



(19) **United States**

(12) **Patent Application Publication**

Prabhavalkar et al.

(10) **Pub. No.: US 2020/0402501 A1**

(43) **Pub. Date: Dec. 24, 2020**

(54) **CONTEXTUAL BIASING FOR SPEECH RECOGNITION**

(71) Applicant: **Google LLC**, Mountain View, CA (US)

(72) Inventors: **Rohit Prakash Prabhavalkar**, Santa Clara, CA (US); **Golan Pundak**, New York, NY (US); **Tara N. Sainath**, Jersey City, NJ (US); **Antoine Jean Bruguier**, Milpitas, CA (US)

(73) Assignee: **Google LLC**, Mountain View, CA (US)

(21) Appl. No.: **16/863,766**

(22) Filed: **Apr. 30, 2020**

Related U.S. Application Data

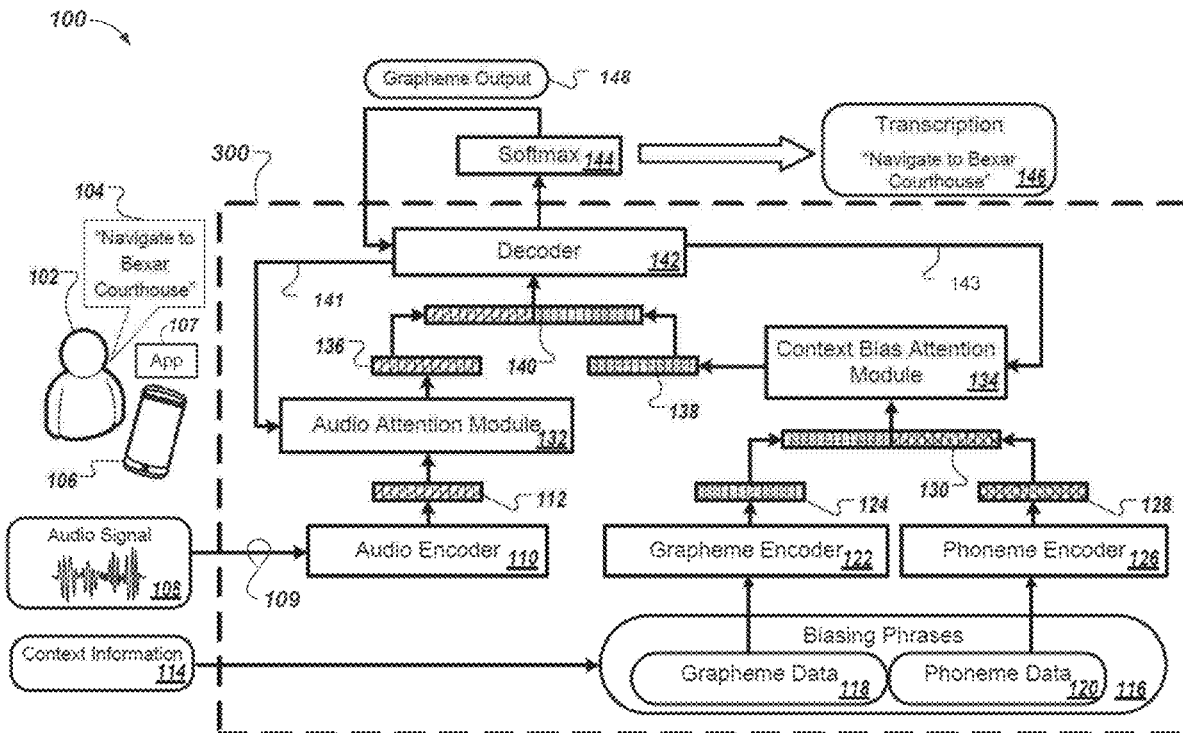
(60) Provisional application No. 62/863,308, filed on Jun. 19, 2019.

Publication Classification

(51) **Int. Cl.**
G10L 15/187 (2006.01)
G10L 19/04 (2006.01)
G06N 20/10 (2006.01)
(52) **U.S. Cl.**
CPC *G10L 15/187* (2013.01); *G10L 2015/088* (2013.01); *G06N 20/10* (2019.01); *G10L 19/04* (2013.01)

(57) **ABSTRACT**

A method of biasing speech recognition includes receiving audio data encoding an utterance and obtaining a set of one or more biasing phrases corresponding to a context of the utterance. Each biasing phrase in the set of one or more biasing phrases includes one or more words. The method also includes processing, using a speech recognition model, acoustic features derived from the audio data and grapheme and phoneme data derived from the set of one or more biasing phrases to generate an output of the speech recognition model. The method also includes determining a transcription for the utterance based on the output of the speech recognition model.



