



- (51) International Patent Classification:
G10L 15/00 (2013.01)
- (21) International Application Number:
PCT/US2015/065617
- (22) International Filing Date:
14 December 2015 (14.12.2015)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
62/092,251 15 December 2014 (15.12.2014) US
62/093,057 17 December 2014 (17.12.2014) US
14/735,002 9 June 2015 (09.06.2015) US
- (71) Applicant: **BAIDU USA LLC** [US/US]; 1195 Bordeaux Dr., Sunnyvale, California 94089 (US).
- (72) Inventors: **HANNUN, Awni**; 416 Leland Ave., Palo Alto, California 94306 (US). **CASE, Carl**; 270 Parnassus Ave., San Francisco, California 94117 (US). **CASPER, Jared**; 661 Winggate Dr., Sunnyvale, California 94087 (US). **CATANZARO, Bryan**; 10439 N. Stelling Rd., Cupertino, California 95014 (US). **DIAMOS, Gregory**; 6191 Cloverhill Dr., San Jose, California 95120 (US). **ELSEN, Erich**; 2541 Betlo Ave., Mountain View, California 94043 (US). **PRENGER, Ryan**; Apt. 13, 14 Whitmore Pl., Oakland,

California 94611 (US). **SATHEESH, Sanjeev**; Apt. 1, 245 Bernardo Ave., Sunnyvale, California 94086 (US). **SHUBHABRATA, Sengupta**; 2125 Gordon Ave., Menlo Park, California 94025 (US). **COATES, Adam**; 1284 Poplar Ave., Sunnyvale, California 94086 (US). **NG, Andrew**; Apt. 201, 2111 Latham St., Mountain View, California 94040 (US).

(74) Agent: **BAUGH, Bradley**; NORTH WEBER & BAUGH LLP, Baugh, Bradley, 2479 E. Bayshore Rd., Ste. 210, Palo Alto, California 94303 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, 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, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(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,

[Continued on next page]

(54) Title: SYSTEMS AND METHODS FOR SPEECH TRANSCRIPTION

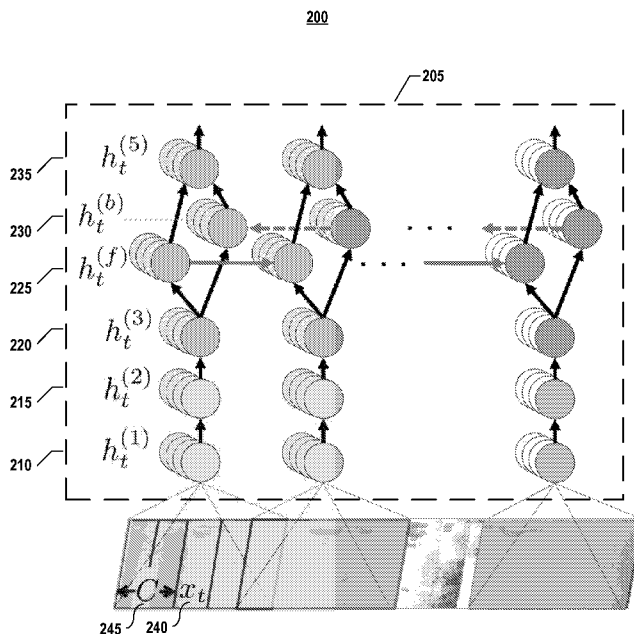


FIG. 2

(57) Abstract: Presented herein are embodiments of state-of-the-art speech recognition systems developed using end-to-end deep learning. In embodiments, the model architecture is significantly simpler than traditional speech systems, which rely on laboriously engineered processing pipelines; these traditional systems also tend to perform poorly when used in noisy environments. In contrast, embodiments of the system do not need hand-designed components to model background noise, reverberation, or speaker variation, but instead directly learn a function that is robust to such effects. A phoneme dictionary, nor even the concept of a "phoneme," is needed. Embodiments include a well-optimized recurrent neural network (RNN) training system that can use multiple GPUs, as well as a set of novel data synthesis techniques that allows for a large amount of varied data for training to be efficiently obtained. Embodiments of the system can also handle challenging noisy environments better than widely used, state-of-the-art commercial speech systems.

WO 2016/100231 A1

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).

Published:

- *with international search report (Art. 21(3))*
- *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))*

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))*

SYSTEMS AND METHODS FOR SPEECH TRANSCRIPTION**Inventors:**

Awni Hannun

Carl Case

5

Jared Casper

Bryan Catanzaro

Gregory Damos

Erich Elsen

Ryan Prenger

10

Sanjeev Satheesh

Shubhabrata Sengupta

Adam Coates

Andrew Y. Ng

CROSS-REFERENCE TO RELATED APPLICATION

15 [0001] This application claims the priority benefit to commonly assigned and co-
pending U.S. Patent Application No. 14/735,002 (Docket No. 28888-1910), filed on 9 June
2015, entitled "SYSTEMS AND METHODS FOR SPEECH TRANSCRIPTION," and
listing Awni Hannun, Carl Case, Jared Casper, Bryan Catanzaro, Gregory Damos, Erich
Elsen, Ryan Prenger, Sanjeev Satheesh, Shubhabrata Sengupta, Adam Coates, and Andrew
20 Y. Ng as inventors, which claimed the priority benefit under 35 USC § 119(e) to commonly
assigned and co-pending: (1) U.S. Patent Application No. 62/092,251 (Docket No. 28888-
1908P), filed on 15 December 2014, entitled "DEEP SPEECH: SCALING UP END-TO-
END SPEECH RECOGNITION," and listing Awni Hannun, Carl Case, Jared Casper, Bryan
Catanzaro, Greg Damos, Erich Elsen, Ryan Prenger, Sanjeev Satheesh, Shubho Sengupta,
25 Adam Coates, and Andrew Y. Ng as inventors; and (2) U.S. Patent Application No.
62/093,057 (Docket No. 28888-1910P), filed on 17 December 2014, entitled "DEEP
SPEECH: SCALING UP END-TO-END SPEECH RECOGNITION," and listing Awni

Hannun, Carl Case, Jared Casper, Bryan Catanzaro, Greg Diamos, Erich Elsen, Ryan Prenger, Sanjeev Satheesh, Shubho Sengupta, Adam Coates, and Andrew Y. Ng as inventors. Each of the aforementioned patent documents is incorporated by reference herein in its entirety.

5 A. Technical Field

[0002] The present disclosure relates to data processing. More particularly, the present disclosure relates to systems and methods for improving the transcription of speech into text.

 B. Description of the Related Art

[0003] Computing devices have become increasingly more prevalent. Computing
10 devices are used in diverse ways, in diverse settings, and appear in diverse form factors. For example, computing devices are used in appliances (such as televisions, refrigerators, and thermostats) and in mobile devices (such as smart phones and tablets), and in wearable devices (such as smart watches).

[0004] The ever-expanding form factors and uses of computing devices have also led to
15 finding better ways to interface with these devices. Interface issues are particularly acute when dealing with computing devices that have limited space or limited input/output abilities.

[0005] Also, regardless of the device's interface capabilities, there is always a need for improved interfaces. Making easier or more natural interfaces can provide a significant competitive advantage. One interface area that has undergone increased attention is the area
20 of speech recognition.

[0006] Top speech recognition systems typically rely on sophisticated pipelines composed of multiple algorithms and hand-engineered processing stages. Traditional speech systems use many heavily engineered processing stages, including specialized input features, acoustic models, and Hidden Markov Models (HMMs). To improve these pipelines, domain

experts generally invest a great deal of effort tuning their features and models. The introduction of deep learning algorithms has improved speech system performance, usually by improving acoustic models. While this improvement has been significant, deep learning still plays only a limited role in traditional speech pipelines. As a result, to improve
5 performance on a task such as recognizing speech in a noisy environment, one must laboriously engineer the rest of the system for robustness.

[0007] Accordingly, what is needed are systems and methods that provide improved speech-to-text conversion.

BRIEF DESCRIPTION OF THE DRAWINGS

[0008] References will be made to embodiments of the invention, examples of which may be illustrated in the accompanying figures. These figures are intended to be illustrative, not limiting. Although the invention is generally described in the context of these
5 embodiments, it should be understood that it is not intended to limit the scope of the invention to these particular embodiments.

[0009] Figure 1 (“FIG. 1”) depicts methods for training a model or set of models according to embodiments of the present invention.

[0010] FIG. 2 graphically represents a neural network model according to embodiments
10 of the present invention.

[0011] FIG. 3 depicts a method for using jitter samples in training a model according to embodiments of the present invention.

[0012] FIG. 4 depicts methods for ensembling a set of neural network models according to embodiments of the present invention.

15 [0013] FIG. 5 depicts a method for constraining output of a model using a language model according to embodiments of the present invention.

[0014] FIG. 6 depicts methods for improving performance of a model according to embodiments of the present invention.

[0015] FIGS. 7A & 7B depict methods for data parallelization according to
20 embodiments of the present invention.

[0016] FIG. 8 depicts methods for model parallelization according to embodiments of the present invention.

[0017] FIG. 9 depicts a method for striding data according to embodiments of the present invention.

[0018] FIG. 10 depicts a method for collecting a set of noise audio clips according to embodiments of the present invention.

[0019] FIG. 11 depicts a method for generating synthesized noise samples according to embodiments of the present invention.

5 [0020] FIG. 12 depicts a method for generating Lombard-effect audio samples according to embodiments of the present invention.

[0021] FIG. 13 depicts methods for using a trained model or set of trained models according to embodiments of the present invention.

10 [0022] FIG. 14 depicts a simplified block diagram of a computing system according to embodiments of the present invention.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

[0023] In the following description, for purposes of explanation, specific details are set forth in order to provide an understanding of the invention. It will be apparent, however, to one skilled in the art that the invention can be practiced without these details. Furthermore, one skilled in the art will recognize that embodiments of the present invention, described below, may be implemented in a variety of ways, such as a process, an apparatus, a system, a device, or a method on a tangible computer-readable medium.

[0024] Components, or modules, shown in diagrams are illustrative of exemplary embodiments of the invention and are meant to avoid obscuring the invention. It shall also be understood that throughout this discussion that components may be described as separate functional units, which may comprise sub-units, but those skilled in the art will recognize that various components, or portions thereof, may be divided into separate components or may be integrated together, including integrated within a single system or component. It should be noted that functions or operations discussed herein may be implemented as components. Components may be implemented in software, hardware, or a combination thereof.

[0025] Furthermore, connections between components or systems within the figures are not intended to be limited to direct connections. Rather, data between these components may be modified, re-formatted, or otherwise changed by intermediary components. Also, additional or fewer connections may be used. It shall also be noted that the terms “coupled,” “connected,” or “communicatively coupled” shall be understood to include direct connections, indirect connections through one or more intermediary devices, and wireless connections.

[0026] Reference in the specification to “one embodiment,” “preferred embodiment,” “an embodiment,” or “embodiments” means that a particular feature, structure, characteristic,

or function described in connection with the embodiment is included in at least one embodiment of the invention and may be in more than one embodiment. Also, the appearances of the above-noted phrases in various places in the specification are not necessarily all referring to the same embodiment or embodiments. Furthermore, the use of
5 certain terms in various places in the specification is for illustration and should not be construed as limiting. Any headings used herein are for organizational purposes only and shall not be used to limit the scope of the description or the claims.

[0027] Furthermore, it shall be noted that: (1) certain steps may optionally be performed; (2) steps may not be limited to the specific order set forth herein; (3) certain steps
10 may be performed in different orders; and (4) certain steps may be done concurrently.

A. Introduction

[0028] Presented herein are embodiments of an end-to-end speech systems, which may be referred to herein as “Deep Speech,” where deep learning supersedes the multiple algorithms and hand-engineered processing stages of prior approaches. In embodiments, this
15 approach, combined with a language model, achieves higher performance than traditional methods on hard speech recognition tasks while also being much simpler. In embodiments, these results may be achieved by training a large recurrent neural network (RNN) using multiple graphics processing units (GPUs) and thousands of hours of data. Because
embodiments of this system learn directly from data, specialized components for speaker
20 adaptation or noise filtering are not required. In fact, in settings where robustness to speaker variation and noise are critical, embodiments of the present invention excels: as discussed further herein, Deep Speech outperforms previously published methods on the Switchboard Hub5’00 corpus, achieving 16.5% error, and performs better than commercial systems in
noisy speech recognition tests.

[0029] It should be noted that tapping the benefits of end-to-end deep learning, however, poses several challenges: (i) innovative ways were found to build large, labeled training sets, and (ii) ways were found to train networks that were large enough to effectively utilize all of this data. One challenge for handling labeled data in speech systems is finding the alignment of text transcripts with input speech. This problem has been previously
5 addressed, thus enabling neural networks to easily consume unaligned, transcribed audio during training. Meanwhile, approaches for rapid training of large neural networks have been suggested, demonstrating the speed advantages of multi-GPU computation. In embodiments, these insights are leveraged to fulfill the vision of a generic learning system, based on large
10 speech datasets and scalable RNN training, that can surpass more complicated traditional methods. This vision is inspired partly by the work that applied early unsupervised feature learning techniques to replace hand-built speech features.

[0030] In embodiments, a recurrent neural network (RNN) model was developed specifically to map well to GPUs. And, in embodiments, a novel model partition scheme may
15 be employed to improve parallelization. Additionally, in embodiments, processes were developed for assembling large quantities of labeled speech data exhibiting the distortions that systems should learn to handle. Using a combination of collected and synthesized data, embodiments of systems disclosed herein learned robustness to realistic noise and speaker variation (including Lombard Effect). These ideas built into embodiments of an end-to-end
20 speech system yielded a system that is at once simpler than traditional pipelines yet also performs better on difficult speech tasks. An embodiment of the Deep Speech system achieved an error rate of 16.5% on the full Switchboard Hub5'00 test set—the best published result. Further, on a new noisy speech recognition dataset of the inventors' own construction, an embodiment of the Deep Speech system achieved a word error rate of 19.1%, where the
25 best commercial systems achieved 30.5% error.

[0031] Presented herein are embodiments of novel speech recognition systems and novel methods for training and for using the systems. Presented below are, first, embodiments of a recurrent neural network model and training framework, followed by a discussion of embodiments of optimizations, and embodiments of data capture and synthesis. Finally, usage and experimental results are presented that demonstrate the state-of-the-art performance of embodiments of a Deep Speech system.

