

CONFIDENCE ESTIMATION AND DELETION PREDICTION USING BIDIRECTIONAL RECURRENT NEURAL NETWORKS

A. Ragni, Q. Li, M. J. F. Gales, Y. Wang

Department of Engineering, University of Cambridge
Trumpington Street, Cambridge CB2 1PZ, UK
{ar527, ql264, mjfg, yw396}@eng.cam.ac.uk

ABSTRACT

The standard approach to assess reliability of automatic speech transcriptions is through the use of confidence scores. If accurate, these scores provide a flexible mechanism to flag transcription errors for upstream and downstream applications. One challenging type of errors that recognisers make are deletions. These errors are not accounted for by the standard confidence estimation schemes and are hard to rectify in the upstream and downstream processing. High deletion rates are prominent in limited resource and highly mismatched training/testing conditions studied under IARPA Babel and Material programs. This paper looks at the use of bidirectional recurrent neural networks to yield confidence estimates in predicted as well as deleted words. Several simple schemes are examined for combination. To assess usefulness of this approach, the combined confidence score is examined for untranscribed data selection that favours transcriptions with lower deletion errors. Experiments are conducted using IARPA Babel/Material program languages.

Index Terms— confidence score, deletion error, bidirectional recurrent neural network

1. INTRODUCTION

Recent years have seen an increase in demand for speech enabled solutions. These range from speech transcription to personal assistants [1, 2] designed to handle ever increasing in complexity human-machine interactions [3]. The perceived usefulness of these applications depends on the quality of the underlying automatic speech recogniser which for some tasks has been reported to approach the agreement level of human annotators [4, 5]. Even in such favourable conditions a measure of reliability in hypothesised words can prove useful for flagging up words that may need to be handled differently. For more challenging scenarios such measure can prove fundamental for achieving high performance [6] and yielding reliable feedback [7]. The development cycle of speech recognisers may itself benefit from the measure of reliability. Rather than expending resources on annotating data that can already be reliably transcribed a more challenging set can be automatically identified in a live data and incorporated into training these recognisers.

This research is based upon work supported in part by the ALTA Institute, Cambridge University and the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via Air Force Research Laboratory (AFRL) contract # FA8650-17-C-9117. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of ODNI, IARPA, AFRL or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein.

Confidence scores [8] have been traditionally used as the measure of reliability of automatic speech transcriptions in both downstream as well as upstream tasks. In the simplest form these scores are posterior probabilities associated with words in hypothesised transcriptions. These probabilities however are believed to be over estimated [9] which could hamper their utility as the confidence measure. A range of approaches have been examined to improve accuracy of confidence estimates. This includes simple piece-wise linear mappings in the form of decision trees [9], generalised linear models [10], highly non-linear feed-forward neural networks [11] and sequence models, such as conditional random fields (CRF) [12, 13] and recurrent neural networks [14]. All these approaches reported gains over posterior based confidence scores using a range of evaluation metrics. This paper examines a neural network model that takes into account information about not only the current word and its history but also about all the following words. The bidirectional nature of this model is expected to be well suited for the task of confidence estimation given the complete knowledge of the entire hypothesised sequence. This work thus combines the benefits of neural networks for feature extraction and context modelling with the whole sequence context of CRFs in a single model.¹

The standard formulation of confidence scores is focused on assigning estimates to hypothesised words only. This leaves any deleted word out of consideration. Deletions, however, are a very challenging type of error to recover from in both downstream and upstream tasks. Moreover, drifts in deployment domains either planned [16] or unplanned [17] typically lead to large increases in deletion errors. Though schemes such as frame weighting [18, 19] could be used to alleviate the impact of deletions to a certain extent, an explicit deletion prediction is expected to be more appropriate. The previous work on deletion prediction [20] has looked at predicting whether a deletion occurs between current and next word using CRFs. This paper extends that work by predicting deletions by means of a bidirectional recurrent neural network (BiRNN) model [21]. For these deletion predictions to be useful a mechanism is needed to incorporate them in upstream and downstream tasks. This paper examines utterance selection for unsupervised training as the target task and shows how simple combinations of confidence and deletion predictions can yield selected data with different deletion error rate characteristics. Experiments with IARPA Babel/Material program languages show promise of this approach for confidence estimation and deletion prediction.

The rest of this paper is organised as follows. Section 2 discusses confidences and evaluation metrics. The following section 3 examines interactions between these scores and deletions. Section 4

¹This part of the work was completed prior to the publication of an independent study on confidence prediction using BiRNNs [15].

describes the proposed BiRNN architecture for confidence and deletion prediction. Experimental results are presented in Section 5. Finally, conclusions drawn from this work are given in Section 6.