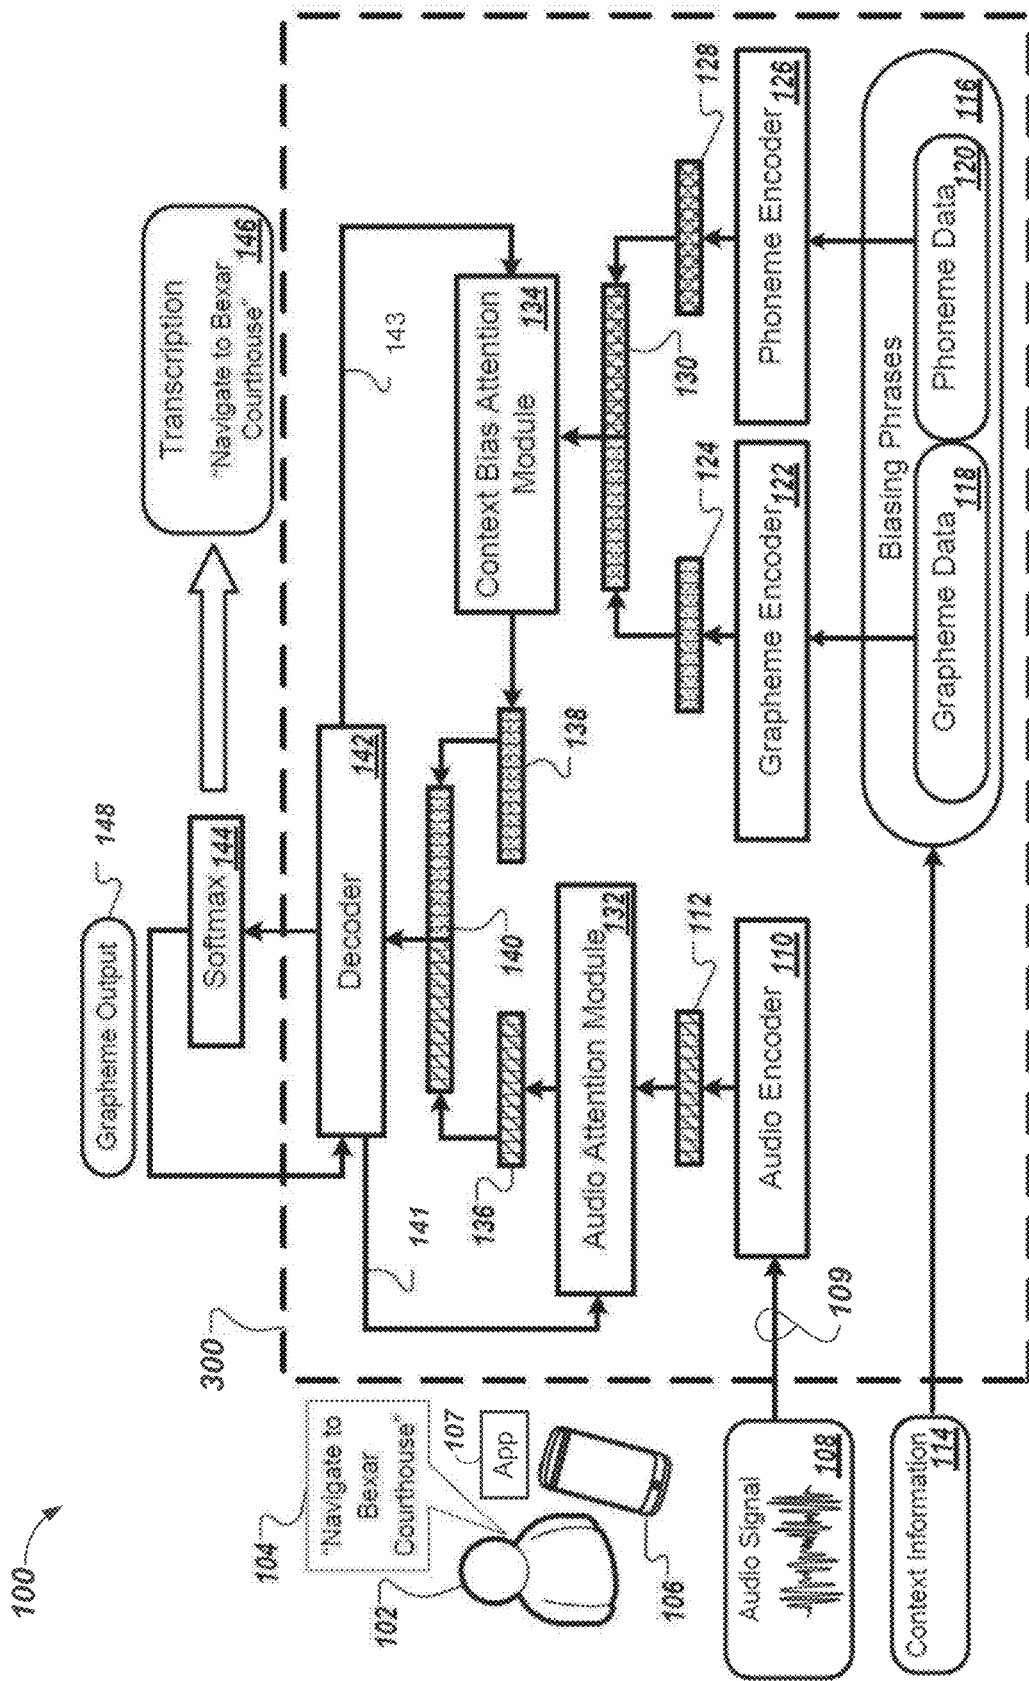


FIG. 1

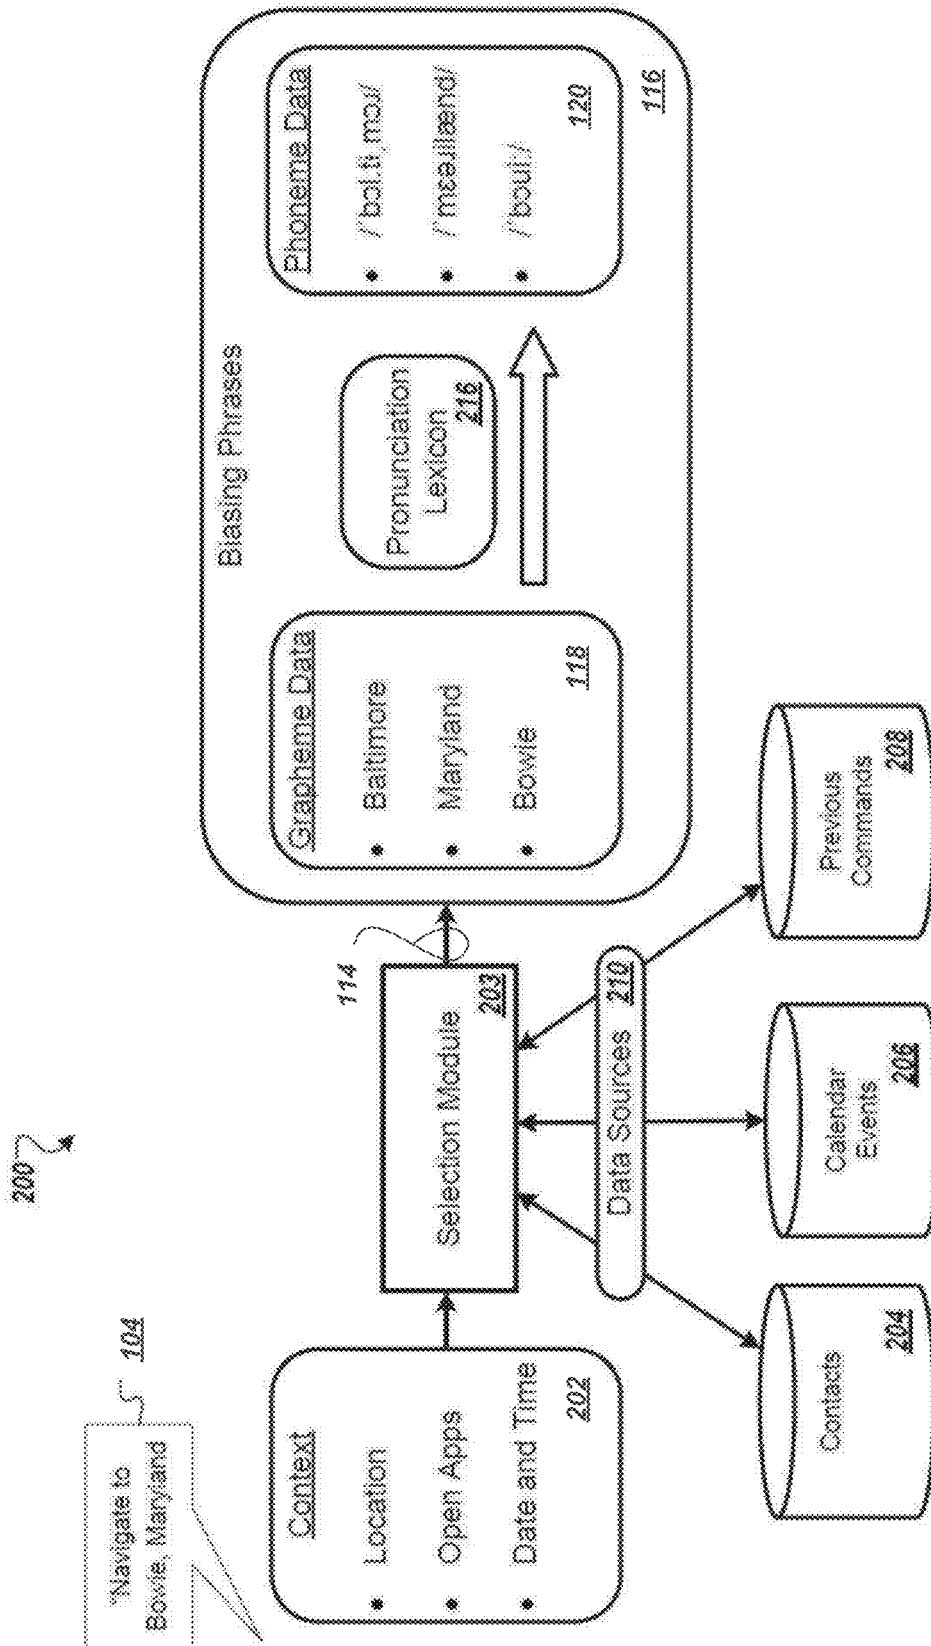


FIG. 2

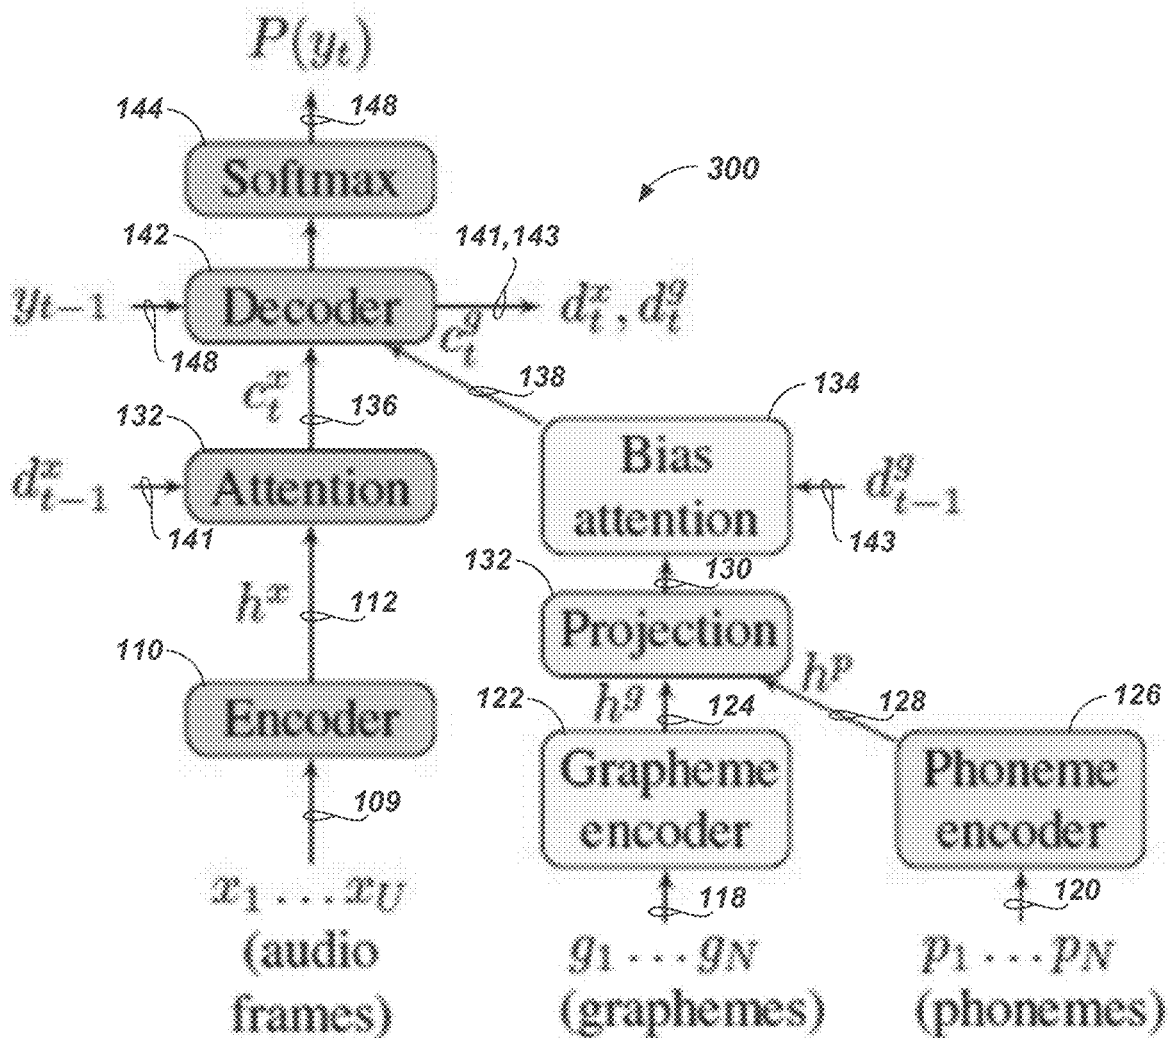


FIG. 3

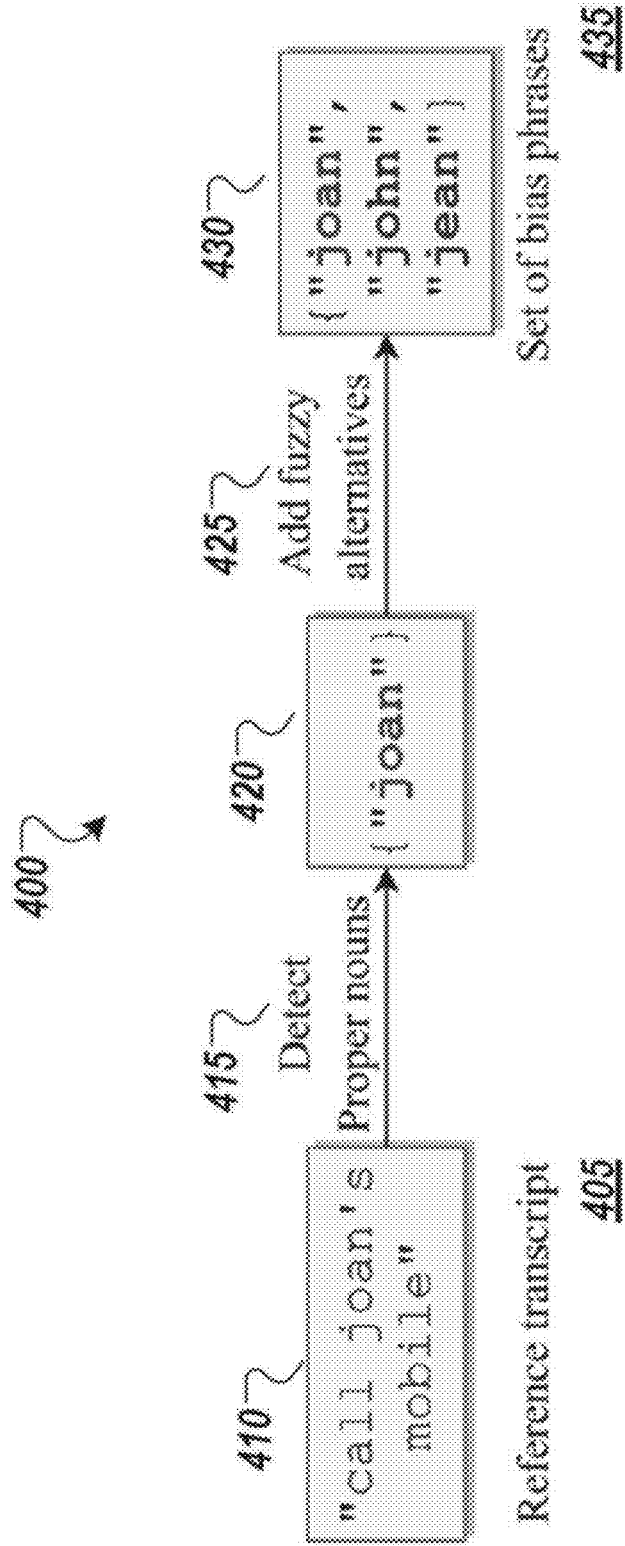


FIG. 4

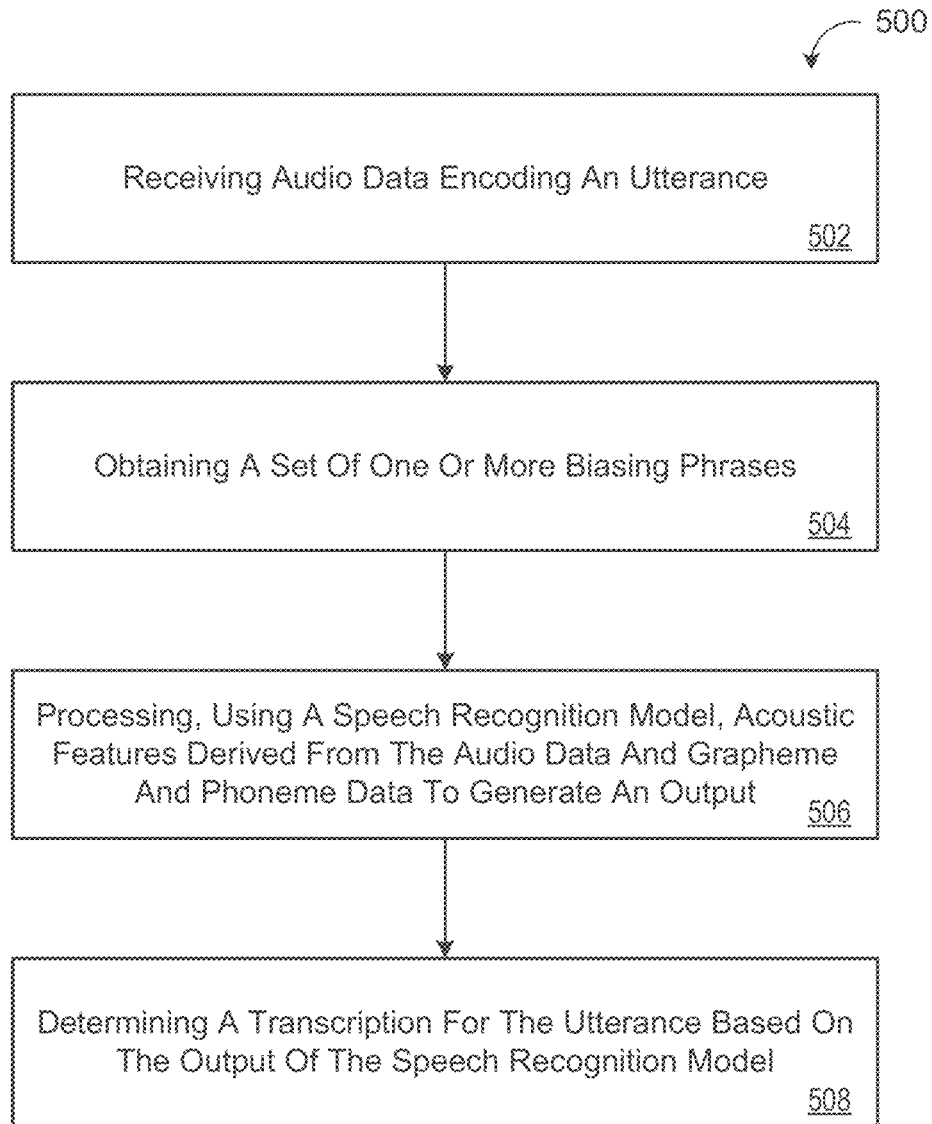


FIG. 5

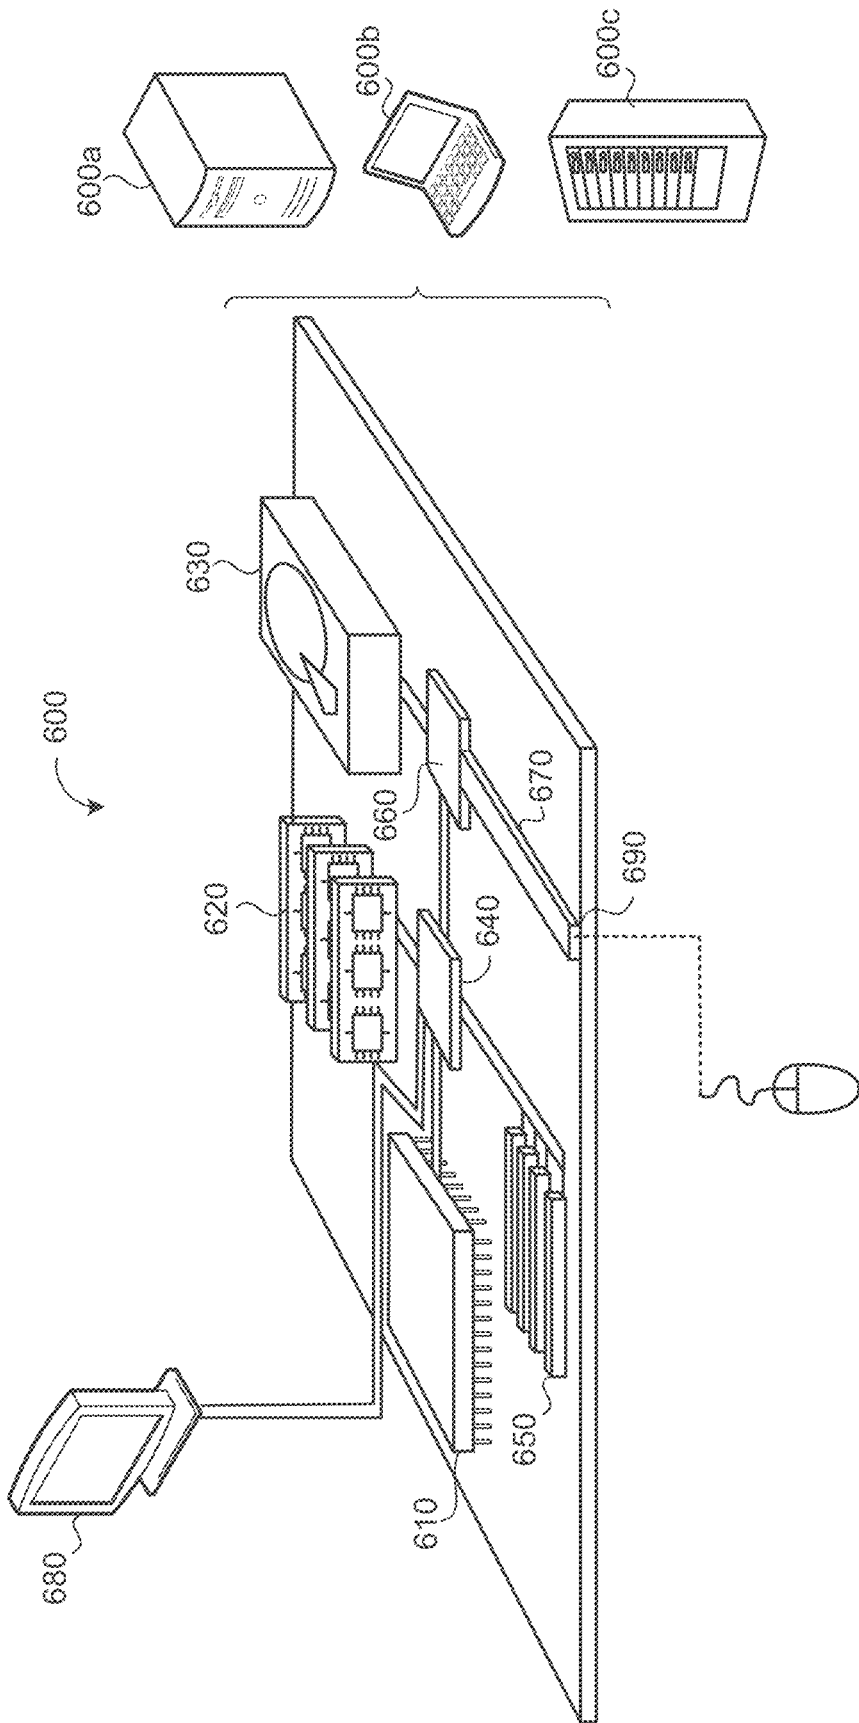


FIG. 6

CONTEXTUAL BIASING FOR SPEECH RECOGNITION

CROSS REFERENCE TO RELATED APPLICATIONS

[0001] This U.S. patent application claims priority under 35 U.S.C. § 119(e) to U.S. Provisional Application 62/863, 308, filed on Jun. 19, 2019. The disclosure of this prior application is considered part of the disclosure of this application and is hereby incorporated by reference in its entirety.

TECHNICAL FIELD

[0002] This disclosure relates to contextual biasing for speech recognition.

BACKGROUND

[0003] Recognizing the context of speech is the goal of automatic speech recognition (ASR) systems. The ability to recognize context in speech, however, is challenging given the wide variety of words that people may speak and the many variations in accents and pronunciation. In many cases, the types of words and phrases that a person speaks varies depending on the context the person finds himself or herself in.

[0004] Contextual automated speech recognition (ASR) involves biasing speech recognition towards a given context, such as towards a user's own playlist, contacts, or geographic place names. Context information usually includes a list of relevant phrases to be recognized, which often includes rare phrases or even foreign words which are seen infrequently in training. To perform contextual biasing, conventional ASR systems sometimes model contextual information in an independent contextual language model (LM), using an n-gram weighted finite state transducer (WFST), and compose the independent contextual LM with a baseline LM for on-the-fly (OTF) rescoring.

[0005] Recently, end-to-end (E2E) models have shown great promise for ASR, exhibiting improved word error rates (WERs) and latency metrics as compared to conventional on-device models. These E2E models, which fold the acoustic model (AM), pronunciation model (PM), and LMs into a single network to directly learn speech-to-text mapping, have shown competitive results compared to conventional ASR systems which have a separate AM, PM, and LMs. Representative E2E models include word-based connectionist temporal classification (CTC) models, recurrent neural network transducer (RNN-T) models, and attention-based models such as Listen, Attend, and Spell (LAS). Because E2E models maintain a limited number of recognition candidates during beam-search decoding, contextual ASR can be challenging for E2E models.