B. Recurrent Neural Network (RNN) Training

[0032] In embodiments, a recurrent neural network (RNN) is trained to ingest speech spectrograms and generate English text transcriptions. Let a single utterance x and label y be sampled from a training set:

$$[0033] \quad \mathcal{X} = \{(x^{(1)}, y^{(1)}), (x^{(2)}, y^{(2)}), \dots\}.$$

[0034] Each utterance, $x^{(i)}$, is a time-series of length $T^{(i)}$ where every time-slice is a vector of audio features, $x_t^{(i)}$, $t = 1, \dots, T^{(i)}$. In embodiments, spectrograms may be used as features, so $x_{t,p}^{(i)}$ denotes the power of the p 'th frequency bin in the audio frame at time t . In embodiments, a goal of the RNN is to convert an input sequence x into a sequence of character probabilities for the transcription y , with $\hat{y}_t = \mathbb{P}(c_t|x)$, wherein $c_t \in \{\text{a, b, c, } \dots, \text{z, space, apostrophe, blank}\}$.

[0035] Figure 1 depicts methods for training a model or set of models according to embodiments of the present invention. In embodiments, an utterance, x , that comprises a time-series of spectrogram frames, $x_{(t)}$, is inputted (105) into a first layer of a recurrent neural network (RNN) model that evaluates at least some of the spectrogram frames with a context of C frames, wherein the utterance, x , has an associated ground-truth, or known, label, y .

[0036] In embodiments, the RNN model may comprise five (5) layers of hidden units. For an input x , the hidden units at layer l are denoted $h^{(l)}$ with the convention that $h^{(0)}$ is the

input. In embodiments, the first three layers are not recurrent. For the first layer, at each time t , the output depends on the spectrogram frame x_t along with a context of C frames, which may be a context from one side or both sides. Thus, in embodiments, a context may comprise one or more spectrogram frames which occur before a spectrogram frame, after the
 5 frame, or both. It shall be noted that, in embodiments, if the spectrogram frame x_t is at the edge of the clip, the context may be from one side or may have data added to the missing side. In embodiments, the context of frames may be $C \in \{5, 7, 9\}$, although other numbers of frames may be used.

[0037] In embodiments, the RNN model calculates (110) through a plurality of hidden
 10 layers. For example, in embodiments, the remaining non-recurrent layers operate on independent data for each time step. Thus, for each time t , the first three (3) layers are computed by:

$$[0038] \quad h_t^{(l)} = g(W^{(l)}h_t^{(l-1)} + b^{(l)})$$

[0039] where $g(z) = \min\{\max\{0, z\}, 20\}$ is the clipped rectified-linear (ReLU)
 15 activation function and $W^{(l)}, b^{(l)}$ are the weight matrix and bias parameters for layer l , respectively. It shall be noted that other activation functions may be employed. In embodiments, the ReLU units are clipped to keep the activations in the recurrent layer from exploding; in practice the units rarely saturate at the upper bound.

[0040] In embodiments, the fourth layer is a bi-directional recurrent network. In
 20 embodiments, this layer includes two sets of hidden units: a set with forward recurrence, $h^{(f)}$, and a set with backward recurrence $h^{(b)}$:

$$[0041] \quad h_t^{(f)} = g(W^{(4)}h_t^{(3)} + W_r^{(f)}h_{t-1}^{(f)} + b^{(4)})$$

$$[0042] \quad h_t^{(b)} = g(W^{(4)}h_t^{(3)} + W_r^{(b)}h_{t+1}^{(b)} + b^{(4)})$$

[0043] Note that, in embodiments, $h^{(f)}$ is computed sequentially from $t = 1$ to $t = T^{(i)}$ for the i 'th utterance, while the units $h^{(b)}$ is computed sequentially in reverse from $t = T^{(i)}$ to $t = 1$.

[0044] In embodiments, a fifth (non-recurrent) layer takes both the forward and
5 backward units as inputs:

$$[0045] \quad h_t^{(5)} = g(W^{(5)}h_t^{(4)} + b^{(5)}), \text{ where } h_t^{(4)} = h_t^{(f)} + h_t^{(b)}.$$

[0046] In embodiments, the output layer is a standard softmax function that yields (115) the predicted character probabilities for each time slice t and character k in the alphabet:

$$[0047] \quad h_{t,k}^{(6)} = \hat{y}_{t,k} \equiv \mathbb{P}(c_t = k|x) = \frac{\exp(W_k^{(6)}h_t^{(5)} + b_k^{(6)})}{\sum_j \exp(W_j^{(6)}h_t^{(5)} + b_j^{(6)})}.$$

10 [0048] Here $W_k^{(6)}$ and $b_k^{(6)}$ denote the k 'th column of the weight matrix and k 'th bias, respectively.

[0049] Once a prediction for $\mathbb{P}(c_t|x)$ has been computed (115), a loss is calculated (120) to measure the error in prediction. In embodiments, a Connectionist Temporal Classification (CTC) loss, $\mathcal{L}(\hat{y}, y)$, may be used to measure the error in prediction, although
15 other error measurement techniques may be employed. During training, the gradient $\nabla_{\hat{y}} \mathcal{L}(\hat{y}, y)$ is evaluated (125) with respect to the network outputs given the ground-truth character sequence y . From this point, computing the gradient with respect to all of the model parameters may be done via back-propagation (130) through the rest of the network. In embodiments, Nesterov's Accelerated gradient method may be used for training, although
20 other techniques may be employed. In embodiments, a momentum of 0.99 may be used and anneal the learning rate by a constant factor, chosen to yield the fastest convergence, after each epoch through the data.

[0050] An embodiment of the complete RNN model is illustrated in Figure 2. As shown in the graphical representation of an embodiment of the RNN model, the model 205 comprises five (5) layers. The first three layers (first layer 210, second layer 215, and third layer 220) are non-recurrent layers, in which the input to the first layer 210 is the spectrogram frames x_i 240 along with a context of C frames (e.g., C 245). The fourth layer 225 & 230 is a bi-directional recurrent network that comprises a set of hidden units with forward recurrence, $h^{(f)}$ 225 and a set of hidden units with backward recurrence, $h^{(b)}$ 230. Finally, as shown in Figure 2, the fifth layer 235 is a non-recurrent layer that receives as inputs the outputs of both the forward and backward units of the fourth layer 225 & 230, and outputs predicted character probabilities.

[0051] It shall be noted that the structure depicted in Figure 2 is considerably simpler than other models—the depicted embodiment has been limited to a single recurrent layer (which is the hardest to parallelize), and the model does not use Long-Short-Term-Memory (LSTM) circuits.

[0052] One disadvantage of LSTM cells is that they require computing and storing multiple gating neuron responses at each step. Since the forward and backward recurrences are sequential, this small additional cost, in embodiments, may become a computational bottleneck. By using a homogeneous model in embodiments, the computation of the recurrent activations are made efficient: computing the ReLu outputs involves only a few highly optimized Basic Linear Algebra Subprograms (BLAS) operations on the GPU and a single point-wise nonlinearity.

1. Regularization

[0053] Embodiments of the models were trained using expanded datasets (embodiments of ways to generating datasets and expand datasets are discussed herein in Section 4), yet

embodiments of the recurrent networks used herein are still adept at fitting the training data. In embodiments, to reduce variance further, several techniques may be employed. In embodiments, during training, a dropout rate (e.g., 5%) was applied. In embodiments, dropout was applied in the feed-forward layers but not to the recurrent hidden activations.

5 [0054] A commonly employed technique in computer vision during network evaluation is to randomly jitter inputs by translations or reflections, feed each jittered version through the network, and vote or average the results. This is not common in speech recognition, however; it was found that translating the raw audio files by 5 milliseconds (ms) (which represented half the filter bank step size used in embodiments herein) to the left and right,
10 forward propagating the recomputed features, and averaging the results beneficial.

[0055] Accordingly, embodiments of the present invention include generating and using jitter sets of audio data. Figure 3 depicts a method for using jitter samples in training a model according to embodiments of the present invention. In embodiments, a jitter set of audio files for the audio file are generated (305) by translating an audio file by one or more
15 time values. For example, the audio file may be moved by a few milliseconds (e.g., 5 ms) forward and back in time. The jitter set of audio files, including the corresponding original audio file, are converted (310) into a set of spectrograms, which are inputted into a model or set of models. In embodiments, the output results from the model or set of models for the set of spectrograms are obtained (315). The output results for the set of spectrograms are
20 blended (320) to obtain an output for the corresponding audio file. In embodiments, the blending may be done by averaging (e.g., taking a mean, median, or mode), weighted average, or voting.

[0056] In embodiments, at test time, an ensemble of several RNNs, averaging probabilities prior to decoding, may also be used. Using a set of models can achieve better
25 results; however, there can be time shifts between the models. Figure 4 depicts methods for

ensembling a set of neural network models according to embodiments of the present invention. As shown in Figure 4, issues of time shifts between neural network models when ensembling outputs of a set of neural network models may be addressed using one or more of the following approaches comprising: (a) using (405) neural network models that exhibit the same temporal shift; (b) shifting (410) the inputs into the various models to have aligned outputs; and (c) checking (415) alignment between output of models and shifting one or more of the outputs to align the outputs.

[0057] For example, it should be noted that naive averaging of the RNN output probabilities may fail in a subtle way. In embodiments, the CTC loss function is agnostic to small time-shifts of the RNN output, and thus several equally good networks can yield shifted versions of essentially the same output probabilities. In isolation these networks perform well, but averaging the shifted probabilities no longer yields any variance reduction for estimates of $\mathbb{P}(c_t|x)$ and, in fact, usually causes performance to degrade. To resolve this, in embodiments, only models where the onset of the output transcriptions begin at the same time were combined. Though this could be achieved by shifting the outputs manually, enough networks have been trained that a collection with matching shifts is available for experiments.

2. Language Model

[0058] When trained from large quantities of labeled speech data, embodiments of the RNN model can learn to produce readable character-level transcriptions. Indeed for many of the transcriptions, the most likely character sequence predicted by embodiments of a trained RNN model is exactly correct without external language constraints. Errors made by the RNN tend to be phonetically plausible renderings of English words—Table 1 shows some examples.

[0059] **Table 1:** Examples of transcriptions directly from an embodiment of a RNN (left) with errors that are fixed by addition of a language model (right).

RNN Output	Decoded Transcription
what is the weather like in bostin right now	what is the weather like in boston right now
prime miniter nerenr modi	prime minister narendra modi
arther n tickets for the game	are there any tickets for the game

[0060] Many of the errors occur on words that rarely or never appear in the training set. In practice, this is hard to avoid: training from enough speech data to hear all of the words or language constructions that the model might need to know is impractical. Therefore, embodiments of the model include, or are integrated with, a language model. In embodiments, an *N*-gram language model is used since these models are easily trained from huge unlabeled text corpora. For comparison, while our speech datasets typically include up to 3 million utterances, the *N*-gram language model used for the experiments in Section 6 is trained from a corpus of 220 million phrases, supporting a vocabulary of 495,000 words.

[0061] Figure 5 depicts a method for constraining output of model using a language model according to embodiments of the present invention. In embodiments, predicted character outputs, $\mathbb{P}(c|x)$, for an utterance, *x*, are obtained (505). Given the output $\mathbb{P}(c|x)$ of the RNN model or models, a search is performed (510) to find the sequence of characters *c*₁, *c*₂, ... that is most probable according to both the RNN output and the language model (where the language model interprets the string of characters as words). In embodiments, the aim is to find a sequence *c* that maximizes the combined objective:

$$Q(c) = \log(\mathbb{P}(c|x)) + \alpha \log(\mathbb{P}_{lm}(c)) + \beta \text{word_count}(c)$$

[0063] where α and β are tunable parameters (set by cross-validation) that control the trade-off between the RNN, the language model constraint, and the length of the sentence.

The term \mathbb{P}_{tm} denotes the probability of the sequence c according to the N -gram model. In embodiments, this objective is maximized using a beam search algorithm, with a typical beam size in the range 1000–8000.

3. Optimizations

5 **[0064]** In embodiments, one or more optimization may be employed to speed-up the operation of the model. Figure 6 depicts methods for improving performance of a model according to embodiments of the present invention. In embodiments, one or more optimizations may be implemented (605) comprising: (a) one or more data parallelisms; (b) model parallelism; and (c) striding the input into the first layer of the model.

10 **[0065]** As noted previously, embodiments include several design decisions to make the networks amenable to high speed execution, and thus fast training. For example, using homogeneous rectified linear networks that are simple to implement and depend on just a few highly-optimized BLAS calls helps the execution speed. When fully unrolled, embodiment of the networks include almost 5 billion connections for a typical utterance; and thus, efficient computation is important. In embodiments, multi-GPU training may be used but
15 doing this effectively may require some additional work, as explained below.

a) Data parallelism

[0066] In embodiments, to help process data efficiently, two levels of data parallelism may be used in which sets of data may be concurrently processed. Figures 7A & 7B depict
20 methods for data parallelization according to embodiments of the present invention.

[0067] First, as depicted in Figure 7A, each GPU may process (705) many examples in parallel. In embodiments, this may be done by concatenating many examples into a single matrix. For instance, rather than performing a single matrix-vector multiplication $W_r h_t$ in the recurrent layer, many may be done in parallel by computing $W_r H_t$, where $H_t =$

$[h_t^{(i)}, h_t^{(i+1)}, \dots]$ and where $h_t^{(i)}$ corresponds to the i 'th example $x^{(i)}$ at time t . In embodiments, the GPU is most efficient when H_t is relatively wide (e.g., 1000 examples or more), and thus it is preferred to process as many examples on one GPU as possible (e.g., up to the limit of GPU memory).

5 [0068] Second, when it is desired to use larger minibatches, or sets of data, than a single GPU can support on its own, data parallelism across multiple GPUs may be used, as illustrated in Figure 7B.

 [0069] As shown in Figure 7B, copies of the model are used (715) on different GPUs, with each GPU processing a separate minibatch of examples. Then, in embodiments, each
10 GPU combines (720) its computed gradient with its peers during each iteration. In embodiments, typically 2× to 4× data parallelism across GPUs were used for experiments by the inventors.

 [0070] Data parallelism is not easily implemented, however, when utterances have different lengths since they cannot be combined into a single matrix multiplication. In
15 embodiments, this problem was resolved by sorting training examples by length and combining only similarly-sized utterances into minibatches, padding with silence when necessary so that all utterances in a batch have the same length.

b) Model parallelism

 [0071] Data parallelism yields training speed-ups for modest multiples of the minibatch
20 size (e.g., 2 to 4), but may face diminishing returns as batching more examples into a single gradient update fails to improve the training convergence rate. That is, processing 2× as many examples on 2× as many GPUs do not always yield a 2× speedup in training. It can also be inefficient to fix the total minibatch size but spread out the examples to 2× as many GPUs: as the minibatch *within* each GPU shrinks, most operations become memory-

bandwidth limited. To scale further, in embodiments, parallelize by partitioning the model (“model parallelism”) may be employed.