2. CONFIDENCE SCORES

Quantifying uncertainty in predictions is a challenge that modern speech recognisers are faced with due to their ubiquitous use in an ever growing number of applications. Confidence scores have been a traditional measure of certainty of a speech recogniser in its prediction [8]. These scores are typically derived for each hypothesised word. The simplest form of a confidence score is a posterior probability of a hypothesised word that can be estimated over a lattice of possible transcriptions generated by the speech recogniser [22, 23]. This simple approach has been found [9] to significantly over-estimate these scores. The reason most often stated is a limited size of lattices over which these posterior probabilities are normalised. In order to adjust confidences so they better reflect the true values a number of approaches have been examined. These can be divided into simple mappings that make use of only original scores and more complex approaches that make use of additional information. The most prominent example of the former are decision trees [9]. This simple approach estimates a piece-wise linear mapping given some held-out data. The latter group, including the work reported in this paper, make use of more complex approaches, such as CRFs [12], feed-forward [11] and recurrent neural networks [14] to estimate confidence scores from manually designed features. These features may include various statistics extracted from audio, acoustic models, language models and lattices [24].

In the ideal, limiting, case, confidence scores of correctly hypothesised words are one and zero otherwise. In order to measure how far any given set of confidence scores is from the ideal scenario a number of metrics could be used. One popular metric is a normalised cross-entropy (NCE) [25]. This metric measures a relative change in the cross-entropy caused when an empirical estimate of ASR correctness is replaced by hypothesised confidences. The estimate of ASR correctness can be computed from reference confidences $\mathbf{c}^* = \{c_1^*, \dots, c_T^*\}$ by

$$P(C = 1) = \frac{1}{T} \sum_{t=1}^T c_t^* \quad (1)$$

where $P(C = 1)$ is the probability of ASR system being correct and $P(C = 0) = 1 - P(C = 1)$ is for the opposite event, T is the number of hypothesised words. The average binary cross-entropy between this estimate and reference confidences is given by

$$\overline{H}(\mathbf{c}^*) = -\frac{1}{T} \sum_{t=1}^T c_t^* \log(P(C = 1)) + (1 - c_t^*) \log(P(C = 0)) \quad (2)$$

On the other hand, the average binary cross-entropy between hypothesised $\mathbf{c} = \{c_1, \dots, c_T\}$ and reference confidences is given by

$$H(\mathbf{c}|\mathbf{c}^*) = -\frac{1}{T} \sum_{t=1}^T c_t^* \log(c_t) + (1 - c_t^*) \log(1 - c_t) \quad (3)$$

If the hypothesised confidences \mathbf{c} are systematically better than the estimate of ASR system being correct, $P(C = 1)$, the relative change in cross-entropy or NCE given by

$$\text{NCE}(\mathbf{c}, \mathbf{c}^*) = \frac{\overline{H}(\mathbf{c}^*) - H(\mathbf{c}|\mathbf{c}^*)}{\overline{H}(\mathbf{c}^*)} \quad (4)$$

is positive. In the opposite case it is negative. The maximum value of NCE is 1 which corresponds to the case where hypothesised confidences match reference confidences exactly.

Though popular NCE may not be the most optimal metric to assess confidence scores. For these scores to be the perfect correct/incorrect predictor, it is not necessary to match reference confidences exactly. It is however sufficient to yield correct rank ordering such that confidences of all incorrectly hypothesised words are below some threshold and those of all correctly hypothesised words are above. This suggests that an area under the curve (AUC) type metric may be more suitable. When choosing which curve to use it is important to know the balance of positive and negative class examples. For balanced data it is common to use a receiver operating curve (ROC) of false positive (FP) rate (FPR) against true positive (TP) rate (TPR) for a range of threshold θ values

$$\text{TPR}(\theta) = \frac{\text{TP}(\theta)}{\text{TP}(\theta) + \text{FN}(\theta)}, \quad \text{FPR}(\theta) = \frac{\text{FP}(\theta)}{\text{FP}(\theta) + \text{TN}(\theta)} \quad (5)$$

where $\text{TN}(\theta)$ and $\text{FN}(\theta)$ are true negatives and false negatives. The AUC of a random guess under the ROC curve is 0.5. Values larger than that signal better than random performance. For imbalanced data it is more appropriate to use a precision-recall curve [26]

$$\text{Precision}(\theta) = \frac{\text{TP}(\theta)}{\text{TP}(\theta) + \text{FP}(\theta)}, \quad \text{Recall}(\theta) = \frac{\text{TP}(\theta)}{\text{TP}(\theta) + \text{FN}(\theta)} \quad (6)$$

The AUC associated with a random guess under the precision recall curve depends on the ratio between positive P and negative N examples and is given by $\frac{P}{P+N}$. For a perfectly balanced data set this yields 0.5 and less and larger than 0.5 for data sets dominated by negative and positive examples respectively. Speech recognisers described in this work operate in the region where there are more correctly than incorrectly hypothesised words. However, as the difference is not large the ROC curve will be used in this work.

