electronic content is attributed.




FIG. 1




FIG. 2




FIG. 3




FIG. 4




FIG. 5




FIG. 6




FIG. 7




FIG. 8




FIG. 9




FIG. 10

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2021/023378

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06Q10/10
ADD.

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)
G06Q

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

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

EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2019/392837 A1 (JUNG CHRISTOPHER ANDREWS [US] ET AL) 26 December 2019 (2019-12-26) paragraphs [0002], [0004], [0010], [0011], [0033] - [0038]; figure 1 paragraphs [0048] - [0050], [0055] paragraphs [0060] - [0067]; figure 4 paragraphs [0074] - [0076]; figure 6 ----- -/--	1-15



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 another citation or other special reason (as specified)

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search

25 May 2021

Date of mailing of the international search report

04/06/2021

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040,
Fax: (+31-70) 340-3016

Authorized officer

Wohner, Wolfgang

INTERNATIONAL SEARCH REPORT

International application No

PCT/US2021/023378

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	HUI LIU ET AL: "ProMetheus", MOBILE AND UBIQUITOUS SYSTEMS: COMPUTING, NETWORKING AND SERVICES, ACM, 2 PENN PLAZA, SUITE 701NEW YORKNY10121-0701USA, 5 November 2018 (2018-11-05), pages 392-401, XP058423158, DOI: 10.1145/3286978.3286995 ISBN: 978-1-4503-6093-7 page 393, right-hand column, paragraph 2-3; figure 1 page 395, left-hand column, paragraph 1 - right-hand column, paragraph 1; figure 2 page 396, left-hand column, paragraph 1 - page 397, left-hand column, paragraph 1 page 397, right-hand column, paragraph 3 - page 398, left-hand column, paragraph 3 -----	1-15
X	US 2015/106091 A1 (WETJEN SPENCE [US] ET AL) 16 April 2015 (2015-04-16) paragraphs [0072] - [0075] paragraphs [0077] - [0081]; figure 4 paragraphs [0082] - [0107]; figure 5 paragraph [0155]; figure 8 -----	1-15

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2021/023378

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2019392837 A1	26-12-2019	EP 3797413 A1	31-03-2021
		US 2019392837 A1	26-12-2019
		WO 2019245770 A1	26-12-2019

US 2015106091 A1	16-04-2015	NONE	