[0072] Some embodiments of the model are challenging to parallelize due to the sequential nature of the recurrent layers. Since the bidirectional layer is comprised of a forward computation and a backward computation that are independent, it is possible to perform the two computations in parallel. Unfortunately, in embodiments, naively splitting the RNN to place $h^{(f)}$ and $h^{(b)}$ on separate GPUs involves significant data transfers when computing $h^{(5)}$, which depends on both $h^{(f)}$ and $h^{(b)}$. Thus, in embodiments, a different partitioning of work may be chosen that requires less communication for the models—the model may be divided in half along the *time* dimension.

[0073] Figure 8 depicts a method for model parallelization according to embodiments of the present invention. In embodiments, all layers except the recurrent layer may be trivially decomposed along the time dimension, with the first half of the time-series, from $t = 1$ to $t = T^{(i)}/2$, assigned to one GPU and the second half to another GPU. In embodiments, when computing the recurrent layer activations, the first GPU begins computing (805) the forward activations $h^{(f)}$, while the second begins computing (805) the backward activations $h^{(b)}$. At the mid-point ($t = T^{(i)}/2$), the two GPUs exchange (810) the intermediate activations, $h_{T/2}^{(f)}$ and $h_{T/2}^{(b)}$ and swap roles. The first GPU then finishes (815) the backward computation of $h^{(b)}$ and the second GPU finishes the forward computation of $h^{(f)}$.

20 c) Striding

[0074] Because the recurrent layers are the hardest to parallelize, embodiments of the present invention may include efforts to reduce the running time of the recurrent layers of embodiments of the RNN by taking “steps” (or strides). Figure 9 depicts a method for striding data according to embodiments of the present invention.

[0075] As shown in Figure 9, processing time may be shortened for the recurrent layers by taking strides of a step size of q time slices (e.g., step size of 2) in the original input so that the unrolled RNN has fewer steps. This approach is similar to a convolutional network with a step-size of 2 in the first layer. In embodiments, the cuDNN library, which is a set of optimized low-level primitives to boost the processing speed of deep neural networks (DNN) on CUDA®-compatible GPUs made by NVIDIA of Santa Clara, California, may be used to implement this first layer of convolution efficiently.

4. Training Data

[0076] Large-scale deep learning systems require an abundance of labeled data. For training embodiments of the system, many recorded utterances and corresponding English transcriptions are needed, but there are few public datasets of sufficient scale. To train embodiments of the models, an extensive dataset consisting of 5000 hours of read speech from 9600 speakers was collected. For comparison, Table 2 summarizes the labeled datasets that are available.

[0077] **Table 2:** A summary of the datasets used to train embodiment of Deep Speech. The Wall Street Journal, Switchboard, and Fisher corpora (see, C. Cieri, D. Miller, and K. Walker, “The Fisher corpus: a Resource for the Next Generations of Speech-to-Text,” in **Language Resources and Evaluation Conference (LREC)**, volume 4, pages 69–71, 2004, which is incorporated by reference herein in its entirety) are all published by the Linguistic Data Consortium.

Dataset	Type	Hours	Speakers
WSJ	read	80	280
Switchboard	conversational	300	4000
Fisher	conversational	2000	23000

Baidu	read	5000	9600
-------	------	------	------

[0078] In embodiments, speech data was recorded from paid volunteers through Amazon Mechanical Turk. Each volunteer (“Turker”) was instructed to read a short text prompt into a computer microphone. The recorded speech and matching text transcription were added to a dataset on the server. Since read speech created in this way sounds different from spontaneous speech heard in other contexts, Turkers were instructed to speak as naturally as possible. A separate team of trusted Turkers was occasionally tasked with quality control, verifying that recorded utterances and transcriptions matched. A cap on contributions of volunteers was enforced to ensure that the dataset contained a variety of speakers and was not overwhelmed by a few big contributors. In embodiments, noise was also played through the Turker’s computer speakers or headset to induce the Lombard effect, which will be discussed in more detail below.

a) Synthesis by superposition

[0079] In embodiments, to expand the potential training data even further, data synthesis was used. In embodiments, a goal was to improve performance in noisy environments where existing systems breakdown. Capturing labeled data (e.g., read speech) from noisy environments is not practical, however, and thus other ways, which are disclosed herein, were found to generate such data.

[0080] To a first order, audio signals are generated through a process of superposition of source signals. This fact was used to synthesize noisy training data. For example, given a speech audio track $x^{(i)}$ and a “noise” audio track $\xi^{(i)}$, then a “noisy speech” track may be formed by $\hat{x}^{(i)} = x^{(i)} + \xi^{(i)}$ to simulate audio captured in a noisy environment. In

embodiments, reverberations, echoes, or other forms of damping may be added to the power spectrum of $x^{(i)}$, $\xi^{(i)}$, or both and then added together to make fairly realistic audio scenes.

[0081] It should be noted that there are, however, some risks in this approach. For example, in order to take 1000 hours clean speech and create 1000 hours noisy speech, 5 unique noise tracks spanning roughly 1000 hours would normally be needed. It is not preferred to settle for, say, 10 hours of repeating noise, since it may become possible for the recurrent network to memorize the noise track and “subtract” it out of the synthesized data. Thus, instead of using a single noise source $\xi^{(i)}$ with a length of 1000 hours, in embodiments, a large number of shorter clips (which are easier to collect from public video sources) may 10 instead be used and treat them as separate sources of noise before superimposing of them:

$$[0082] \quad \hat{x}^{(i)} = x^{(i)} + \xi_1^{(i)} + \xi_2^{(i)} + \dots$$

[0083] Figure 10 depicts a method for collecting a set of noise audio clips according to embodiments of the present invention. In embodiments, a set of candidate “noise” clips are collected (1005), which may be collected from any of a number of sources, such as the 15 Internet or other datasets.

[0084] When superimposing many signals collected from audio or video clips, one can end up with “noise” sounds that are different from the kinds of noise recorded in real environments. To ensure a good match between synthetic data and real data, any candidate noise clips where the average power in each frequency band differed significantly from the 20 average power observed in real noisy recordings may be rejected. Thus, in embodiments, a candidate noise clip from the set of candidate noise clips may be selected (1010) for evaluation. A check is made (1015) to determine whether the candidate noise clip’s average power in frequency bands differed significantly from average powers observed in real noise recordings. If the candidate noise clip’s average powers in frequency bands differ, it is

discarded (1030). If the candidate noise clip's average powers in frequency bands do not differ, it is added (1020) to a set of approved noise clips.

[0085] In embodiments, a determination is made (1025) whether there are more candidate noise clips to evaluate. If there are one or more candidate noise clips to evaluate, the process (steps 1010 – 1020) may be repeated. However, if there are not any more
5 candidate noise clips to evaluate, the process ends (1035).

[0086] Given the set of approved noise clips, a set of synthesized noise may be generated. Figure 11 depicts a method for generating synthesized noise samples according to embodiments of the present invention. In embodiments, one or more noise clips selected
10 from the set of approved noise clips may be added (1105) together to form a noise track. In embodiments, the signal-to-noise ratio of the noise track is adjusted (1110) relative to an audio sample. The adjusted noise track is added (1115) to the audio sample to form a synthesized noise sample, and that synthesized noise sample is included (1120) in a set of training data. In embodiments, this process may be repeated (1125) until a satisfactory
15 number of samples have been made, at which point the process ends (1130).

b) Capturing Lombard Effect

[0087] One challenging effect encountered by speech recognition systems in noisy environments is the “Lombard Effect,” in which speakers actively change the pitch or inflections of their voice to overcome noise around them. This (involuntary) effect does not
20 show up in recorded speech datasets since they are collected in quiet environments. To ensure that the effect is represented in training data used in embodiments herein, the Lombard effect was intentionally induced during data collection. Figure 12 depicts a method for generating Lombard-Effect audio samples according to embodiments of the present invention.

computed over windows of p milliseconds (e.g., 20 ms) strided by q milliseconds (e.g., 10 ms).

[0092] In embodiments, each of the spectrogram frame, $x_{(t)}$, from the set of spectrogram frames is inputted (1325) into the trained neural network or networks. In embodiments, a spectrogram frame is input (1325) along with a context of C frames (e.g., 9 frames). The context of frames may be from one side of the spectrogram frame or from both sides.

[0093] In embodiments, a plurality of trained neural networks are used and the predicted character probabilities outputs from the set of trained neural networks for the inputted spectrogram frames may be combined (1330), or ensembled, together. Ensembling may be performed using one or more of the methods described above with respect to Figure 4—although other methods may also be employed.

[0094] Finally, in embodiments, a transcription of the input audio is obtained by decoding (1335) the ensemble predicted character probabilities outputs. In embodiments, the decoding processing may comprise using a trained N -gram language model to constrain, at a word level, the decoding of the ensemble predicted character probabilities outputs.

6. Example Experiments

[0095] Two sets of experiments were performed to evaluate embodiments of the model system. In both cases, an embodiment of the model as described in Section B was trained from a selection of the datasets in Table 2 to predict character-level transcriptions. The predicted probability vectors and language model were then fed into a decoder to yield a word-level transcription, which was compared with the ground truth transcription to yield the word error rate (WER). It shall be noted that these experiments and results are provided by way of illustration and were performed under specific conditions using an embodiment of the

present invention; accordingly, neither these experiments nor their results shall be used to limit the scope of the present invention.

a) Conversational speech: Switchboard Hub5'00 (full)

[0096] To compare an embodiment system of the present invention (which may be referred to herein as the Deep Speech embodiment, Deep Speech model, Deep Speech network, Deep Speech system, or simply, Deep Speech) to prior research, an accepted but highly challenging test set, Hub5'00 (LDC2002S23), was used. Some researchers split this set into "easy" (Switchboard) and "hard" (CallHome) instances, often reporting new results on the easier portion alone. In the experiments presented herein, the full set was used, which is the most challenging case and report the overall word error rate.

[0097] The Deep Speech system was trained on only the 300 hour Switchboard conversational telephone speech dataset and trained on both Switchboard (SWB) and Fisher (FSH), a 2000 hour corpus collected in a similar manner as Switchboard. Many researchers evaluate models trained only with 300 hours from Switchboard conversational telephone speech when testing on Hub5'00. In part, this is because training on the full 2000 hour Fisher corpus is computationally difficult. Using the optimization techniques mentioned in Section 3, the Deep Speech system was able perform a full pass over the 2300 hours of data in just a few hours.

[0098] Since the Switchboard and Fisher corpora are distributed at a sample rate of 8kHz, spectrograms of 80 linearly spaced log filter banks and an energy term were computed. The filter banks were computed over windows of 20 ms strided by 10 ms. More sophisticated features, such as the mel-scale log filter banks or the mel-frequency cepstral coefficients, were not evaluated.

[0099] Speaker adaptation can be important to the success of other current automatic speech recognition (ASR) systems, particularly when trained on 300 hour Switchboard. For the models tested on Hub5'00, a simple form of speaker adaptation was applied by normalizing the spectral features on a per speaker basis. Other than this, the input features were not modified in any way.

[00100] For decoding, a 4-gram language model with a 30,000 word vocabulary trained on the Fisher and Switchboard transcriptions was used. Hyperparameters for the decoding objective were chosen via cross-validation on a held-out development set.

[00101] The tested Deep Speech SWB model was a network of 5 hidden layers, each with 2048 neurons trained on only 300 hour switchboard. The Deep Speech SWB + FSH model was an ensemble of 5 RNNs, each with 5 hidden layers of 2304 neurons trained on the full 2300 hour combined corpus. All networks were trained on inputs of +/- 9 frames of context.

[00102] The results are reported in Table 3. The model from Vesely *et al.* (DNN-HMM sMBR) used a sequence-based loss function on top of a DNN after using a typical hybrid DNN-HMM system to realign the training set. The performance of this model on the combined Hub5'00 test set was the best previously published result. When trained on the combined 2300 hours of data, the Deep Speech system improved upon this baseline by 1.9% absolute WER and 10.3% relative. The model from Maas *et al.* (DNN-HMM FSH) achieved 19.9% WER when trained on the Fisher 2000 hour corpus. That system was built using Kaldi, a state-of-the-art open source speech recognition software. This result was included to demonstrate that Deep Speech, when trained on a comparable amount of data, is competitive with the best existing ASR systems.

[00103] **Table 3:** Published error rates (%WER) on Switchboard dataset splits. The columns labeled "SWB" and "CH" are respectively the easy and hard subsets of Hub5'00.

Model	SWB	CH	Full
Vesely <i>et al.</i> (GMM-HMM BMMI) [Ref. A]	18.6	33.0	25.8
Vesely <i>et al.</i> (DNN-HMM sMBR) [Ref. A]	12.6	24.1	18.4
Maas <i>et al.</i> (DNN-HMM SWB) [Ref. B]	14.6	26.3	20.5
Maas <i>et al.</i> (DNN-HMM FSH) [Ref. B]	16.0	23.7	19.9
Seide <i>et al.</i> (CD-DNN) [Ref. C]	16.1	n/a	n/a
Kingsbury <i>et al.</i> (DNN-HMM sMBR HF) [Ref. D]	13.3	n/a	n/a
Sainath <i>et al.</i> (CNN-HMM) [Ref. E]	11.5	n/a	n/a
Deep Speech SWB	20.0	31.8	25.9
Deep Speech SWB + FSH	13.1	19.9	16.5

[00104] **Ref. A:** K. Vesely, A. Ghoshal, L. Burget, and D. Povey, “Sequence-Discriminative Training of Deep Neural Networks,” in **Proc. of Interspeech 2013**, Lyon, France, Aug. 2013, pp. 2345-2349 (which is incorporated by reference herein in its entirety).

5 [00105] **Ref. B:** A. L. Maas, A. Y. Hannun, C. T. Lengerich, P. Qi, D. Jurafsky, and A. Y. Ng, “Increasing Deep Neural Network Acoustic Model Size for Large Vocabulary Continuous Speech Recognition,” abs/1406.7806, 2014 (available at //arxiv.org/abs/1406.7806) (which is incorporated by reference herein in its entirety).

10 [00106] **Ref. C:** F. Seide, G. Li, X. Chen, D. Yu, “Feature Engineering in Context-Dependent Deep Neural Networks for Conversational Speech Transcription,” in **Proc. of IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)** (Waikoloa, HI, USA, 11 December 2011), pp. 24–29 (which is incorporated by reference herein in its entirety).