SUMMARY

[0006] Implementations herein are directed toward biasing a speech recognition model toward a set of phrases relevant to a current context. The set of phrases are dynamic in that the phrases may change as the context changes. The speech recognition model may further take into account the pronunciations of the phrases used for biasing in addition to written representations of the phrases. For example, the model can use text information (e.g., grapheme sequences) and pronunciation information (e.g., phoneme sequences)

representing biasing phrases to select which phrases are most applicable to speech being recognized. This allows the contextual biasing aspect of the speech model to be aware of the pronunciations of individual biasing phrases, thereby enabling the model to achieve much higher accuracy when recognizing rare words and words with unusual pronunciations. The contextual biasing techniques discussed herein can be used with various types of speech recognition models, including end-to-end models that can generate transcription data without a separate acoustic model, pronunciation model, and language model.

[0007] To provide high speech recognition accuracy, general purpose automatic speech recognition (ASR) systems may use information indicating the context in which speech occurs. The context may be related to, for example, the user's contacts, calendar appointments, open apps, and location. One example of an ASR system contains separate acoustic, pronunciation, and language models. Other ASR systems combine the acoustic, pronunciation, and language models as a single neural network. A single neural network model improves simplicity and quality, and optimizes word error rate (WER).

[0008] For conventional neural network-based ASR models, it is challenging to recognize words that infrequently appear in a language. The distribution of words in a language typically follows a Zipfian distribution, where a small number of words are used very frequently, and vast numbers of words are rarely used. It can be difficult to obtain a large enough set of training data with audio and corresponding text to effectively provide examples of rare words. Even if sufficient training data is available, adding more and more training examples often yields improvements of lower and lower magnitude. Additionally, it is challenging for these models to recognize words with unusual pronunciations relative to their spelling.