15 [00107] **Ref. D:** B. Kingsbury, T. N. Sainath, and H. Soltau, “Scalable Minimum Bayes Risk Training of Deep Neural Network Acoustic Models Using Distributed Hessian-free

Optimization,” in **Proc. INTERSPEECH**, September 2012 (which is incorporated by reference herein in its entirety).

[00108] **Ref. E:** T. N. Sainath, A.-R. Mohamed, B. Kingsbury, and B. Ramabhadran, “Deep convolutional neural networks for LVCSR,” in **2013 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)**, 2013 (which is incorporated by
5 reference herein in its entirety).

b) Noisy speech

[00109] Few standards exist for testing noisy speech performance, so an evaluation set of 100 noisy and 100 noise-free utterances from 10 speakers was constructed. The noise
10 environments included a background radio or TV; washing dishes in a sink; a crowded cafeteria; a restaurant; and inside a car driving in the rain. The utterance text came primarily from web search queries and text messages, as well as news clippings, phone conversations, Internet comments, public speeches, and movie scripts. Precise control over the signal-to-noise ratio (SNR) of the noisy samples was not possible, but the SNR ratio target was
15 between 6 and 2 dB.

[00110] For the following experiments, embodiments of the RNNs were trained on all the datasets (more than 7000 hours) listed in Table 2. Since training was for 15 to 20 epochs with newly synthesized noise in each pass, the model learned from over 100,000 hours of novel data. Also, an ensemble of 6 networks, each with 5 hidden layers of 2560 neurons, was
20 used. No form of speaker adaptation was applied to the training or evaluation sets. Training examples were normalized on a per utterance basis in order to make the total power of each example consistent. The features were 160 linearly spaced log filter banks computed over windows of 20 milliseconds strided by 10 milliseconds and an energy term. Audio files were resampled to 16 kHz prior to the featurization. Finally, from each frequency bin the

experimenters removed the global mean over the training set and divided by the global standard deviation, primarily so the inputs are well scaled during the early stages of training.

[00111] As described in Section B.2, a 5-gram language model was used for the decoding. The language model was trained on 220 million phrases of the Common Crawl
 5 (available at commoncrawl.org), selected such that at least 95% of the characters of each phrase were in the alphabet. Only the most common 495,000 words were kept, the rest were remapped to an “UNKNOWN” token.

[00112] The trained embodiment of the Deep Speech system was compared to several commercial speech systems: (1) wit.ai, (2) Google Speech API, (3) Bing Speech, and (4)
 10 Apple Dictation. Note that wit.ai and Google Speech each have HTTP-based APIs; to test Apple Dictation and Bing Speech, a kernel extension was used to loop audio output back to audio input and the OS X Dictation service as well as the Windows 8 Bing speech recognition API.

[00113] The test was designed to benchmark performance in noisy environments. This
 15 situation created challenges for evaluating the web speech APIs: these systems will give no result at all when the SNR is too low or in some cases when the utterance is too long. Therefore, the comparison was restricted to the subset of utterances for which all systems returned a non-empty result. This leads to much higher accuracies than would be reported if 100% error was attributed in cases where an API failed to respond. The results of evaluating
 20 each system on the test files appear in Table 4.

[00114] **Table 4:** Results (%WER) for 3 systems evaluated on the original audio. All systems were scored only on utterances with predictions given by all systems. The number in parenthesis next to each dataset, e.g. Clean (94), is the number of utterances scored.

System	Clean (94)	Noisy (83)	Combined (177)
--------	------------	------------	----------------

Apple Dictation	14.24	43.76	26.73
Bing Speech	11.73	36.12	22.05
Google API	6.64	30.47	16.72
wit.ai	7.94	35.06	19.41
Deep Speech	6.56	19.06	11.85

[00115] To evaluate the efficacy of the noise synthesis techniques described in Section B.4, embodiments of two RNNs were trained, one on 5000 hours of raw data and the other trained on the same 5000 hours plus noise. On the 100 clean utterances both models performed about the same, 9.2% WER and 9.0% WER for the clean trained model and the noise trained model, respectively. However, on the 100 noisy utterances the noisy model achieves 22.6% WER over the clean model’s 28.7% WER, a 6.1% absolute and 21.3% relative improvement.

c) Conclusion

[00116] Several parts of embodiments presented herein expanded upon or were inspired by previous results. Neural network acoustic models and other connectionist approaches were first introduced to speech pipelines in the early 1990s. These systems, similar to deep neural network (DNN) acoustic models, replace only one stage of the speech recognition pipeline. Mechanically, embodiments of the systems presented herein may be considered to include some similarity to other efforts to build end-to-end speech systems from deep learning algorithms. For example, Graves *et al.* (A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, “Connectionist temporal classification: Labelling unsegmented sequence data with recurrent neural networks,” in **Proceedings of the 23rd International Conference on Machine Learning (ICML)**, pp. 369–376. ACM (2006), which is incorporated by reference

herein in its entirety) previously introduced the “Connectionist Temporal Classification” (CTC) loss function for scoring transcriptions produced by RNNs and, with Long Short-Term Memory (LSTM) networks, previously applied this approach to speech (see, A. Graves and N. Jaitly, “Towards end-to-end speech recognition with recurrent neural networks,” in **Proceedings of the 31st International Conference on Machine Learning (ICML)**, pp. 1764–1772, 2014, which is incorporated by reference herein in its entirety). Embodiments presented herein similarly adopt the CTC loss for part of the training procedure but use much simpler recurrent networks with rectified-linear activations. Embodiment of the recurrent network may be considered to include some similarity to the bidirectional RNN used by Hannun *et al.* (A. Y. Hannun, A. L. Maas, D. Jurafsky, and A. Y. Ng, “First-pass large vocabulary continuous speech recognition using bi-directional recurrent DNNs,” in *The Computing Research Repository (CoRR)* abs/1408.2873, 2014 (available at arxiv.org/abs/1408.2873), which is incorporated by reference herein in its entirety), but with multiple changes to enhance its scalability. By focusing on scalability, it is shown herein that simpler networks can be effective even without the more complex LSTM machinery.

[00117] Scalability to improve performance of deep learning (DL) algorithms has been previously considered, and the use of parallel processors (including GPUs) has been used in recent large-scale DL results. Early ports of DL algorithms to GPUs revealed significant speed gains. Researchers have also begun choosing designs that map well to GPU hardware to gain even more efficiency, including convolutional and locally connected networks, especially when optimized libraries like cuDNN (see, S. Chetlur, C. Woolley, P. Vandermersch, J. Cohen, J. Tran, B. Catanzaro, and E. Shelhamer, “cuDNN: Efficient primitives for deep learning,” **CoRR**, abs/1410.0759, 2104 (available at arxiv.org/abs/1410.0759), which is incorporated by reference herein in its entirety) and BLAS are available. Indeed, using high-performance computing infrastructure, it is possible today to

train neural networks with more than 10 billion connections using clusters of GPUs. These results inspired at least some of the embodiments herein to be developed by focusing on making scalable design choices to efficiently utilize many GPUs before trying to engineer the algorithms and models themselves.

5 [00118] With the potential to train large models, there is a need for large training sets as well. In other fields, such as computer vision, large labeled training sets have enabled significant leaps in performance as they are used to feed larger and larger DL systems. In speech recognition, however, such large training sets are not common, with typical benchmarks having training sets ranging from tens of hours (e.g., the Wall Street Journal
10 corpus with 80 hours) to several hundreds of hours (e.g. Switchboard and Broadcast News). Larger benchmark datasets, such as the Fisher corpus with 2000 hours of transcribed speech, are rare and only recently being studied. To fully utilize the expressive power of embodiments of the networks described herein, large sets of labeled utterances were not the only datasets that were relied upon, but also synthesis techniques were used to generate novel
15 examples. A similar approach is known in computer vision, but it was found that, when done properly, synthesizing data can be especially convenient and effective for speech training.

[00119] It shall be noted that the novel disclosures herein provided novel and significant advances over the prior these approaches. Presented herein are novel systems and methods for creating, training, and using end-to-end deep learning-based speech systems capable of
20 outperforming existing state-of-the-art recognition pipelines in at least two challenging scenarios: clear, conversational speech, and speech in noisy environments. Embodiments include multi-GPU training and also include unique data collection and synthesis strategies to build large training sets exhibiting the distortions systems should be able to handle (such as background noise and Lombard Effect). These solutions allow for the building of a data-

driven speech system that is at once better performing than existing methods while no longer relying on the complex processing stages that had stymied further progress.

C. Exemplary Computing System Embodiments

[00120] Aspects of the present patent document are directed to a computing system. For purposes of this disclosure, a computing system may include any instrumentality or aggregate of instrumentalities operable to compute, calculate, determine, classify, process, transmit, receive, retrieve, originate, route, store, display, communicate, manifest, detect, record, reproduce, handle, or utilize any form of information, intelligence, or data for business, scientific, control, or other purposes. For example, a computing may be a personal computer (e.g., desktop or laptop), tablet computer, mobile device (e.g., personal digital assistant (PDA) or smart phone), server (e.g., blade server or rack server), a network device, or any other suitable device and may vary in size, shape, performance, functionality, and price. The computing system may include random access memory (RAM), one or more processing resources such as a central processing unit (CPU) or hardware or software control logic, ROM, and/or other types of memory. Additional components of the computing system may include one or more disk drives, one or more network ports for communicating with external devices as well as various input and output (I/O) devices, such as a keyboard, a mouse, touchscreen and/or a video display. The computing system may also include one or more buses operable to transmit communications between the various hardware components.

[00121] Figure 14 depicts a block diagram of a computing system 1400 according to embodiments of the present invention. It will be understood that the functionalities shown for system 1400 may operate to support various embodiments of a computing system—although it shall be understood that a computing system may be differently configured and include different components. As illustrated in Figure 14, system 1400 includes one or more central

processing units (CPU) 1401 that provides computing resources and controls the computer. CPU 1401 may be implemented with a microprocessor or the like, and may also include one or more graphics processing units (GPU) 1417 and/or a floating point coprocessor for mathematical computations. System 1400 may also include a system memory 1402, which
5 may be in the form of random-access memory (RAM), read-only memory (ROM), or both.

[00122] A number of controllers and peripheral devices may also be provided, as shown in Figure 14. An input controller 1403 represents an interface to various input device(s) 1404, such as a keyboard, mouse, or stylus. There may also be a scanner controller 1405, which communicates with a scanner 1406. System 1400 may also include a storage
10 controller 1407 for interfacing with one or more storage devices 1408 each of which includes a storage medium such as magnetic tape or disk, or an optical medium that might be used to record programs of instructions for operating systems, utilities, and applications, which may include embodiments of programs that implement various aspects of the present invention. Storage device(s) 1408 may also be used to store processed data or data to be processed in
15 accordance with the invention. System 1400 may also include a display controller 1409 for providing an interface to a display device 1411, which may be a cathode ray tube (CRT), a thin film transistor (TFT) display, or other type of display. The computing system 1400 may also include a printer controller 1412 for communicating with a printer 1413. A communications controller 1414 may interface with one or more communication devices
20 1415, which enables system 1400 to connect to remote devices through any of a variety of networks including the Internet, an Ethernet cloud, an FCoE/DCB cloud, a local area network (LAN), a wide area network (WAN), a storage area network (SAN) or through any suitable electromagnetic carrier signals including infrared signals.

[00123] In the illustrated system, all major system components may connect to a bus
25 1416, which may represent more than one physical bus. However, various system

components may or may not be in physical proximity to one another. For example, input data and/or output data may be remotely transmitted from one physical location to another. In addition, programs that implement various aspects of this invention may be accessed from a remote location (e.g., a server) over a network. Such data and/or programs may be conveyed
5 through any of a variety of machine-readable medium including, but are not limited to: magnetic media such as hard disks, floppy disks, and magnetic tape; optical media such as CD-ROMs and holographic devices; magneto-optical media; and hardware devices that are specially configured to store or to store and execute program code, such as application specific integrated circuits (ASICs), programmable logic devices (PLDs), flash memory
10 devices, and ROM and RAM devices.

[00124] Embodiments of the present invention may be encoded upon one or more non-transitory computer-readable media with instructions for one or more processors or processing units to cause steps to be performed. It shall be noted that the one or more non-transitory computer-readable media shall include volatile and non-volatile memory. It shall
15 be noted that alternative implementations are possible, including a hardware implementation or a software/hardware implementation. Hardware-implemented functions may be realized using ASIC(s), programmable arrays, digital signal processing circuitry, or the like. Accordingly, the “means” terms in any claims are intended to cover both software and hardware implementations. Similarly, the term “computer-readable medium or media” as
20 used herein includes software and/or hardware having a program of instructions embodied thereon, or a combination thereof. With these implementation alternatives in mind, it is to be understood that the figures and accompanying description provide the functional information one skilled in the art would require to write program code (i.e., software) and/or to fabricate circuits (i.e., hardware) to perform the processing required.

[00125] It shall be noted that embodiments of the present invention may further relate to computer products with a non-transitory, tangible computer-readable medium that have computer code thereon for performing various computer-implemented operations. The media and computer code may be those specially designed and constructed for the purposes of the present invention, or they may be of the kind known or available to those having skill in the relevant arts. Examples of tangible computer-readable media include, but are not limited to: magnetic media such as hard disks, floppy disks, and magnetic tape; optical media such as CD-ROMs and holographic devices; magneto-optical media; and hardware devices that are specially configured to store or to store and execute program code, such as application specific integrated circuits (ASICs), programmable logic devices (PLDs), flash memory devices, and ROM and RAM devices. Examples of computer code include machine code, such as produced by a compiler, and files containing higher level code that are executed by a computer using an interpreter. Embodiments of the present invention may be implemented in whole or in part as machine-executable instructions that may be in program modules that are executed by a processing device. Examples of program modules include libraries, programs, routines, objects, components, and data structures. In distributed computing environments, program modules may be physically located in settings that are local, remote, or both.

[00126] One skilled in the art will recognize no computing system or programming language is critical to the practice of the present invention. One skilled in the art will also recognize that a number of the elements described above may be physically and/or functionally separated into sub-modules or combined together.