[0009] Incorporating contextual biasing into a neural network-based ASR model can improve recognition for rare words and words with unusual pronunciations. One useful application is to better recognize proper names (i.e., proper nouns such as people names, song names, city names, etc.), which may be relatively rare in training data and/or may not follow typical pronunciation rules. For example, a smart phone or other user device often stores a user's contacts. When a user is using a messaging or phone calling application, this information can be used as context to help the ASR system recognize names spoken by a user. In a similar manner, a list of songs in a music library of a user can be used to bias speech recognition, for example, when the user is using a media player application. Implementations herein are directed toward applying contextual biasing to an ASR model by injecting information associated with both written forms and pronunciations of biasing phrases into the ASR model. Contextually biasing the ASR model incorporates knowledge of rare word pronunciations even if the words were not observed during training of the ASR model.

[0010] One aspect of the disclosure provides a method for biasing speech recognition that includes receiving, at data processing hardware, audio data encoding an utterance, and obtaining, by the data processing hardware, a set of one or more biasing phrases corresponding to a context of the utterance, each biasing phrase in the set of one or more biasing phrases includes one or more words. The method also includes processing, by the data processing hardware, using a speech recognition model, acoustic features derived

from the audio data and grapheme and phoneme data derived from the set of one or more biasing phrases to generate an output of the speech recognition model. The method also includes determining, by the data processing hardware, a transcription for the utterance based on the output of the speech recognition model.

[0011] Implementations of the disclosure may include one or more of the following optional features. In some implementations, the speech recognition model includes a first encoder, a first attention module, a grapheme encoder, a phoneme encoder, a second attention module, and a decoder. The first encoder is configured to receive, as input, the acoustic features and generate, as output, audio encodings from the acoustic features. The first attention module is configured to receive, as input, the audio encodings output from the first encoder and generate, as output, first attention outputs. The grapheme encoder is configured to receive, as input, grapheme data indicating graphemes of each word in the set of one or more biasing phrases and generate, as output, grapheme encodings. The phoneme encoder is configured to receive, as input, phoneme data indicating phonemes of each word in the set of one or more biasing phrases and generate, as output, phoneme encodings. The second attention module is configured to receive, as input, a representation of the grapheme encodings output from the grapheme encoder and the phoneme encodings output from the phoneme encoder and generate, as output, second attention outputs. The decoder is configured to determine likelihoods of sequences of speech elements based on the first attention outputs and the second attention outputs.

[0012] In some examples, for each particular word of each biasing phrase in the set of one or more biasing phrases: the grapheme encoder is configured to generate a corresponding grapheme encoding for the particular word; the phoneme encoder is configured to generate a corresponding phoneme encoding for the particular word; and the second attention module is configured to encode a corresponding second attention output that comprises a corresponding contextual biasing vector for the particular word based on the corresponding grapheme and phoneme encodings for the particular word. In addition examples, the representation of the grapheme encodings output from the grapheme encoder and the phoneme encodings output from the phoneme encoder includes a projection vector representing a concatenation between the grapheme encodings and the phoneme encodings. Additionally, the first encoder, the first attention module, the grapheme encoder, the phoneme encoder, the second attention module, and the decoder may be trained jointly to predict a sequence of graphemes from a sequence of acoustic feature frames.

[0013] In some implementations, the first attention module is configured to compute attention as a function of a previous hidden state of the decoder and a full sequence of audio encodings output by the first encoder, while the second attention module is configured to compute attention as a function of the previous hidden states of the decoder and a full sequence of projection vectors representing characteristics of both the grapheme and phoneme data derived from the biasing phrases in the set of one or more biasing phrases. In additional implementations, the second attention module is configured to receive a contextual biasing vector that does not correspond to any of the biasing phrases in the set of one or more biasing phrases, the contextual biasing vector representing an option to not bias the output of the speech

recognition model. In some examples, the speech elements include graphemes. In other examples, the speech elements include words or wordpieces. Optionally, the speech elements may include phonemes.

[0014] The set of one or more biasing phrases may include one or more contact names personalized for a particular user. Additionally or alternatively, the set of one or more biasing phrases may include one or more calendar events personalized for the particular user. In some examples, the method also includes determining, by the data processing hardware, the context of the utterance based on at least one of: a location of a user that spoke the utterance; one or more applications open on a user device associated with the user that spoke the utterance; or a current date and/or time of the utterance. The context of the utterance may additionally or alternatively be based on one or more previous commands issued by the user.

[0015] In some implementations, the speech recognition model includes a decoder configured to determine a hidden state and the output of the speech recognition model based on an embedding vector, a previous hidden state of the decoder; a first vector; and a second vector. Here, the embedding vector is for a previous grapheme output by the speech recognition model, while the first vector is output by a first attention module and the second vector is output by a second attention module.

[0016] Another aspect of the disclosure provides a system for biasing speech recognition based on a context of an utterance. The system includes data processing hardware and memory hardware in communication with the data processing hardware and storing instructions that when executed on the data processing hardware cause the data processing hardware to perform operations. The operations include receiving audio data encoding an utterance and obtaining a set of one or more biasing phrases corresponding to a context of the utterance, each biasing phrase in the set of one or more biasing phrases includes one or more words. The operations also include processing, using a speech recognition model, acoustic features derived from the audio data and grapheme and phoneme data derived from the set of one or more biasing phrases to generate an output of the speech recognition model. The operations also include determining a transcription for the utterance based on the output of the speech recognition model.

[0017] This aspect may include one or more of the following optional features. In some implementations, the speech recognition model includes a first encoder, a first attention module, a grapheme encoder, a phoneme encoder, a second attention module, and a decoder. The first encoder is configured to receive, as input, the acoustic features and generate, as output, audio encodings from the acoustic features. The first attention module is configured to receive, as input, the audio encodings output from the first encoder and generate, as output, first attention outputs. The grapheme encoder is configured to receive, as input, grapheme data indicating graphemes of each word in the set of one or more biasing phrases and generate, as output, grapheme encodings. The phoneme encoder is configured to receive, as input, phoneme data indicating phonemes of each word in the set of one or more biasing phrases and generate, as output, phoneme encodings. The second attention module is configured to receive, as input, a representation of the grapheme encodings output from the grapheme encoder and the phoneme encodings output from the phoneme encoder

and generate, as output, second attention outputs. The decoder is configured to determine likelihoods of sequences of speech elements based on the first attention outputs and the second attention outputs.

[0018] In some examples, for each particular word of each biasing phrase in the set of one or more biasing phrases: the grapheme encoder is configured to generate a corresponding grapheme encoding for the particular word; the phoneme encoder is configured to generate a corresponding phoneme encoding for the particular word; and the second attention module is configured to encode a corresponding second attention output that comprises a corresponding contextual biasing vector for the particular word based on the corresponding grapheme and phoneme encodings for the particular word. In addition examples, the representation of the grapheme encodings output from the grapheme encoder and the phoneme encodings output from the phoneme encoder includes a projection vector representing a concatenation between the grapheme encodings and the phoneme encodings. Additionally, the first encoder, the first attention module, the grapheme encoder, the phoneme encoder, the second attention module, and the decoder may be trained jointly to predict a sequence of graphemes from a sequence of acoustic feature frames.

[0019] In some implementations, the first attention module is configured to compute attention as a function of a previous hidden state of the decoder and a full sequence of audio encodings output by the first encoder, while the second attention module is configured to compute attention as a function of the previous hidden states of the decoder and a full sequence of projection vectors representing characteristics of both the grapheme and phoneme data derived from the biasing phrases in the set of one or more biasing phrases. In additional implementations, the second attention module is configured to receive a contextual biasing vector that does not correspond to any of the biasing phrases in the set of one or more biasing phrases, the contextual biasing vector representing an option to not bias the output of the speech recognition model. In some examples, the speech elements include graphemes. In other examples, the speech elements include words or wordpieces. Optionally, the speech elements may include phonemes.

[0020] The set of one or more biasing phrases may include one or more contact names personalized for a particular user. Additionally or alternatively, the set of one or more biasing phrases may include one or more calendar events personalized for the particular user. In some examples, the operations also include determining the context of the utterance based on at least one of: a location of a user that spoke the utterance; one or more applications open on a user device associated with the user that spoke the utterance; or a current date and/or time of the utterance. The context of the utterance may additionally or alternatively be based on one or more previous commands issued by the user.

[0021] In some implementations, the speech recognition model includes a decoder configured to determine a hidden state and the output of the speech recognition model based on an embedding vector, a previous hidden state of the decoder; a first vector; and a second vector. Here, the embedding vector is for a previous grapheme output by the speech recognition model, while the first vector is output by a first attention module and the second vector is output by a second attention module.

[0022] The details of one or more implementations of the disclosure are set forth in the accompanying drawings and the description below. Other aspects, features, and advantages will be apparent from the description and drawings, and from the claims.

DESCRIPTION OF DRAWINGS

[0023] FIG. 1 is a schematic view of an example automatic speech recognition system including a speech recognition model that biases speech recognition results based on a context relevant to an utterance.

[0024] FIG. 2 is a schematic view of an example selection process for selecting biasing phrases.

[0025] FIG. 3 is a schematic view of an example architecture of the speech recognition model of FIG. 1.

[0026] FIG. 4 is a schematic view of an example training process for training the speech recognition model of FIG. 1.

[0027] FIG. 5 is a flowchart of an example arrangement of operations for a method of contextual-biasing a speech recognition model.

[0028] FIG. 6 is a schematic view of an example computing device that may be used to implement the systems and methods described herein.

[0029] Like reference symbols in the various drawings indicate like elements.

DETAILED DESCRIPTION

[0030] Conventional automated speech recognition (ASR) systems rely on three models: an acoustic model, a pronunciation model, and a language model. Design decisions for the pronunciation model influence both the acoustic model and the language model. For the pronunciation model, a set of phonemes is chosen to represent all of the possible distinct units of sound, thus determining the set of labels the acoustic model has to predict. All of the words in the language model have a pronunciation.

[0031] An accurate pronunciation model has a large coverage. For a given variant of a language, the phoneme set varies from region to region. For example, in American English, native speakers do not agree whether words like “pen” and “pin” are pronounced the same way. Additionally, certain words can have reduced pronunciations, where speakers may or may not pronounce the letter “t” in words like “twenty.” For extremely common words like “the,” it is unnatural to hard-code pronunciations.

[0032] End-to-end (E2E) speech recognition models combine the acoustic, pronunciation, and language models into a single neural network. A single neural network model improves simplicity and quality, and optimizes word error rate (WER). However, a challenge in E2E speech recognition models is optimizing performance on recognizing words that appear infrequently in a language and/or have unusual pronunciations relative to their spelling. While training data can include both human-transcribed voice data and text-only data, the use of large training data sets for training these E2E speech recognition models is inefficient. As the distribution of words in a language typically follow a Zipfian distribution, where a small number of words are used very frequently, and vast numbers of words are rarely used, increasing the number of training examples typically yields improvements of lower and lower magnitude.

[0033] Incorporating contextual biasing into a neural network ASR model can improve recognition for rare words

and words with unusual pronunciations. For example, since a user's contacts are often stored on a smart phone, they can be used as context to help the ASR system recognize the names of contacts spoken by a user. Contextual biasing can be applied to ASR models by injecting both biasing context and pronunciation into the model. The contextually biased model retains the advantages of neural network models, including simple, unified training and implicit learning of the pronunciation of rare words. The contextually biased model incorporates knowledge of rare word pronunciation even if the words have never been present during training.

[0034] Referring to FIG. 1, an example ASR system 100 includes a speech recognition model 300 that incorporates contextual biasing to improve speech recognition accuracy. The speech recognition model 300 includes an audio encoder 110, a grapheme encoder 122, a phoneme encoder 126, an audio attention module 132, a contextual biasing attention module 134, and a decoder 142 that receives, as input, outputs from both of the audio and contextual biasing attention modules 132, 134. In some examples, an output of the decoder 142 is normalized with a softmax layer 144, whereby the softmax layer 144 outputs a probability distribution over sequences of output graphemes. The softmax layer 144 may provide a probability distribution over a set of language units (e.g., speech elements), such as a set of graphemes. For example, the output for the softmax layer 144 can be a vector having a probability value for each of the graphemes in a language, as well as potentially for other symbols (e.g., punctuation, space, etc.). The sequence of vectors from the softmax layer 144 is used to produce a transcription 150 of a speech utterance 104 recorded by a user device 106. In some examples, the speech recognition model 300 resides on a user device 106 associated with a user 102. In other examples, the speech recognition model 300 resides on a remote server in communication with the user device 106 or functionality of the speech recognition model 200 is split among the remote server and the user device 106.

[0035] In the example shown, the user 102 speaks an utterance 104 captured by one or more microphones of the user device 106. The user device 106 may include a mobile device, such as a smart phone, tablet, smart headphones, smart watch, etc. The utterance 104 spoken by the user 102 may be a command, "Navigate to Bexar Courthouse", in which "Bexar" is a name of a county pronounced /ber/, similar to the words "bear" and "bare". The one or more microphones of the user device 106 generate an audio signal 108 from the spoken command. The audio signal (interchangeably referred to as "audio data") encodes the utterance 104 and can be processed in any of various ways to determine inputs for the neural network elements of the speech recognition model 300. For example, a feature extraction module (not shown) can generate acoustic features 109 derived from the audio signal 108, such as Mel-frequency cepstral components for different windows or frames of the audio signal 108.

[0036] The user device 106 also provides context information 114 to the speech recognition model 300. Context information 114 can include, for example, a user's location, the apps that are open on the device 104, contacts of the user, and upcoming calendar appointments. In this case, context information 114 may include the user's current GPS location within Bexar County, a navigation application 107 open on the user device 106, and the user's upcoming calendar

appointment at the courthouse. The calendar appointment could be obtained by accessing a calendar application 107 executable on the user device 106. The user 102 may explicitly opt-in to consent to sharing each type of context information 114 individually, and may opt-out at any time to revoke the previously consented to sharing of one or more types of context information 114.

[0037] The speech recognition model 300 uses the context information 114 received from the user device 106 to select particular biasing phrases 116, each of which may include both grapheme data 118 and phoneme data 120. The grapheme data 118 is input to the grapheme encoder 122, while the phoneme data 120 is input to the phoneme encoder 126. Additional detail regarding the selection of biasing phrases 116 is provided in FIG. 2.

[0038] In some implementations, the audio, grapheme, and phoneme encoders 110, 122, 126 include sequence models having long short-term memory (LSTM) elements or other recurrent neural network (RNN) elements. The audio encoder 110 may include a first encoder configured to generate an encoded audio vector (e.g., audio encodings) 112 based on the audio signal 108. For instance, the audio encoder 110 may receive acoustic features 109 derived from the audio signal 108 encoding the utterance 104 and generate the encoded audio vector 112 from the acoustic features 109. The encoded audio vector 112 is input to the audio attention module 132. The grapheme encoder 122 is configured to generate an encoded grapheme vector 124 from grapheme data 118 associated with the selected biasing phrases 116, while the phoneme encoder 126 is configured to generate an encoded phoneme vector 128 from the phoneme data 120 associated with the selected biasing phrases 116. In the example shown, the speech recognition model 300 concatenates the encoded grapheme vector 124 and the encoded phoneme vector 128 into a projection vector 130 input to the contextual biasing attention module 134. Accordingly, the projection vector 130 may correspond to a representation of the encoded grapheme and phoneme vectors 124, 128.

[0039] The audio attention module 132 is configured to compute attention weights and generate a weighted audio encoding vector 136 that includes a vector summary of which audio frames the model 300 is currently attending to. For instance, the audio encoder 110 can generate an audio vector 112 for each window or frame of audio data 108 as additional acoustic features 109 are received, and so continues to provide additional audio vectors 112 as more audio is received. In some examples, the audio attention module 132 computes attention weights over the sequence of audio vectors 112 using a decoder audio state 141 from a previous time step and the encoded audio vector 112. The decoder audio state 141 may represent a portion of the state of the decoder 142 when outputting the most recent decoder output, e.g., a portion of the decoder state for the previous time step representing the immediately previous window or frame of audio. The decoder state 141 helps inform the audio attention module 132 of the relative importance of the audio vector 112 input to the audio attention module 132 at the current time step. The decoder audio state 141 can be a portion of the output of the decoder 142, such as a designated subset of the outputs of the decoder 142, typically a different set of outputs than used in a decoder context state 143. In some implementations, the state vectors 141, 143 represent non-overlapping portions of the output of the

decoder 142. In other implementations, the state vectors 141, 143 include one or more overlapping sections or both state vectors 141, 143 each represent the whole output of the decoder 142

[0040] At the same time, the contextual biasing attention module 134 is configured to compute attention weights using both of the encoded grapheme and phoneme vectors 124, 128 included in the projection vector 130. The contextual biasing attention module 134 uses these attention weights to only compute a weighted sum of the grapheme embeddings, which are output as a contextual biasing vector 138. More specifically, the contextual biasing attention module 134 processes the projection vectors 130 and a decoder context state 143 from a previous time step to generate the contextual biasing vector 138. The decoder context state 143 includes the state of the decoder 142 at the previous time step with respect to the previous contextual biasing vectors 138. The decoder context state 143 may represent a portion of the output of the decoder 142, such as a designated subset of the outputs of the decoder 142. The decoder context state 143 informs the contextual biasing attention module 134 of the relative importance of the projection vector 130 input to the contextual biasing attention module 134 at the current time step. Thus, as the decoder context state 143 changes, the contextual biasing attention module 134 calculates a different summary or aggregation of the projection vectors 130, thus changing the probabilities that the decoder 142 will indicate for elements of the different biasing phrases 116.

[0041] By having the contextual biasing attention module 134 compute only the weighted sum of the grapheme embeddings, the model 300 does not learn an a-priori relationship between spelling and pronunciation, thereby permitting the model to assign an arbitrary sequence of phonemes to any word. However, since the contextual biasing attention module 134 does use the encoded phoneme vector 128, in addition to the encoded grapheme vector 124, to compute the attention weights, the weighted sum of the grapheme embeddings computed from the attention weights contains phoneme information in addition to grapheme information.

[0042] Notably, the contextual biasing vector 138 output from the contextual biasing attention module 134 represents contextual biasing of the speech recognition model 300, and includes both grapheme and phoneme information. Accordingly, concatenating the contextual biasing vector 138 with the weighted audio encoding vector 136 into a weighted vector “injects” contextual biasing into the speech recognition model 300. The weighted vector 140 collectively represents the audio, grapheme, and phoneme information. The weighted vector 140 is input to the decoder 142.

[0043] In some configurations, the decoder 142 includes a sequence model such as an LSTM network and is configured to extract an output sequence from the weighted vector 140. The output of the decoder 142 is also normalized with a softmax layer 144 to produce a probability distribution over a set of output targets, which are graphemes in the illustrated example. In other examples, the set of output targets in the probability distribution may include words, wordpieces or phonemes. A beam search process may use the various output vectors from the decoder 142 to produce the transcription 146 for the utterance 104. The decision made about which grapheme represents the audio at the current time step is fed back to the decoder 142 as a speech element output

148 and is used to compute the next output from the decoder 142. In the example shown, the speech element output 148 includes a grapheme output. In other examples, the speech element output 148 includes a word output or a wordpiece output. In yet other examples, the speech element output 148 includes a phoneme output.

[0044] In the example of FIG. 1, the context information 114 of the user’s 102 location, open apps, and upcoming calendar appointment biases the speech recognition model 300 to the biasing phrase “Bexar” and the pronunciation /ber/. This enables the speech recognition model 300 to correctly transcribe the user’s 102 command as “Navigate to Bexar Courthouse.” The user device 106 can then perform an action based on this transcription, such as mapping a route to the courthouse within a navigation app.

[0045] FIG. 2 shows a schematic view of an example process 200 for generating a set of one or more biasing phrases 116 relevant to a context of an utterance 104. The process 200 includes identifying a context 202 of the speech utterance 104, and selecting context information 114 from data sources 210 to generate the biasing phrases 116, including both grapheme data 118 and phoneme data 120. Each biasing phrase 116 in the set of one or more biasing phrases includes one or more words, such that the grapheme data 118 includes the graphemes associated with each word of the biasing phrases 116 and the phoneme data 118 includes the phonemes associated with each word of the biasing phrases 116.

[0046] In more detail, the process 200 begins when a user speaks an utterance 104. When an ASR system receives audio input from the spoken utterance 104, one or more contexts 202 of the utterance 104 are determined. The context 202 determines the biasing phrases 116 that are selected for use in contextual biasing.

[0047] The context 202 can be, for example, the location of the user. The location of a user can be determined by GPS coordinates. For example, the context could be that the user is at home, work, or a geographic location near the user’s home or work. The user’s location could also be identified as a venue such as a restaurant, concert hall, or stadium. If the user is traveling, the context could be that the user is on vacation or traveling for work, perhaps to an unfamiliar or foreign city.

[0048] Another example of context 202 is the apps that are open on a user’s mobile device. If the user opens the telephone app, the user may be about to make a telephone call to a contact. If the user opens a navigation app, the user may be about to navigate to a destination.

[0049] An example of a context 202 can also be the date and time of the spoken utterance. For example, if it is early on a Monday morning, the user may be preparing to drive to work. Thus, the user may speak an utterance 104 that will help him or her prepare for the work day, such as to ask for weather or traffic conditions or to navigate to the office.

[0050] A selection module 203 references the context 202 to select information 114 from various data sources 210. The data sources 210 can include, for example, contacts 204, calendar events 206, and previously spoken commands 208. In one example, if the context 202 includes an open telephone app, the selection module 203 may select information 114 from stored contacts 204. If the context 202 includes an open navigation app, the selection module 203 may select information 114 regarding the location of upcoming calendar events. The selection module 203 may also reference

previous commands **208**. For example, the user may ask to navigate to work every weekday morning. Therefore, an utterance **104** spoken on a Monday morning will likely include words that have recently been spoken by the user related to the user's work location. The context information **114** may also specify the one or more contexts **202** relevant to the spoken utterance **104**.

[0051] The information **114** obtained/selected by the selection module **203** is used to create the biasing phrases **116**. Biasing phrases **116** include both grapheme data **118** and phoneme data **120**. The phoneme data **120** may be generated using a pronunciation lexicon **216**.