[00127] It shall be noted that elements of the claims, below, may be arranged differently including having multiple dependencies, configurations, and combinations. For example, in embodiments, the subject matter of various claims may be combined with other claims. For example, the subject matter of claim 4 may be combined with the subject matter of any of

claims 1 to 3; the subject matter of claim 6 may be combined with the subject matter of claims 4 or 5; the subject matter of claim 8 may be combined with the subject matter of claims 1 to 7; the subject matter of claim 10 may be combined with the subject matter of claims 8 or 9; the subject matter of claim 13 may be combined with the subject matter of
5 claims 11 or 12; the subject matter of claim 15 may be combined with the subject matter of claims 11 or 14; the subject matter of claim 16 may be combined with the subject matter of claims 11 to 15; and the subject matter of claim 19 may be combined with the subject matter of claims 17 or 18.

[00128] It will be appreciated to those skilled in the art that the preceding examples and
10 embodiments are exemplary and not limiting to the scope of the present invention. It is intended that all permutations, enhancements, equivalents, combinations, and improvements thereto that are apparent to those skilled in the art upon a reading of the specification and a study of the drawings are included within the true spirit and scope of the present invention.

CLAIMS

WHAT IS CLAIMED IS:

1. A computer-implemented method for training a transcription model, the method comprising:

5 for each of a set of utterances:

inputting an utterance that comprises a set of spectrogram frames into a first layer of the transcription model that evaluates each of the spectrogram frames from the set of spectrogram frames with a context of one or more spectrogram frames;

10 outputting from the transcription model a predicted character or character probabilities for the utterance; and

computing a loss to measure error in prediction for the utterance;

evaluating a gradient of predicted outputs of the transcription model given the ground-truth characters; and

15 updating the transcription model using back-propagation.

2. The computer-implemented method of Claim 1 further comprising:

jittering at least some of the set of utterances prior to inputting into the transcription model.

3. The computer-implemented method of Claim 2 wherein the step of jittering at least some
20 of the set of utterances prior to inputting into the transcription model comprises:

generating a jitter set of utterances for an utterance by translating an audio file of the utterance by one or more time values;

converting the jitter set of utterances and the utterance into a set of spectrograms;

obtaining output results from the model or from a set of models for the set of spectrograms; and
averaging the output results for the set of spectrograms to obtain an output for the audio file.

5 4. The computer-implemented method of Claim 3 further comprising:
generating one or more utterances for a set of training data for use in training the transcription model.

5. The computer-implemented method of Claim 4 wherein generating one or more utterances for a dataset for using in training comprises:

10 having a person wear headphones as the person records an utterance;
intentionally inducing a Lombard effect during data collection of the utterance by playing background noise through the headphones worn by the person; and capturing the Lombard-effected utterance of the person via a microphone without capturing the background noise.

15 6. The computer-implemented method of Claim 4 wherein generating one or more utterances for a dataset for using in training comprises:

adding one or more noise clips selected from a set of approved noise clips to form a noise track;
adjusting a signal-to-noise ratio of the noise track relative to an audio file;
20 adding the adjusted noise track to the audio file to form a synthesized noise audio file; and
adding the synthesized noise audio file to the set of training data.

7. The computer-implemented method of Claim 6 wherein the set of approved noise clips are generated by performing the steps comprising:

collecting a set of candidate noise clips; and

repeating the following steps until each noise clip from the set of candidate noise

5 clips has been evaluated:

selecting a candidate noise clip from the set of candidate noise clips for
evaluation;

responsive to the candidate noise clip's average powers in frequency bands not

differing significantly from average powers in frequency bands

10 observed in real noise recordings, adding the candidate noise clip to the

set of approved noise clips; and

responsive to the candidate noise clip's average powers in frequency bands

differing significantly from average powers observed in real noise

recordings, discarding the candidate noise clip.

- 15 8. The computer-implemented method of Claim 1 further comprising incorporating one or more optimizations in the training of the transcription model, comprising:

incorporating one or more data parallelisms;

incorporating a model parallelism; and

striding the input into the first layer of the transcription model.

- 20 9. The computer-implemented method of Claim 8 wherein the step of incorporating one or more data parallelism comprises:

using several copies of the transcription model across multiple processing units

with each processing unit processing a separate minibatch of utterances;

and

combining a computed gradient from a processing unit with its peers during each iteration.

10. The computer-implemented method of Claim 8 wherein the step of incorporating one or more data parallelism comprises:

5 having each processing unit process many utterances in parallel by concatenating many utterances into a single matrix; and
 sorting utterances by length and combining similarly-sized utterances into minibatches and padding utterances with silence so that all utterances in a minibatch have the same length.

10 11. A computer-implemented method for transcribing speech comprising:

 receiving an input audio from a user;
 normalizing the input audio to make a total power of the input audio consistent with a set of training samples used to train a trained neural network model;
 generating a jitter set of audio files from the normalized input audio by translating
15 the normalized input audio by one or more time values;
 for each audio file from the jitter set of audio files, which includes the normalized input audio:
 generating a set of spectrogram frames for each audio file;
 inputting the audio file along with a context of spectrogram frames into a
20 trained neural network;
 obtaining predicted character probabilities outputs from the trained neural network; and
 decoding a transcription of the input audio using the predicted character probabilities outputs from the trained neural network constrained by a

language model that interprets a string of characters from the predicted character probabilities outputs as a word or words.

12. The computer-implemented method of Claim 11 wherein the step of generating a set of spectrogram frames for each audio file comprises:

5 generating spectrogram frames wherein a spectrogram frame comprises a set of linearly spaced log filter banks computed over windows of a first value of milliseconds strided by a second value of milliseconds.

13. The computer-implemented method of Claim 11 wherein:

10 the step of inputting the audio file along with a context of spectrogram frames into a trained neural network comprises:

 inputting the audio file along with a context of spectrogram frames into a plurality of trained neural network; and

 the step of obtaining predicted character probabilities outputs from the trained neural network comprises:

15 ensembling predicted character probabilities outputs from the set of trained neural networks to obtain the predicted character probabilities.

14. The computer-implemented method of Claim 13 wherein the step of ensembling predicted character probabilities outputs from the set of trained neural networks to obtain the predicted character probabilities comprises:

20 addressing time shifts between trained neural network models by using one or more of the following comprising:

 using neural network models that exhibit the same temporal shift;

checking alignment between output of neural network models and shifting one or more of the outputs to align the outputs; and shifting the inputs into one or more of the neural network models to have aligned outputs.

5 15. The computer-implemented method of Claim 11 wherein the step of decoding a transcription of the input audio using the predicted character probabilities outputs from the trained neural network constrained by a language model that interprets the string of characters as words comprises:

10 given the predicted character probabilities outputs from the trained neural network, performing a search to find a sequence of characters that is most probable according to both the predicted character probabilities outputs and a trained *N*-gram language model output that interprets a string of characters from the predicted character probabilities outputs as a word or words.

15 16. The computer-implemented method of Claim 11 wherein the neural network model comprises a five-layer model comprising:

a first set of three layers that are non-recurrent;

a fourth layer that is a bi-directional recurrent network, which includes two sets of hidden units comprising a set with forward recurrence and a set with backward recurrence; and

20 a fifth layer that is a non-recurrent layer, which takes forward and backward units from the fourth layer as inputs and outputs the predicted character probabilities.

17. A non-transitory computer-readable medium or media comprising one or more sequences of instructions which, when executed by one or more processors, causes the steps to be performed comprising:

receiving an input audio from a user;

5 generating a set of spectrogram frames from the audio file;

inputting the audio file along with a context of spectrogram frames into a set of trained neural networks;

obtaining predicted character probabilities outputs from the set of trained neural network; and

10 decoding a transcription of the input audio using the predicted character probabilities outputs from the set of trained neural networks constrained by a language model that interprets a string of characters from the predicted character probabilities outputs as a word or words.

18. The non-transitory computer-readable medium or media of Claim 17 wherein the step of
15 generating a set of spectrogram frames comprises:

generating spectrogram frames wherein a spectrogram frame comprises a set of linearly spaced log filter banks computed over windows of a first value of milliseconds strided by a second value of milliseconds.

19. The non-transitory computer-readable medium or media of Claim 17 wherein the step of
20 obtaining predicted character probabilities outputs from the set of trained neural network comprises:

ensembling predicted character probabilities outputs from the set of trained neural networks to obtain the predicted character probabilities.

20. The non-transitory computer-readable medium or media of Claim 19 wherein the step of ensembling predicted character probabilities outputs from the set of trained neural networks to obtain the predicted character probabilities comprises:

addressing time shifts between trained neural network models by using one or

5 more of the following comprising:

using neural network models that exhibit the same temporal shift;

checking alignment between output of neural network models and shifting one

or more of the outputs to align the outputs; and

shifting the inputs into one or more of the neural network models to have

10 aligned outputs.