[0052] As an example, a user may leave home most weekday mornings, drive to Bowie, Md. to drop off a child at daycare, then continue to Baltimore, Md., for work. When the user speaks the utterance **104** "Navigate to Bowie, Md.," the context **202** for the utterance **104** can include the location of the user, which is at home, in Maryland. The context **104** can also include an open navigation app, and the date and time, which may be a Monday morning. The selection module **203** accesses the data sources **210**. The calendar events **206** show that the user has a meeting in Baltimore, Md., later that day. The previous commands **208** show that on most weekday mornings, the user asks to navigate to Bowie and then Baltimore. The selection module **203** uses this context information **114** to create biasing phrases **212**.

[0053] In this example, the grapheme biasing data **118** may include the words "Baltimore," "Maryland," and "Bowie," as the user has a navigation app open, and these three words are related to likely destinations of the user. The pronunciation lexicon **216** is used to create phoneme biasing data **120**. The user device **106** accesses the pronunciation lexicon **216** and retrieves pronunciations, e.g., sequences of phonemes, which represent common pronunciation of words. A pronunciation lexicon may include variations of pronunciations of words. For example, someone who is from Baltimore may pronounce "Baltimore" differently from someone who is visiting from another state. Likewise, if a traveler is visiting a country where the spoken language is not their native language, the traveler's pronunciation of a place name is likely different from a native speaker's pronunciation.

[0054] This process **200** generates grapheme biasing data **118** and phoneme biasing data **120**. The biasing phrases **116** are then fed to the grapheme encoder **124** and phoneme encoder **126** for injection into the ASR model **300**, as illustrated in FIG. 1.

[0055] FIG. 3 illustrates an example architecture of the speech recognition model **300** of the ASR system **10** of FIG. 1 for contextual biasing using graphemes and phonemes. The example architecture for the model **300** depicted FIG. 3 includes a listen, attend, and spell (LAS) model. The LAS architecture includes three main components: the audio encoder (e.g., first encoder) **110**; the audio attention module (e.g., first attention module) **132**; and the decoder **142**. The audio encoder **110** is configured to receive a time-frequency representation (i.e., acoustic features) **109** of the input speech audio signal **108** and uses a set of neural network layers to map the input **109** to a higher-level feature representation, i.e., the encoded audio vector (e.g., audio encodings) **112**. The encoded audio vector **112** output of the audio encoder **110** is passed to the audio attention module **132**, which uses the encoded audio vector **112** to learn an alignment between the input acoustic features **109** and predicted

graphemes **148**. The output **136** of the attention audio module **132** is passed to the decoder **142**, which produces a probability distribution over a set of hypothesized graphemes. The components of the LAS model are trained jointly as a single end-to-end neural network.

[0056] The example LAS architecture is enhanced to include additional components to enable the model **300** to bias speech recognition results **146**, **148** toward biasing phrases **116** relevant to a given context **202**. These additional components include the grapheme encoder (e.g., second encoder) **122**, the phoneme encoder (e.g., third encoder) **126**, and the bias attention module (e.g., second attention module) **134**.

[0057] The enhanced architecture of the speech recognition model **300** accepts audio, grapheme, and phoneme inputs x_u , g_n , and p_m , where the subscripts correspond to the U audio frames **109** and the N biasing phrases **116**. The encoders **110**, **122**, **126** encode the respective inputs into encodings h^x , h^g , and h^p , which are then fed into the corresponding attention modules **132**, **134**. The attentions by the attention modules **132**, **134**, in turn, output corresponding context attention vectors c_f^x and c_f^g fed to the decoder **142**. The decoder **142** outputs d_t^x and d_t^g that are used for the next decoding steps to compute the next attentions' outputs. The output of the example LAS model **300** is a vector y_t for every decoding step t from 1 to T .

[0058] In some implementations, the audio encoder **110** includes a bidirectional LSTM configured to receive, as input, audio frames $x_1 \dots x_U$ (e.g., acoustic features **109** such as log-mel features). In one example, the audio encoder **110** includes ten LSTM layers, each having 256 nodes and a hidden size of 1400. Additionally, the grapheme and phoneme encoders **122**, **126** may each include a single layer LSTMs with 512 dimensions. The grapheme encoder **122** and the phoneme encoder **126** may each receive, as input, respective grapheme data **118** and phoneme data **120** associated with N randomly shuffled biasing phrases **116**. Accordingly, the number of phoneme inputs (e.g., phoneme data) **120** that convey pronunciation information may be identical to the number of grapheme inputs (e.g., grapheme data) **118**, and the corresponding embedding/encodings **124**, **128** output from the encoders **122**, **128** are computed independently from one another.

[0059] In some examples, the audio attention module **132** includes four heads with 512 coordinates. In the example shown, the audio attention module **132** receives two inputs: the audio encoding **112**, h^x , that includes the encoding of the audio frames **109**; and the output **141** of the decoder **142** at the previous time step, d_{t-1}^x . The superscript x informs the audio attention module **132** that the output **141** corresponds to a decoder audio state representing a portion of the state of the decoder **142** when outputting the most recent decoder output, e.g., a portion of the decoder state for the previous time step representing the immediately previous window or frame of audio. Thus, the output **141** associated with the decoder audio state helps inform the audio attention module **132** of the relative importance of the audio vector **112**, h^x , input to the audio attention module **132** at the current time step. The output **141** corresponding to the decoder audio state can be a portion of the output of the decoder **142**, such as a designated subset of the outputs of the decoder **142**, typically a different set of outputs than used in the output **143** corresponding to the decoder context state.

[0060] On the other hand, the bias attention module 134 is configured to receive, as input, an encoded projection 130 corresponding to the concatenation between the grapheme encoding 124, h^g , output from grapheme encoder 122 and the phoneme encoding 128, h^p , output from the phoneme encoder 126 based on the grapheme and phoneme inputs 118, 120. A projection layer 132 may concatenate the encodings 124, 128 and output the corresponding encoded projection 130. More specifically, the bias attention module 134 attends, at every decoding step t , on a weighted sum of the encoded projection 130. Here, the bias attention module 134 may use both the grapheme and phoneme encodings 124, 128 to compute the attention weights. Once the attention weights are computed, they are used to compute a weighted sum c_t^g of the grapheme embeddings. Thus, output 138 of the bias attention module 134 includes only the weighted sum c_t^g of the grapheme encodings 124, h^g .

[0061] In more detail, the output 138 of the bias attention module 134 is computed in three steps. In the first step, the bias attention modules 134 receives, as input, the grapheme encodings h_t^g , phoneme encodings h_t^p , and the outputs 143 from the previous decoding steps, d_{t-1}^g . The superscript g informs the bias attention module 134 that the output 143 corresponds to a decoder grapheme state representing a portion of the state of the decoder 142 when outputting the most recent decoder output. In some examples, the bias attention module 134 computes a set of weights for every biasing word and every decoding step u_{it}^g as follows:

$$u_{it}^g = v^g T \cdot \tan h(W_h^g h_t^g + W_h^p h_t^p + W_d^g d_{t-1}^g + b^g) \quad (1)$$

[0062] In the second step, the output is normalized with a softmax 144 to obtain α_t^g as follows:

$$\alpha_t^g = \text{softmax}(u_{it}^g)$$

[0063] In the third step, the output 138 is computed as weighted sum of the embeddings h_t^g .

$$c_t^g = \sum_{i=1}^N \alpha_{it}^g h_i^g$$

[0064] The bias attention output 138, c_t^g , represents an embedding of graphemes. Thus, in some examples the decoder 142 is configured to operate on graphemes only, therefore the decoder input c_t^g is in terms of graphemes only. This enables the injection of words for which the pronunciation is difficult to predict from the spelling. With this approach, there is no a-priori relationship between spelling and pronunciation, and an arbitrary sequence of phonemes can be assigned to any word.

[0065] However, in configurations when the decoder 142 is configured to produce a probability distribution over a set of output targets other than graphemes, such as wordpieces, the bias attention output 138 may be in terms of the type of output target the decoder 142 is configured to produce. Optionally, the output of the bias attention module may represent embeddings of phonemes, e.g., a weighted sum of phonemes, and the decoder 142 may instead operate on phonemes only to produce a set of output targets associated with phonemes.

[0066] In some examples, the decoder 142 includes four LSTM layers, each having a size of 256 with a hidden dimension of 1,024. In the example shown, the decoder 142

receives three inputs. The first input is the previous grapheme output 138, y_{t-1} . When the model is used to make predictions, y_{t-1} is the output of the softmax 144. When the model is training, y_{t-1} can either be a one-hot encoding of the true grapheme, or a sampling of the output and the one-hot encoding of the true grapheme. Another input 136 to the decoder 142, c_t^x , includes a summary of the audio frames 109 the model 300 is currently paying attention to. The third input 138 to the decoder 142 includes the bias attention output cry.

[0067] With continued reference to FIG. 3, the output of the example LAS model 300 includes a vector y_t for every decoding step t from 1 to T . In ASR, the vector y_t corresponds to the one-hot encoding of the output symbols (graphemes, space, and end-of-sequence). The model 300 outputs $T-1$ symbols until it predicts an end-of-sequence symbol at the 7^{th} step.

[0068] In some implementations, a training process trains the speech recognition model 300 on multiple batches of training examples. Here, for each batch, a reference transcript is randomly kept with probability P_{keep} . This allows examples for which there is no matching biasing. For the remaining reference transcripts, n-grams are randomly selected, where n is uniformly sampled from $[1:N_{order}]$. The number of selected n-grams is N , where N is uniformly sampled from $[1:N_{phrases}]$. Example values are $P_{keep}=0.5$, $N_{phrases}=1$, and $N_{order}=4$. Accordingly, the amount of biasing phrases 116 in each batch is variable and not specified at the time of training the model 300. Moreover, one or more biasing phrases 116 applied during training may have different lengths than the other biasing phrases 116. The model 300 can be trained using log-mel acoustic features of every 10 ms, computed for windows of 25 ms, and stacked by groups of 3. An example acoustic feature vector has 80 coordinates.

[0069] A no-biasing option can be added, where the biasing graphemes and phonemes are empty. The no-biasing option can be applied by applying a contextual biasing vector 138 that does not include grapheme and biasing data 118, 120 associated with any biasing phrase 116. This sampling approach ensures that some examples do not have matching words, while others do. A special token is introduced to help the model converge. After each matching bias token, a special character $\langle \text{bias} \rangle$ is added. It introduces a biasing error that can only be corrected by using the biasing phrase.

[0070] For training the phonemes, a mix of lexicon and grapheme to phoneme (G2P) can be used. If a word is in the lexicon, its corresponding pronunciation can be used. If a word is not in the lexicon, G2P can be used to predict the pronunciation.

[0071] Injecting biasing context, including pronunciation knowledge, into a neural network model results in a 14% relative improvement than if no biasing context is injected. The contextually biased model retains the advantages of neural network models such as simple, unified training and implicit learning of pronunciation of rare words. At the same time, the context biased-model incorporates knowledge of rare word pronunciation even if they have never been present during training. However, if the context contains similarly sounding phrases (e.g., when both “Joan” and “John” are in the user’s contact list), disambiguation of the

correct phrase remains challenging. Disambiguation can be improved by training an ASR context mechanism with difficult negative examples.

[0072] In the above training example, training bias phrases are created by randomly sampling n-grams from a reference transcript and other utterances in the training data. At test time, the bias phrases are fixed in advance and are from the same semantic category (e.g., contact names). In this case, the test task is harder than the training task, since distinguishing between first names such as Joan and John is usually more challenging than distinguishing between random unrelated n-grams.

[0073] FIG. 4 illustrates an example method 400 that can be used to close the train-test discrepancy that results from training with randomly sampled n-grams. The method 400 detects proper nouns from a reference transcript instead of sampling random n-grams, and augments them with phonetically similar, or “fuzzy,” alternatives, to create sets of bias phrases for training.

[0074] Method 400 focuses on proper nouns (e.g., unique entities such as names of people and places) in reference transcripts, and uses phonetically similar phrases as negative examples, encouraging the neural network model to learn more discriminative representations. This approach can be applied to a neural network contextual ASR model that jointly learns to transcribe and select the correct context items in order to improve WER by up to 53.1%.

[0075] In ASR, two phrases can be very similar to each other phonetically (i.e., in the way they sound), but be unquestionably different (e.g., “call Joan” and “call John”). For a neural network ASR model, the learned representations for these names might be very similar, leading the model to predict the wrong one. This problem is especially challenging for ASR models presented with rare and difficult-to-spell words, as the model might not observe these words at all during training, and will thus fail to spell them correctly at test time.

[0076] Training using negative examples is a way to teach the model to distinguish these rare and difficult-to-spell phrases from phonetically similar ones. Using this approach, the model learns more robust representations that help it to perform better at test time. The method for training using negative examples includes: (i) detecting difficult-to-transcribe and rare words in the input utterance as the target of focus; and (ii) training harder on these words by providing the model with difficult negative examples.

[0077] For detection of the phrases to focus on ((i)), proper nouns (also tagged as “NNP”) are the general category of phrases that are rare and usually more difficult to transcribe, while being relatively easy to detect. For training harder on these phrases ((ii)), phonetically similar alternative phrases are extracted and fed to the model as negative examples.

[0078] This approach can be thought of as data augmentation for speech. While data augmentation is usually used in machine learning for generating mutated positive examples, this approach is used for generating mutated negative examples. This approach is applied to the contextually biased LAS model to improve WER.

[0079] In the example of FIG. 4, the utterance “Call Joan’s mobile” is provided to an ASR model. When the example utterance 410 is presented, the method 400 detects proper nouns 415 from the reference transcript 405.

[0080] In this example, the word “Joan” 420 is selected as a bias phrase. “Joan” 420 is easily identified as a proper noun

because it is a name of a person. To identify proper nouns, the reference transcript is analyzed with a part-of-speech tagger.

[0081] The method 400 then adds phonetically similar, or fuzzy, phrases as alternatives 425. In this case, the words “John” and “Jean” 430 are selected as alternative bias phrases. A word n-gram is identified as a fuzzy alternative of a second word n-gram if both are phonetically similar and co-occur often in different decoding beams of the same utterances in training data. The process of identifying fuzzy alternatives is described as follows.

[0082] First, a fuzzy inventory is built that stores for each n-gram, a set of fuzzy n-gram alternatives. The fuzzy inventory is built in an unsupervised manner, by using an external conventional model to decode a large corpus of utterances. Each decoded utterance results in a set of hypotheses. All of the co-occurrences of word n-gram pairs that appear in different hypotheses where the rest of their hypotheses are identical are counted, and each n-gram pair is scored according to that count. For instance, “John Lemon” will get a high score for the target n-gram “John Lennon” if these can be often found in different hypotheses of the same utterances.

[0083] Next, the fuzzy inventory is used during contextually biased LAS training. Specifically, given a word n-gram (representing a bias phrase in our case), fuzzy alternatives are selected from the fuzzy inventory and are sorted by the co-occurrence score. The top ranking candidates are filtered by keeping only those that are phonetically similar to the target n-gram, where phonetic similarity is measured using the Hixon metric. The selection of fuzzy bias alternatives is done as part of the data preparation phase.

[0084] The set of extracted proper nouns with their fuzzy alternatives are fed as bias phrases 435 to the grapheme and phoneme encoders. In this way, the bias phrases 435 are used to train the ASR model on the example in the reference transcript 405. This approach allows the model to train on a difficult task of distinguishing between phonetically similar names.

[0085] From each training example, at most three proper nouns are selected at random and three fuzzy alternative are added for each proper noun. Other random proper nouns are added from the rest of the training data and result in up to 64 bias phrases per example.

[0086] When presented with phonetically similar phrases at test time, the fuzzy model captures the subtle phonetic differences better than the non-fuzzy model. This is expressed by both a more accurate prediction. It also results in more attention on the bias phrase that actually appears in the reference transcript rather than its fuzzy alternatives. Thus, training using fuzzy distractors makes the model discriminate phonetically similar phrases better. Additionally, the fuzzy model attends more sharply and its distribution is much cleaner than the non-fuzzy model, which includes incorrect phrases in its attention.

[0087] Training contextualized neural speech recognition models with difficult negative examples results in improved WER. The core idea is to detect and focus on proper nouns (“NNP”) in the reference transcript, and present the model with phonetically similar (“fuzzy”) phrases as their negative examples. When applied to a speech biasing task, the approach improves WER by up to 53.1%.

[0088] FIG. 5 is a flowchart of an example arrangement of operations for a method 500 of biasing a speech recognition

model **300** toward one or more biasing phrases **116** relevant to a context **202** of a spoken utterance **104**. The method **500** may be executed on data processing hardware **610** (FIG. 6) residing on a user device **106** associated with a user **102** that spoke the utterance **104**. Optionally, the data processing hardware **610** may reside on a remote device (e.g., server of a cloud-based computing environment) in communication with the user device **106**, e.g., over a network. At operation **502**, the method **500** includes receiving, at the data processing hardware **610**, audio data **106** encoding the utterance **104**, and at operation **504**, the method **500** includes obtaining, by the data processing hardware **610**, a set of one or more biasing phrases **116** corresponding to the context **202** of the utterance **104**. Here, each biasing phrase **116** in the set of one or more biasing phrases **116** includes one or more words. As used herein, the one or more “words” of each biasing phrase **116** may include proper nouns such, as without limitation, people names in a contact list **204** of the user, city names, music artist names, album/movie (or other media content type) titles, etc. The context **202** may be determined based on at least one of a current date and/or time of the utterance, one or more applications open on the user device **106**, a location of the user **102**.

[**0089**] At operation **506**, the method **500** includes processing, by the data processing hardware **610**, using the speech recognition model **300**, acoustic features **109** derived from the audio data and grapheme and phoneme data **118**, **120** derived from the set of one or more biasing phrases **116** to generate an output **148** of the speech recognition model **300**. The output **148** may include a grapheme output for each frame of the acoustic features **109**. At operation **508**, the method **500** includes determining, by the data processing hardware **610**, a transcription **146** for the utterance **104** based on the output **148** of the speech recognition model **300**. The transcription **146** may be displayed on a graphical user interface **106** of the user device **106** and/or one or more other devices in communication with the data processing hardware **610**. In addition examples, the transcription **146** is provided as a command to a search engine or software application to perform an operatino.

[**0090**] A software application (i.e., a software resource) may refer to computer software that causes a computing device to perform a task. In some examples, a software application may be referred to as an “application,” an “app,” or a “program.” Example applications include, but are not limited to, system diagnostic applications, system management applications, system maintenance applications, word processing applications, spreadsheet applications, messaging applications, media streaming applications, social networking applications, and gaming applications.

[**0091**] The non-transitory memory may be physical devices used to store programs (e.g., sequences of instructions) or data (e.g., program state information) on a temporary or permanent basis for use by a computing device. The non-transitory memory may be volatile and/or non-volatile addressable semiconductor memory. Examples of non-volatile memory include, but are not limited to, flash memory and read-only memory (ROM)/programmable read-only memory (PROM)/erasable programmable read-only memory (EPROM)/electronically erasable programmable read-only memory (EEPROM) (e.g., typically used for firmware, such as boot programs). Examples of volatile memory include, but are not limited to, random access memory (RAM), dynamic random access memory (DRAM), static

random access memory (SRAM), phase change memory (PCM) as well as disks or tapes.

[**0092**] FIG. 6 is schematic view of an example computing device **600** that may be used to implement the systems and methods described in this document. The computing device **600** is intended to represent various forms of digital computers, such as laptops, desktops, workstations, personal digital assistants, servers, blade servers, mainframes, and other appropriate computers. The components shown here, their connections and relationships, and their functions, are meant to be exemplary only, and are not meant to limit implementations of the inventions described and/or claimed in this document.

[**0093**] The computing device **600** includes a processor **610**, memory **620**, a storage device **630**, a high-speed interface/controller **640** connecting to the memory **620** and high-speed expansion ports **650**, and a low speed interface/controller **660** connecting to a low speed bus **670** and a storage device **630**. Each of the components **610**, **620**, **630**, **640**, **650**, and **660**, are interconnected using various busses, and may be mounted on a common motherboard or in other manners as appropriate. The processor **610** can process instructions for execution within the computing device **600**, including instructions stored in the memory **620** or on the storage device **630** to display graphical information for a graphical user interface (GUI) on an external input/output device, such as display **680** coupled to high speed interface **640**. In other implementations, multiple processors and/or multiple buses may be used, as appropriate, along with multiple memories and types of memory. Also, multiple computing devices **600** may be connected, with each device providing portions of the necessary operations (e.g., as a server bank, a group of blade servers, or a multi-processor system).

[**0094**] The memory **620** stores information non-transitorily within the computing device **600**. The memory **620** may be a computer-readable medium, a volatile memory unit(s), or non-volatile memory unit(s). The non-transitory memory **620** may be physical devices used to store programs (e.g., sequences of instructions) or data (e.g., program state information) on a temporary or permanent basis for use by the computing device **600**. Examples of non-volatile memory include, but are not limited to, flash memory and read-only memory (ROM)/programmable read-only memory (PROM)/erasable programmable read-only memory (EPROM)/electronically erasable programmable read-only memory (EEPROM) (e.g., typically used for firmware, such as boot programs). Examples of volatile memory include, but are not limited to, random access memory (RAM), dynamic random access memory (DRAM), static random access memory (SRAM), phase change memory (PCM) as well as disks or tapes.

[**0095**] The storage device **630** is capable of providing mass storage for the computing device **600**. In some implementations, the storage device **630** is a computer-readable medium. In various different implementations, the storage device **630** may be a floppy disk device, a hard disk device, an optical disk device, or a tape device, a flash memory or other similar solid state memory device, or an array of devices, including devices in a storage area network or other configurations. In additional implementations, a computer program product is tangibly embodied in an information carrier. The computer program product contains instructions that, when executed, perform one or more methods, such as

those described above. The information carrier is a computer- or machine-readable medium, such as the memory 620, the storage device 630, or memory on processor 610.

[0096] The high speed controller 640 manages bandwidth-intensive operations for the computing device 600, while the low speed controller 660 manages lower bandwidth-intensive operations. Such allocation of duties is exemplary only. In some implementations, the high-speed controller 640 is coupled to the memory 620, the display 680 (e.g., through a graphics processor or accelerator), and to the high-speed expansion ports 650, which may accept various expansion cards (not shown). In some implementations, the low-speed controller 660 is coupled to the storage device 630 and a low-speed expansion port 690. The low-speed expansion port 690, which may include various communication ports (e.g., USB, Bluetooth, Ethernet, wireless Ethernet), may be coupled to one or more input/output devices, such as a keyboard, a pointing device, a scanner, or a networking device such as a switch or router, e.g., through a network adapter.