1/14

100

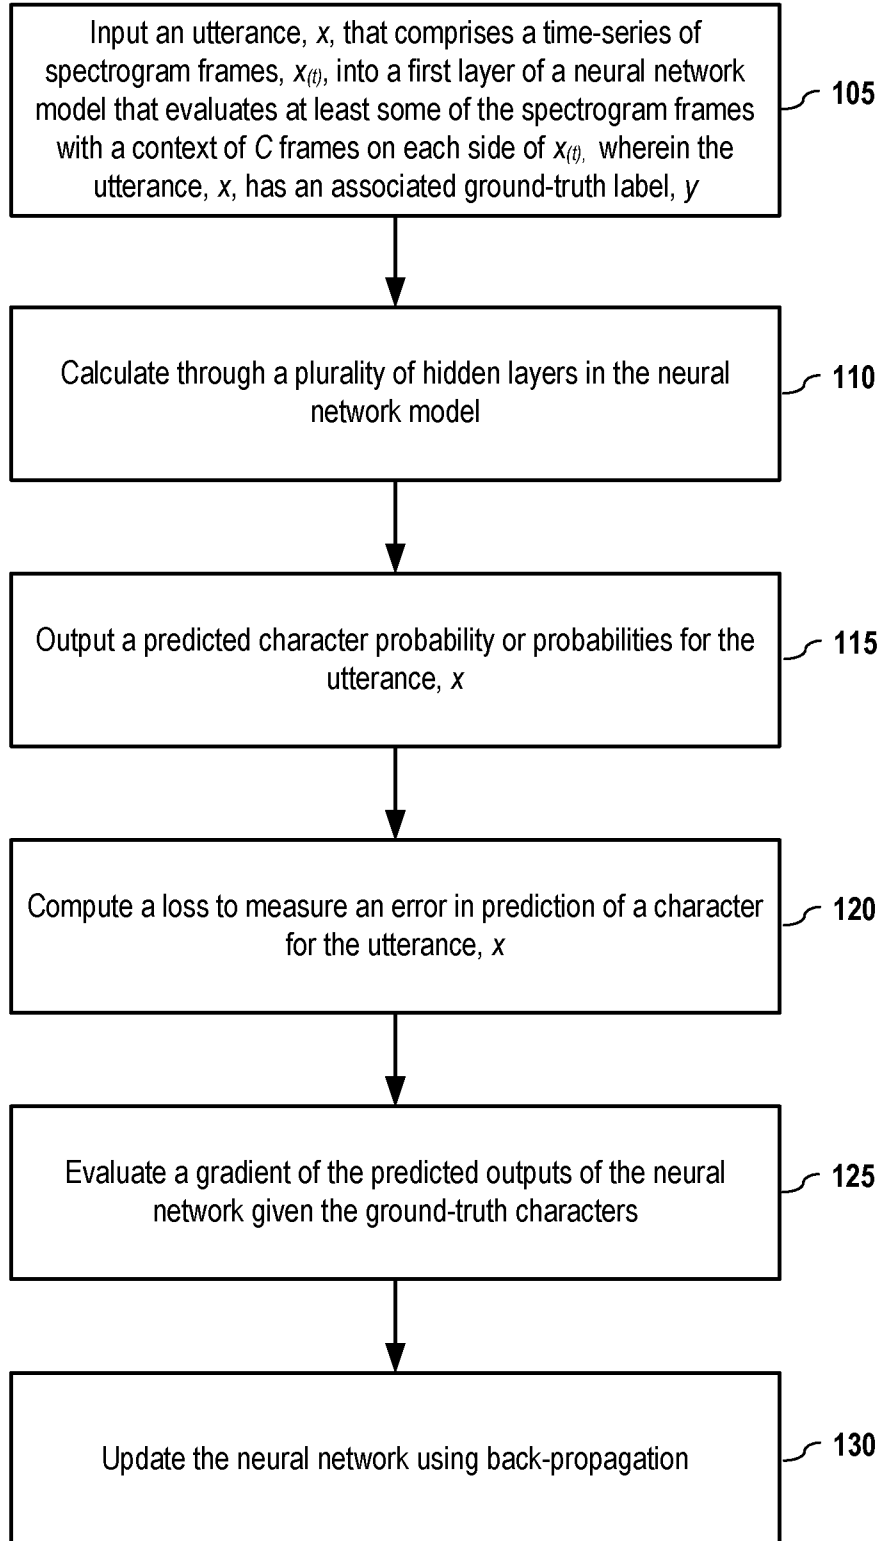


FIG. 1

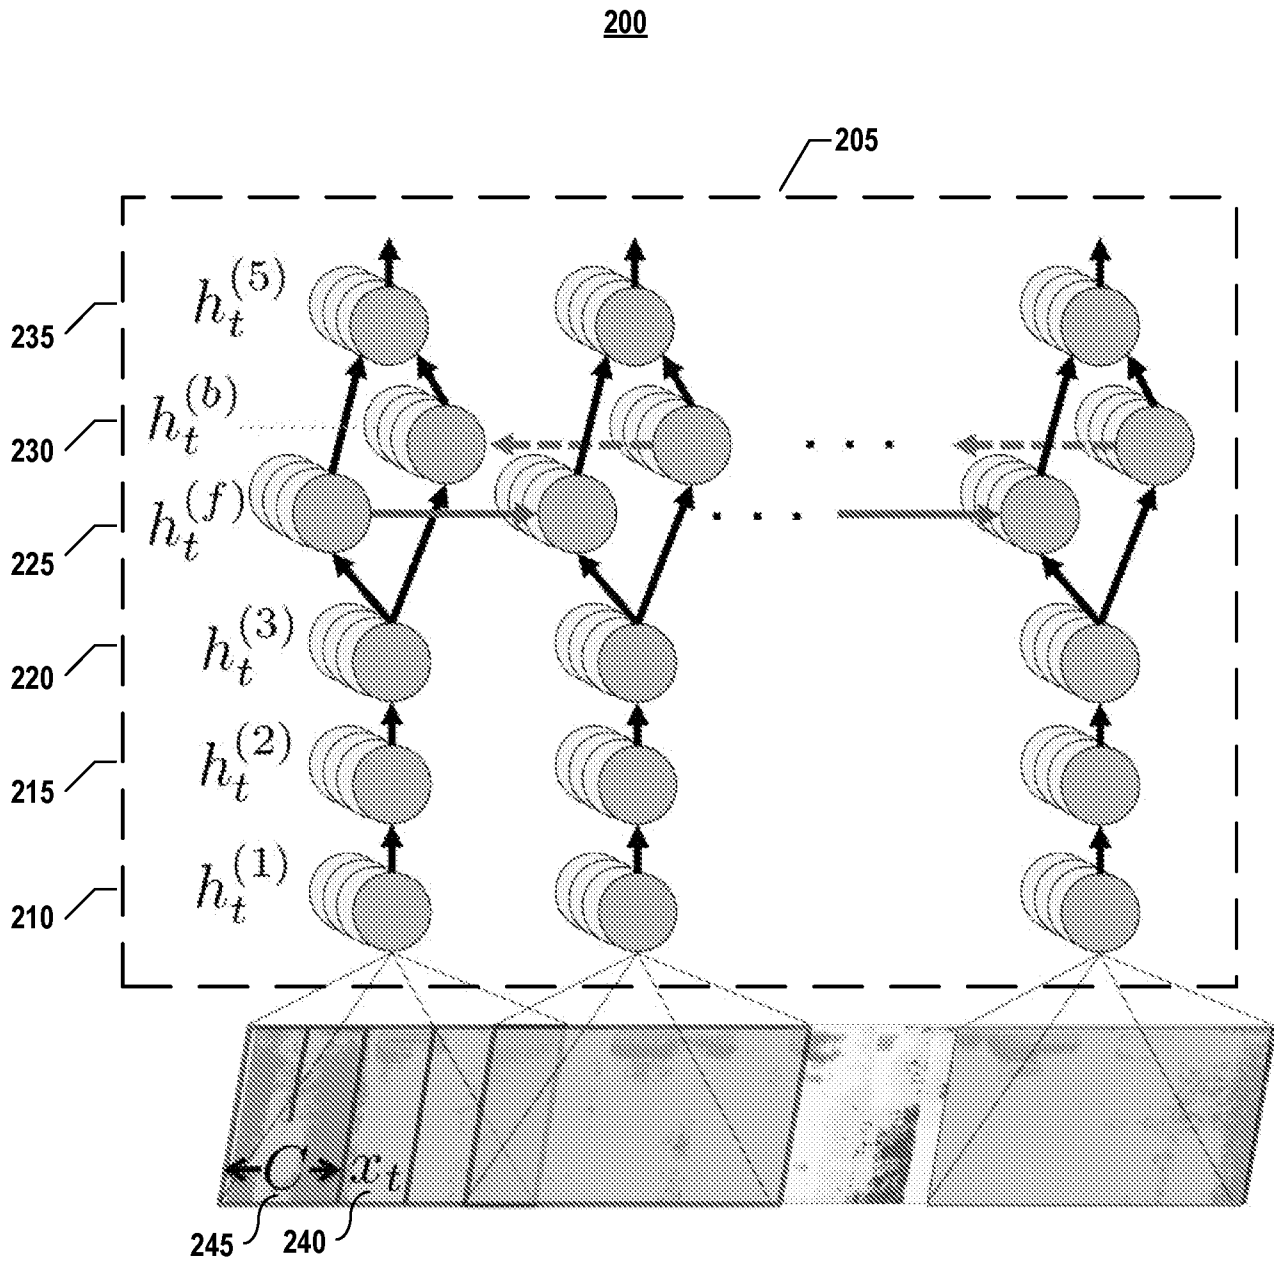


FIG. 2

3/14

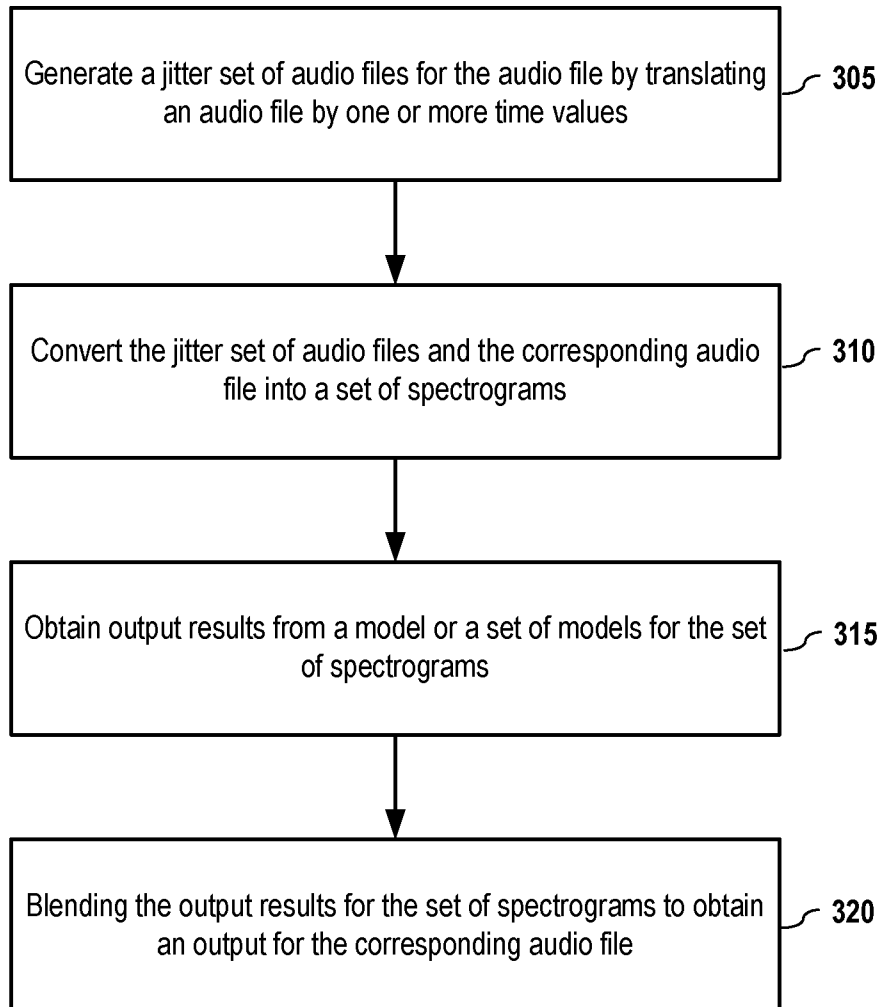
300

FIG. 3

4/14

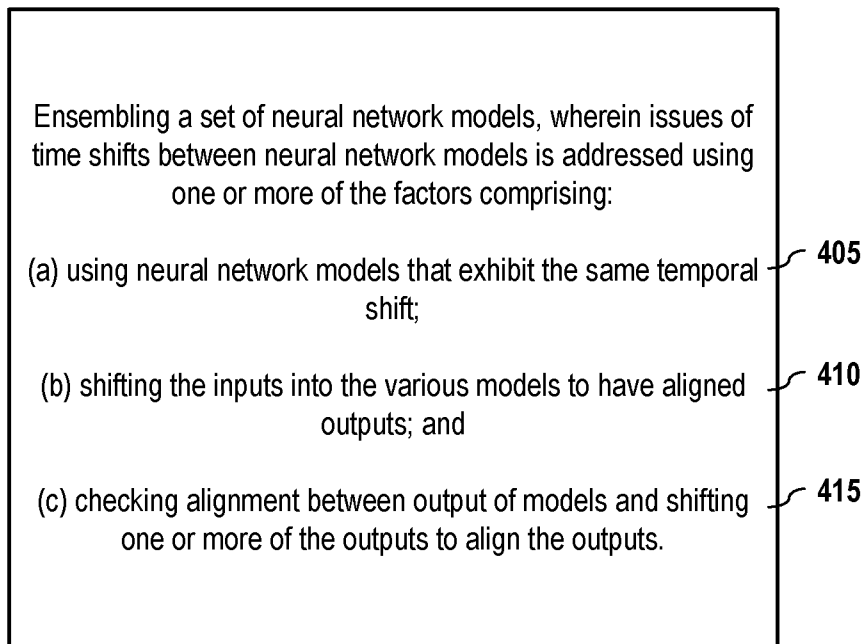
400

FIG. 4

5/14

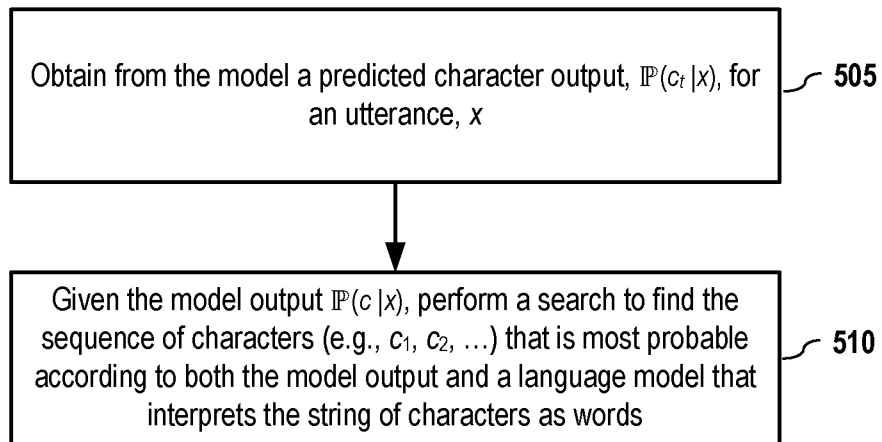
500

FIG. 5

6/14

600




FIG. 6

7/14

700

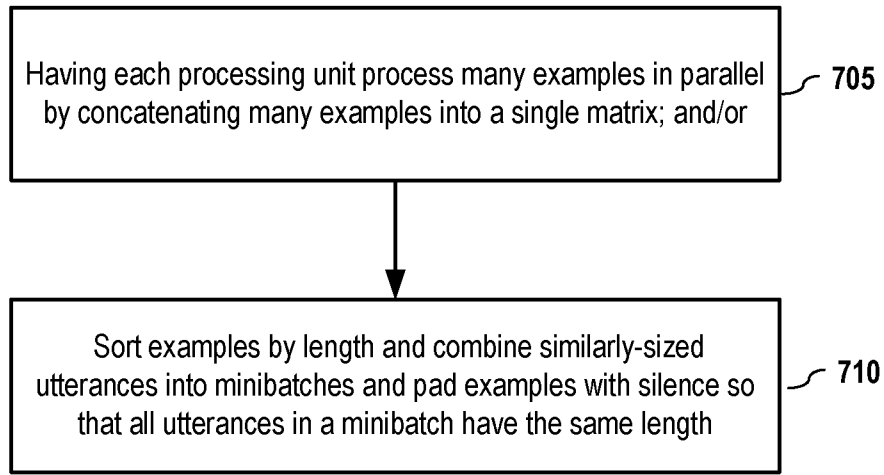


FIG. 7A

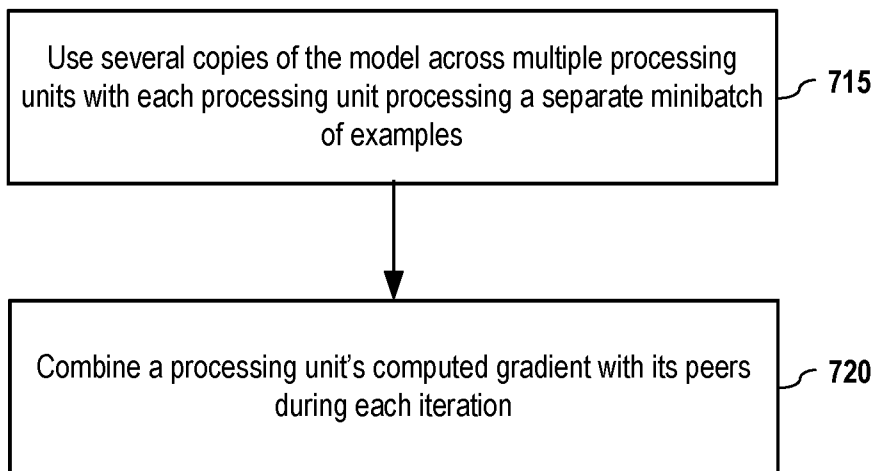
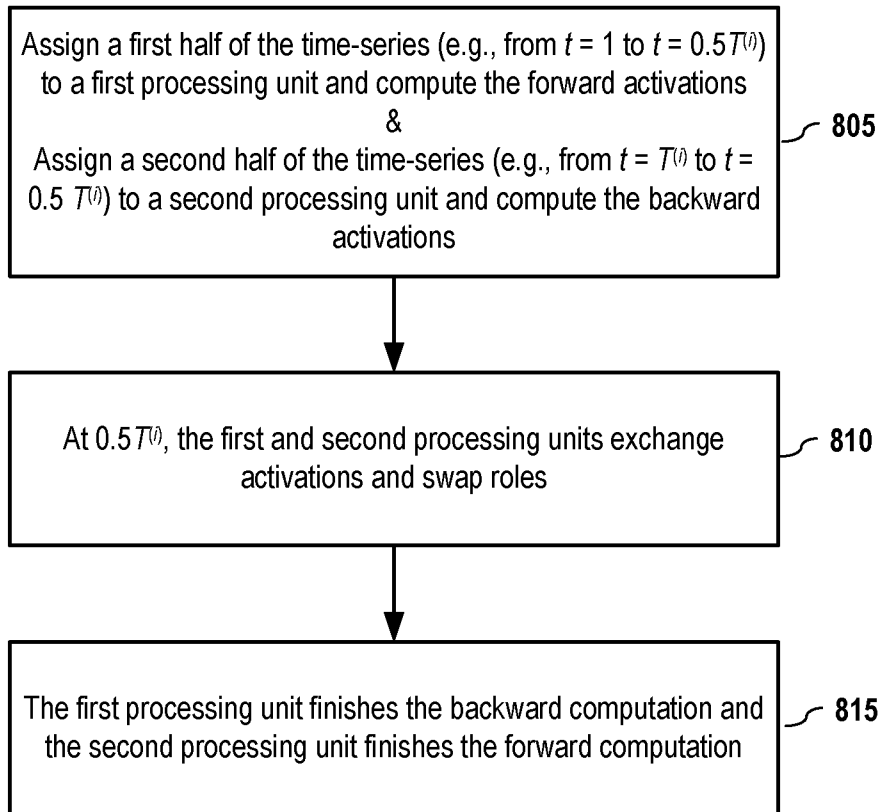


FIG. 7B

8/14

800**FIG. 8**

9/14

900




FIG. 9

10/14

1000

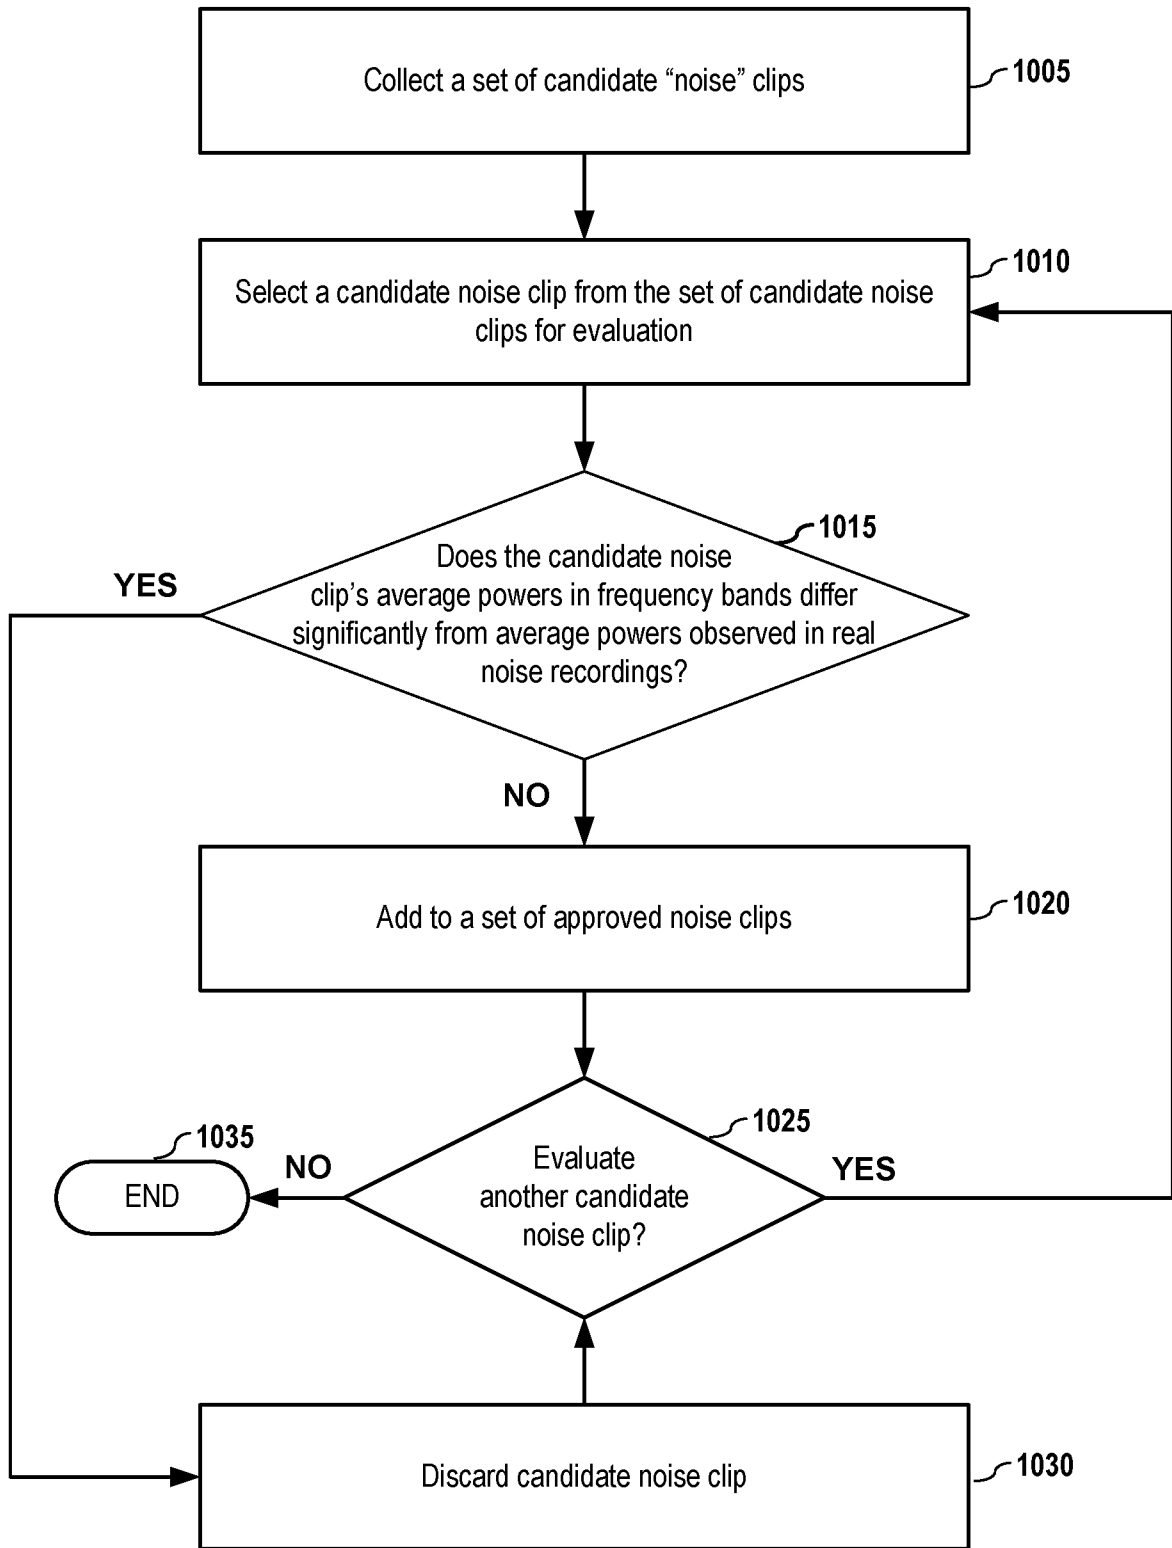


FIG. 10

11/14

1100

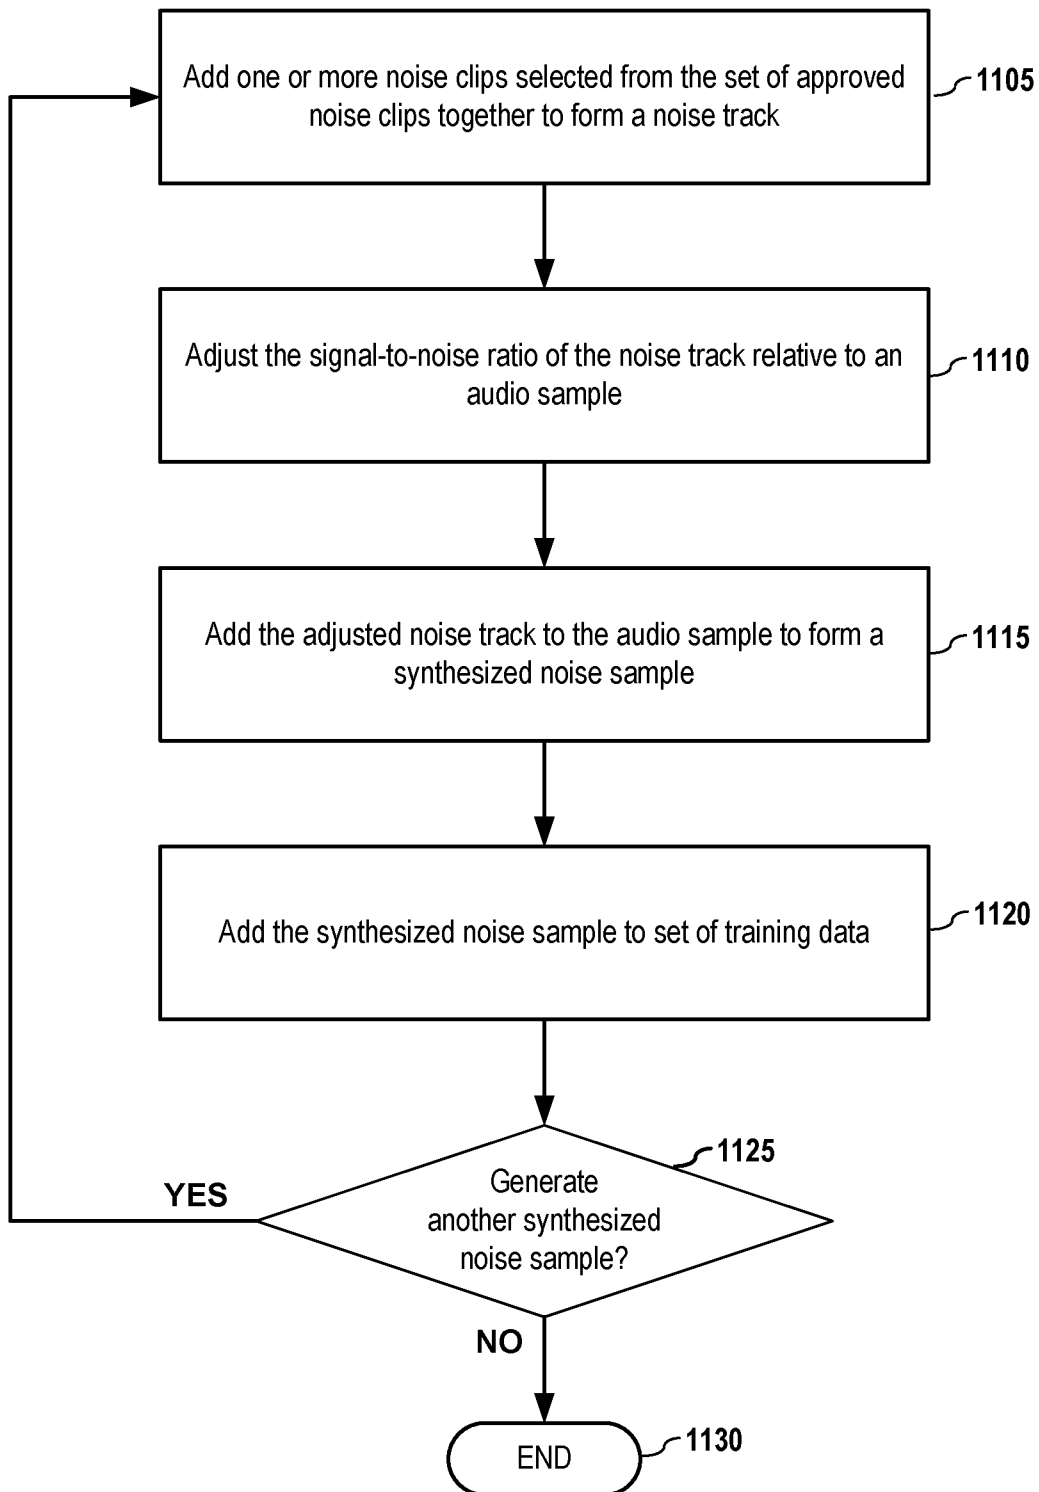


FIG. 11

12/14

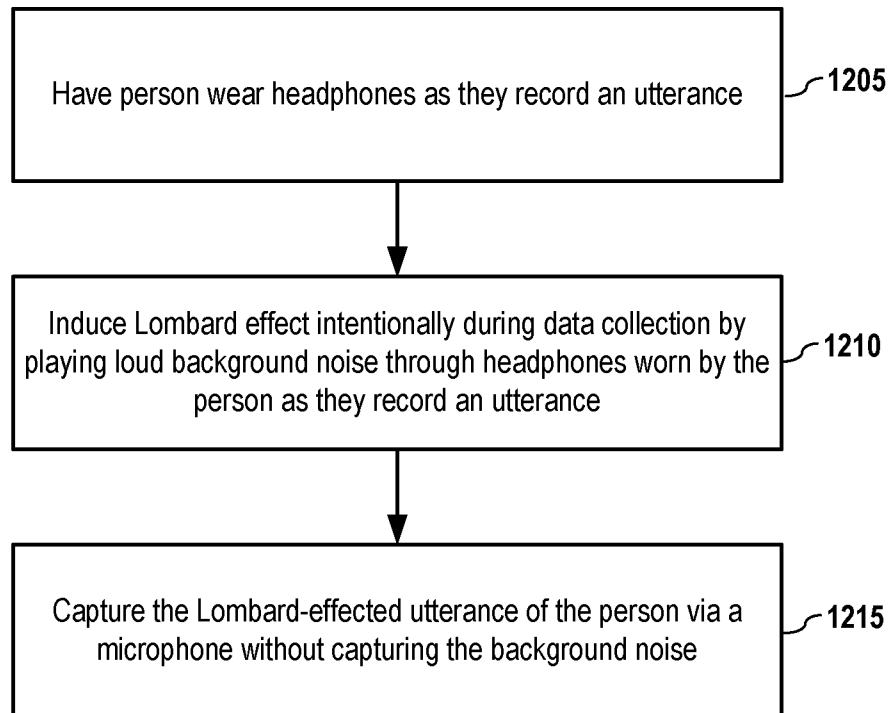
1200

FIG. 12

13/14

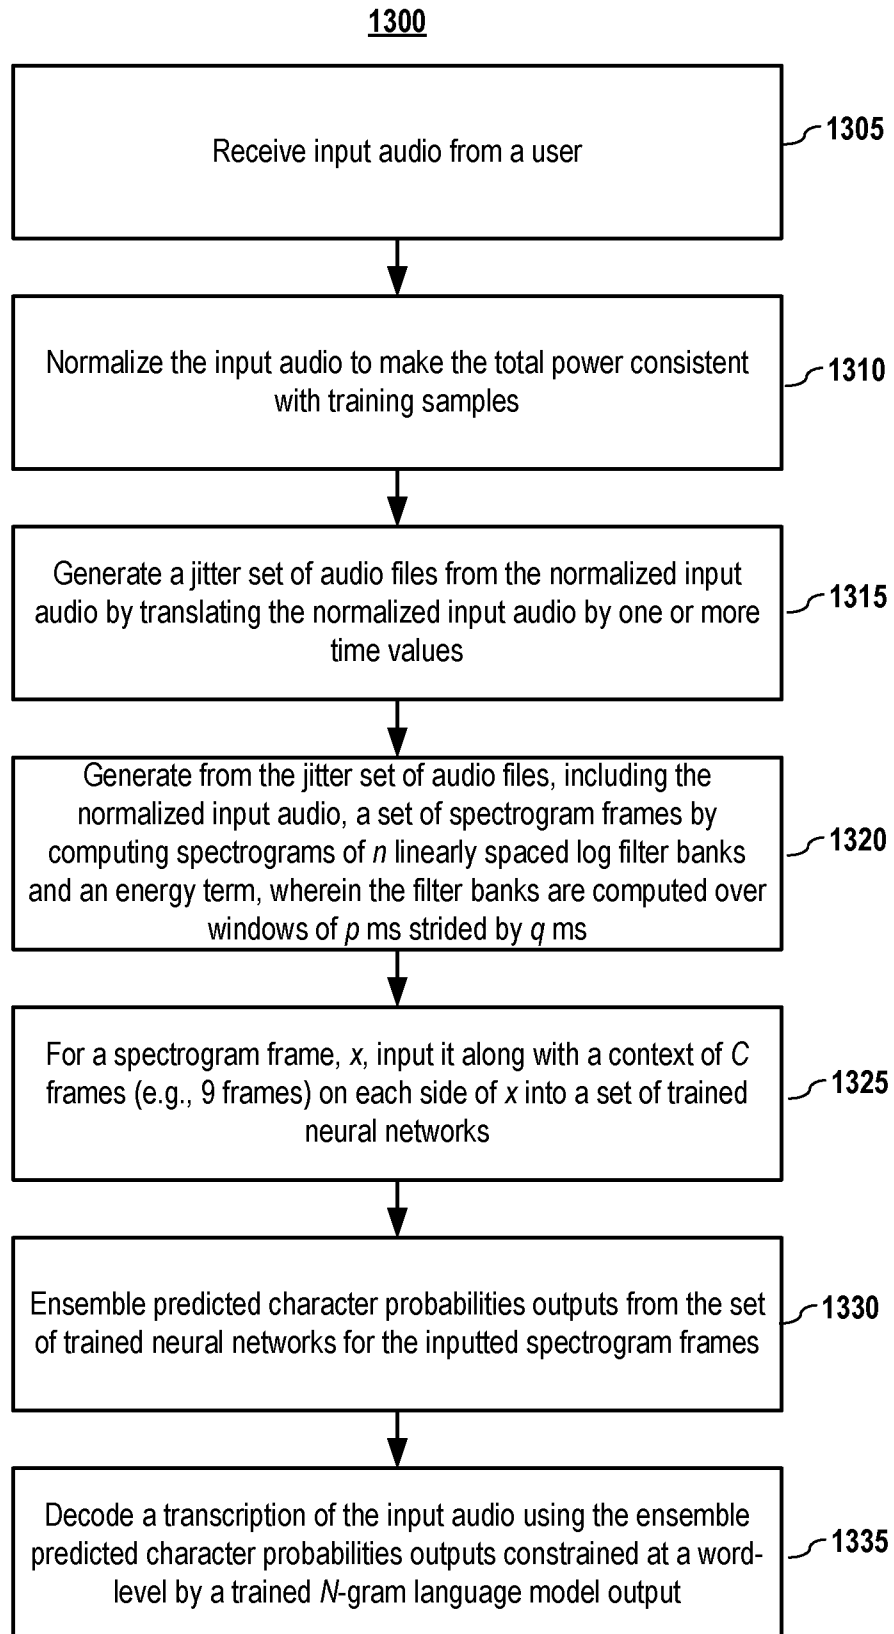


FIG. 13

14/14

1400

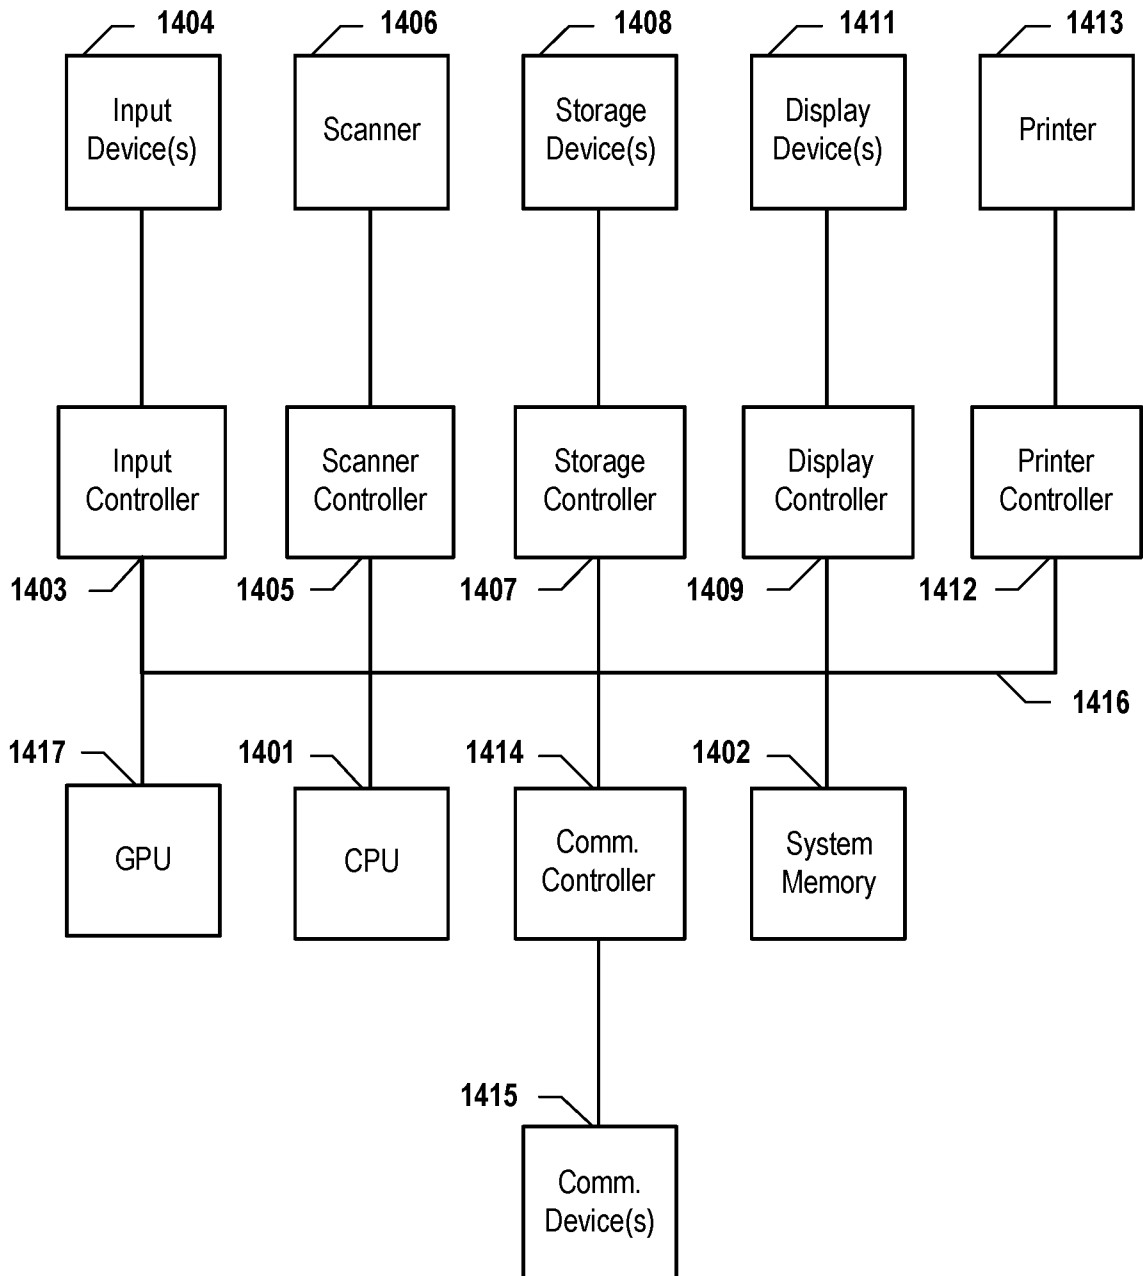


FIG. 14

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US 15/65617

A. CLASSIFICATION OF SUBJECT MATTER

IPC(8) - G10L 15/00 (2016.01)

CPC - G10L 15/22

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)

IPC (8) - G10L 15/00 (2016.01)

CPC - G10L 15/22

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

CPC - G10L 15/30; G10L 15/265 (See Keywords Below)

USPC - 704/231, 704/E15.001, 704/232

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

Thomsoninnovation.com; Patbase; Google Scholar; Google Patents; Gogole.com; Freepatentsonline; ProQuest Dialog
Search Terms: Speech, utterance, spoken, voice, recognition, transcription, model, learn, train, parallel, character, context, frame, spectrogram, layer, gradient, loss, error, back propagation, noise, jitter, etc.

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2014/0257803 A1 (YU et al.), 11 September 2014 (11.09.2014), entire document, especially Abstract; Para [0022], [0024], [0026], [0030], [0041]-[0044], [0052]-[0054]	1-10
Y	Graves et al. "Towards End-to-End Speech Recognition with Recurrent Neural Networks" In: Proceedings of the 31st International Conference on Machine Learning, June 21-26, 2014 [online][retrieved on 04 April 2016 (04.04.2016)] Retrieved from the Internet <URL:http://www.jmlr.org/proceedings/papers/v32/graves14.pdf> entire document, especially Abstract; page 2, page 5	1-10
Y	US 2011/0035215 A1 (SOMPOLINSKY et al.), 10 February 2011 (10.02.2011), entire document, especially Abstract; Para [0040]-[0043], [0082]	2-7
Y	US 2003/0110033 A1 (SHEIKHZADEH-NADJAR et al.), 12 June 2003 (12.06.2003), entire document, especially Abstract; Para [0030]-[0035], [0058]-[0060]	8-10
Y	US 2003/0236672 A1 (AARON et al.), 25 December 2003 (25.12.2003), entire document, especially Abstract; Para [0023]-[0026], [0034]	5
Y	US 2007/0106507 A1 (CHAROENRUENGKIT et al.), 10 May 2007 (10.05.2007), entire document, especially Abstract; Para [0033]-[0035], [0042]-[0047], [0057]-[0060]	6-7
A	US 2006/0217968 A1 (BURGES et al.), 28 September 2006 (28.09.2006), entire document	1-10
A	US 2012/0072215 A1 (YU et al.), 22 March 2012 (22.03.2012), entire document	1-10

 Further documents are listed in the continuation of Box C.

* 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

05 April 2016 (05.04.2016)

Date of mailing of the international search report

22 APR 2016

Name and mailing address of the ISA/US

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents

P.O. Box 1450, Alexandria, Virginia 22313-1450

Facsimile No. 571-273-8300

Authorized officer:

Lee W. Young

PCT Helpdesk: 571-272-4300
PCT OSP: 571-272-7774

Continuation of:
Box III. Observations where unity of invention is lacking

The inventions listed as Groups I-II do not relate to a single general inventive concept under PCT Rule 13.1 because, under PCT Rule 13.2, they lack the same or corresponding special technical features for the following reasons:

Special Technical Features:

The invention of Group I included the features of for each of a set of utterances: inputting an utterance that comprises a set of spectrogram frames into a first layer of the transcription model that evaluates each of the spectrogram frames from the set of spectrogram frames with a context of one or more spectrogram frames; outputting from the transcription model a predicted character or character probabilities for the utterance; and computing a loss to measure error in prediction for the utterance; evaluating a gradient of predicted outputs of the transcription model given the ground-truth characters; and updating the transcription model using back-propagation, not required by group II.