[0097] The computing device 600 may be implemented in a number of different forms, as shown in the figure. For example, it may be implemented as a standard server 600a or multiple times in a group of such servers 600a, as a laptop computer 600b, or as part of a rack server system 600c.

[0098] Various implementations of the systems and techniques described herein can be realized in digital electronic and/or optical circuitry, integrated circuitry, specially designed ASICs (application specific integrated circuits), computer hardware, firmware, software, and/or combinations thereof. These various implementations can include implementation in one or more computer programs that are executable and/or interpretable on a programmable system including at least one programmable processor, which may be special or general purpose, coupled to receive data and instructions from, and to transmit data and instructions to, a storage system, at least one input device, and at least one output device.

[0099] These computer programs (also known as programs, software, software applications or code) include machine instructions for a programmable processor, and can be implemented in a high-level procedural and/or object-oriented programming language, and/or in assembly/machine language. As used herein, the terms “machine-readable medium” and “computer-readable medium” refer to any computer program product, non-transitory computer readable medium, apparatus and/or device (e.g., magnetic discs, optical disks, memory, Programmable Logic Devices (PLDs)) used to provide machine instructions and/or data to a programmable processor, including a machine-readable medium that receives machine instructions as a machine-readable signal. The term “machine-readable signal” refers to any signal used to provide machine instructions and/or data to a programmable processor.

[0100] The processes and logic flows described in this specification can be performed by one or more programmable processors, also referred to as data processing hardware, executing one or more computer programs to perform functions by operating on input data and generating output. The processes and logic flows can also be performed by special purpose logic circuitry, e.g., an FPGA (field programmable gate array) or an ASIC (application specific integrated circuit). Processors suitable for the execution of a computer program include, by way of example, both general

and special purpose microprocessors, and any one or more processors of any kind of digital computer. Generally, a processor will receive instructions and data from a read only memory or a random access memory or both. The essential elements of a computer are a processor for performing instructions and one or more memory devices for storing instructions and data. Generally, a computer will also include, or be operatively coupled to receive data from or transfer data to, or both, one or more mass storage devices for storing data, e.g., magnetic, magneto optical disks, or optical disks. However, a computer need not have such devices. Computer readable media suitable for storing computer program instructions and data include all forms of non-volatile memory, media and memory devices, including by way of example semiconductor memory devices, e.g., EPROM, EEPROM, and flash memory devices; magnetic disks, e.g., internal hard disks or removable disks; magneto optical disks; and CD ROM and DVD-ROM disks. The processor and the memory can be supplemented by, or incorporated in, special purpose logic circuitry.

[0101] To provide for interaction with a user, one or more aspects of the disclosure can be implemented on a computer having a display device, e.g., a CRT (cathode ray tube), LCD (liquid crystal display) monitor, or touch screen for displaying information to the user and optionally a keyboard and a pointing device, e.g., a mouse or a trackball, by which the user can provide input to the computer. Other kinds of devices can be used to provide interaction with a user as well; for example, feedback provided to the user can be any form of sensory feedback, e.g., visual feedback, auditory feedback, or tactile feedback; and input from the user can be received in any form, including acoustic, speech, or tactile input. In addition, a computer can interact with a user by sending documents to and receiving documents from a device that is used by the user; for example, by sending web pages to a web browser on a user's client device in response to requests received from the web browser.

[0102] A number of implementations have been described. Nevertheless, it will be understood that various modifications may be made without departing from the spirit and scope of the disclosure. Accordingly, other implementations are within the scope of the following claims.

What is claimed is:

1. A method comprising:
 - receiving, at data processing hardware, audio data encoding an utterance;
 - obtaining, by the data processing hardware, a set of one or more biasing phrases corresponding to a context of the utterance, each biasing phrase in the set of one or more biasing phrases comprising one or more words;
 - processing, by the data processing hardware, using a speech recognition model, acoustic features derived from the audio data and grapheme and phoneme data derived from the set of one or more biasing phrases to generate an output of the speech recognition model; and
 - determining, by the data processing hardware, a transcription for the utterance based on the output of the speech recognition model.
2. The method of claim 1, wherein the speech recognition model comprises:
 - a first encoder configured to receive, as input, the acoustic features and generate, as output, audio encodings from the acoustic features;

- a first attention module configured to receive, as input, the audio encodings output from the first encoder and generate, as output, first attention outputs;
- a grapheme encoder configured to receive, as input, grapheme data indicating graphemes of each word in the set of one or more biasing phrases and generate, as output, grapheme encodings;
- a phoneme encoder configured to receive, as input, phoneme data indicating phonemes of each word in the set of one or more biasing phrases and generate, as output, phoneme encodings;
- a second attention module configured to receive, as input, a representation of the grapheme encodings output from the grapheme encoder and the phoneme encodings output from the phoneme encoder and generate, as output, second attention outputs; and
- a decoder configured to determine likelihoods of sequences of speech elements based on the first attention outputs and the second attention outputs.
3. The method of claim 2, wherein, for each particular word of each biasing phrase in the set of one or more biasing phrases:
- the grapheme encoder is configured to generate a corresponding grapheme encoding for the particular word;
 - the phoneme encoder is configured to generate a corresponding phoneme encoding for the particular word; and
 - the second attention module is configured to encode a corresponding second attention output that comprises a corresponding contextual biasing vector for the particular word based on the corresponding grapheme and phoneme encodings for the particular word.
4. The method of claim 2, wherein the representation of the grapheme encodings output from the grapheme encoder and the phoneme encodings output from the phoneme encoder comprises a projection vector representing a concatenation between the grapheme encodings and the phoneme encodings.
5. The method of claim 2, wherein the first encoder, the first attention module, the grapheme encoder, the phoneme encoder, the second attention module, and the decoder are trained jointly to predict a sequence of graphemes from a sequence of acoustic feature frames.
6. The method of claim 2, wherein:
- the first attention module is configured to compute attention as a function of a previous hidden state of the decoder and a full sequence of audio encodings output by the first encoder; and
 - the second attention module is configured to compute attention as a function of the previous hidden states of the decoder and a full sequence of projection vectors representing characteristics of both the grapheme and phoneme data derived from the biasing phrases in the set of one or more biasing phrases.
7. The method of claim 2, wherein the second attention module is configured to receive a contextual biasing vector that does not correspond to any of the biasing phrases in the set of one or more biasing phrases, the contextual biasing vector representing an option to not bias the output of the speech recognition model.
8. The method of claim 2, wherein the speech elements comprise graphemes.
9. The method of claim 2, wherein the speech elements comprise words or wordpieces.
10. The method of claim 1, wherein the set of one or more biasing phrases comprises one or more contact names personalized for a particular user.
11. The method of claim 1, wherein the set of one or more biasing phrases comprises one or more calendar events personalized for a particular user.
12. The method of claim 1, further comprising determining, by the data processing hardware, the context of the utterance based on a location of a user that spoke the utterance.
13. The method of claim 1, further comprising determining, by the data processing hardware, the context of the utterance based on one or more applications open on a user device associated with a user that spoke the utterance.
14. The method of claim 1, further comprising determining, by the data processing hardware, the context of the utterance based on a current date and/or time of the utterance.
15. The method of claim 1, wherein the speech recognition model comprises a decoder configured to determine a hidden state and the output of the speech recognition model based on:
- an embedding vector for a previous grapheme output by the speech recognition model;
 - a previous hidden state of the decoder;
 - a first vector output by a first attention module; and
 - a second vector output by a second attention module.
16. A system comprising:
- data processing hardware; and
 - memory hardware in communication with the data processing hardware, the memory hardware storing instructions that when executed on the data processing hardware cause the data processing hardware to perform operations comprising:
 - receiving audio data encoding an utterance;
 - obtaining a set of one or more biasing phrases corresponding to a context of the utterance, each biasing phrase in the set of one or more biasing phrases comprising one or more words;
 - processing, using a speech recognition model, acoustic features derived from the audio data and grapheme and phoneme data derived from the set of one or more biasing phrases to generate an output of the speech recognition model; and
 - determining a transcription for the utterance based on the output of the speech recognition model.
17. The system of claim 16, wherein the speech recognition model comprises:
- a first encoder configured to receive, as input, the acoustic features and generate, as output, audio encodings from the acoustic features;
 - a first attention module configured to receive, as input, the audio encodings output from the first encoder and generate, as output, first attention outputs;
 - a grapheme encoder configured to receive, as input, grapheme data indicating graphemes of each word in the set of one or more biasing phrases and generate, as output, grapheme encodings;
 - a phoneme encoder configured to receive, as input, phoneme data indicating phonemes of each word in the set of one or more biasing phrases and generate, as output, phoneme encodings;
 - a second attention module configured to receive, as input, a representation of the grapheme encodings output

from the grapheme encoder and the phoneme encodings output from the phoneme encoder and generate, as output, second attention outputs; and

a decoder configured to determine likelihoods of sequences of speech elements based on the first attention outputs and the second attention outputs.

18. The system of claim 17, wherein, for each particular word of each biasing phrase in the set of one or more biasing phrases:

the grapheme encoder is configured to generate a corresponding grapheme encoding for the particular word; the phoneme encoder is configured to generate a corresponding phoneme encoding for the particular word; and

the second attention module is configured to encode a corresponding second attention output that comprises a corresponding contextual biasing vector for the particular word based on the corresponding grapheme and phoneme encodings for the particular word.

19. The system of claim 17, wherein the representation of the grapheme encodings output from the grapheme encoder and the phoneme encodings output from the phoneme encoder comprises a projection vector representing a concatenation between the grapheme encodings and the phoneme encodings.

20. The system of claim 17, wherein the first encoder, the first attention module, the grapheme encoder, the phoneme encoder, the second attention module, and the decoder are trained jointly to predict a sequence of graphemes from a sequence of acoustic feature frames.

21. The system of claim 17, wherein:

the first attention module is configured to compute attention as a function of a previous hidden state of the decoder and a full sequence of audio encodings output by the first encoder; and

the second attention module is configured to compute attention as a function of the previous hidden states of the decoder and a full sequence of projection vectors representing characteristics of both the grapheme and

phoneme data derived from the biasing phrases in the set of one or more biasing phrases.

22. The system of claim 17, wherein the second attention module is configured to receive a contextual biasing vector that does not correspond to any of the biasing phrases in the set of one or more biasing phrases, the contextual biasing vector representing an option to not bias the output of the speech recognition model.

23. The system of claim 17, wherein the speech elements comprise graphemes.

24. The system of claim 17, wherein the speech elements comprise words or wordpieces.

25. The system of claim 16, wherein the set of one or more biasing phrases comprises one or more contact names personalized for a particular user.

26. The system of claim 16, wherein the set of one or more biasing phrases comprises one or more calendar events personalized for a particular user.

27. The system of claim 16, wherein the operations further comprise determining the context of the utterance based on a location of a user that spoke the utterance.

28. The system of claim 16, wherein the operations further comprise determining the context of the utterance based on one or more applications open on a user device associated with a user that spoke the utterance.

29. The system of claim 16, wherein the operations further comprise determining the context of the utterance based on a current date and/or time of the utterance.

30. The system of claim 16, wherein the speech recognition model comprises a decoder configured to determine a hidden state and the output of the speech recognition model based on:

an embedding vector for a previous grapheme output by the speech recognition model;

a previous hidden state of the decoder;

a first vector output by a first attention module; and

a second vector output by a second attention module.

* * * * *