The invention of Group II included the features of receiving an input audio from a user; normalizing the input audio to make a total power of the input audio consistent with a set of training samples used to train a trained neural network model; generating a jitter set of audio files from the normalized input audio by translating the normalized input audio by one or more time values; for each audio file from the jitter set of audio files, which includes the normalized input audio: generating a set of spectrogram frames for each audio file; inputting the audio file along with a context of spectrogram frames into a trained neural network; obtaining predicted character probabilities outputs from the trained neural network; and decoding a transcription of the input audio using the predicted character probabilities outputs from the trained neural network constrained by a language model that interprets a string of characters from the predicted character probabilities outputs as a word or words, not required by group I.

Common Technical Features

Groups I-II share the technical features of set of spectrogram frames; transcription model; a context of one or more spectrogram frames; predicted character probabilities.

However, the shared technical features does not represent a contribution over prior art as anticipated by Graves et al. "Towards End-to-End Speech Recognition with Recurrent Neural Networks" In: Proceedings of the 31st International Conference on Machine Learning, June 21-26, 2014 [online][retrieved on 08 February 2016 (08.02.2016)] Retrieved from the Internet
<URL:<http://www.jmlr.org/proceedings/papers/v32/graves14.pdf>>

Graves et al. teaches the set of spectrogram frames (page 2, col 2; page 7, col 2 - individual frames of acoustic data; input was 15 frames of Mel-scale log filterbanks (1 centre frame ?7 frames of context) with 40 coefficients, deltas and acceleration); transcription model (Abstract - deep bidirectional LSTM recurrent neural network architecture and the Connectionist Temporal Classification); a context of one or more spectrogram frames (page 7, col 2 - The input was 15 frames of Mel-scale log filterbanks (1 centre frame ?7 frames of context) with 40 coefficients, deltas and accelerations); predicted character probabilities (figure 4 - the frame-level character probabilities emitted by the CTC layer).

As the common features were known in the art at the time of the invention, this cannot be considered a common technical feature that would otherwise unify the groups.

Therefore, Groups I-II lack unity under PCT Rule 13.

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US 15/65617

Box No. II Observations where certain claims were found unsearchable (Continuation of item 2 of first sheet)

This international search report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. Claims Nos.:
because they relate to subject matter not required to be searched by this Authority, namely:

2. Claims Nos.:
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:

3. Claims Nos.:
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box No. III Observations where unity of invention is lacking (Continuation of item 3 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

This application contains the following inventions or groups of inventions which are not so linked as to form a single general inventive concept under PCT Rule 13.1. In order for all inventions to be examined, the appropriate additional examination fees must be paid.

Group I: Claims 1-10 directed to a method for training a transcription model.

Group II: Claims 11-20 directed to a method for transcribing speech.

--- (See Continuation in Supplemental Box) ---

1. As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.
2. As all searchable claims could be searched without effort justifying additional fees, this Authority did not invite payment of additional fees.
3. As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:
4. No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:
1-10

- Remark on Protest**
- The additional search fees were accompanied by the applicant's protest and, where applicable, the payment of a protest fee.
- The additional search fees were accompanied by the applicant's protest but the applicable protest fee was not paid within the time limit specified in the invitation.
- No protest accompanied the payment of additional search fees.