Confidence scores are used not only in downstream applications but also in upstream tasks. For instance in semi-supervised training confidence scores are routinely used for selecting high-confidence hypotheses for training acoustic models [19]. This requires extending the definition of confidence scores to an utterance (segment) and audio recording level. Given a sequence of confidence scores $\mathbf{c}_r = \{c_{r,1}, \dots, c_{r,T_r}\}$ associated with recording r , a common solution is to compute frame-weighted confidence score [18, 19]

$$c_r = \frac{\sum_{t=1}^{T_r} \lambda_{r,t} c_{r,t}}{\sum_{t=1}^{T_r} \lambda_{r,t}} \quad (7)$$

where T_r is the number of hypothesised words and $\lambda_{r,t}$ is the number of frames associated with hypothesised word $w_{r,t}$. This enables the complete dataset to be rank ordered and manipulated to yield subsets of the data with suitable confidence characteristics. The frame-weighting applied in equation (7) ensures that long, high confidence, regions do not get penalised unfairly by short, low confidence, regions. The fundamental problem with using equation (7) for data selection is that it does not take into account deletion errors and may select transcriptions embedding large numbers of deletions.

3. DELETION ERRORS

Errors made by a speech recogniser are traditionally partitioned into 3 types: substitution, insertion and deletion. The former two have a direct realisation in the form of hypothesised words whereas deletions do not. The precise nature of partitioning into these error types

is a complex issue. First, the errors are computed by a Levenshtein alignment that weighs different types unequally [27]. This balance is further affected non-uniformly by the task at hand, the forms of acoustic and language models used, decoding approaches. Any mismatch between training and testing conditions typically would lead to an additional imbalance. Although some forms of imbalance may not necessarily be harmful, high levels of deletion errors do not interact well with a range of techniques such as discriminative adaptation/training, semi-supervised training, forms of minimum Bayes' risk decoding. Therefore, keeping deletion error under control is an important attribute of building speech recognisers.

The standard confidence estimation and confidence-based approaches do not explicitly address the presence of deletions [20]. This may not necessarily lead to a problem however. Figure 1 shows an example of the impact that the average frame-weighted confidence-based data selection has on word error rate (WER). The underlying is a narrow-band conversational telephone speech (CTS) recogniser for a limited resource language. As can be seen from

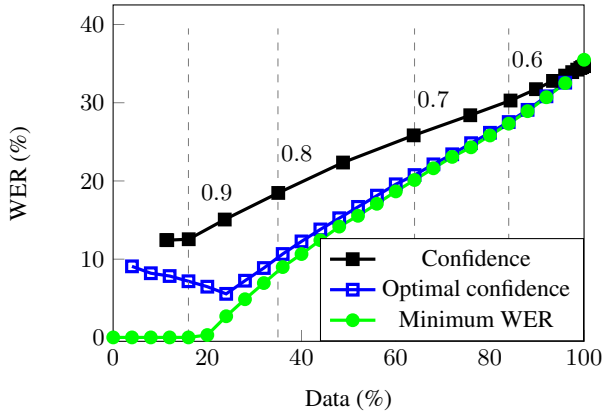


Fig. 1. Example impact of data selection approaches on word error rate (WER) in selected transcriptions

Figure 1, the WER of the full data set is above 30%. The top, filled squares, line shows that as the amount of selected data is decreased, by imposing ever increasing minimum average frame-weighted confidence constraints, the WER decreases approaching 10%. The middle, hollow squares, oracle line shows what would have happened if correct confidences were available. The bottom, filled circles, oracle line shows the impact of also using the knowledge of deletion errors. These lines show that large improvements are possible if better confidence scores were used and unless limited quantities of data are selected the impact of deletion errors can be small. The latter observation however does not take into account the presence of correlation between substitutions, insertions and deletion errors that can be exploited if deletions were taken into account by the data selection scheme. Figure 2 provides an insight into the individual error distributions during the selection process. As the amount of confidence required to be selected increases the relative reduction in all error types increases up to 0.8 average frame-weighted confidence level where the speech recogniser starts to delete progressively more words. Thus, the transcriptions that the speech recogniser is most confident about may have worst deletion characteristics. Table 1 shows what happens when this narrow-band speech recogniser is used to bootstrap a wide-band speech recogniser on some untranscribed news broadcast data. The first line shows the narrow-band development

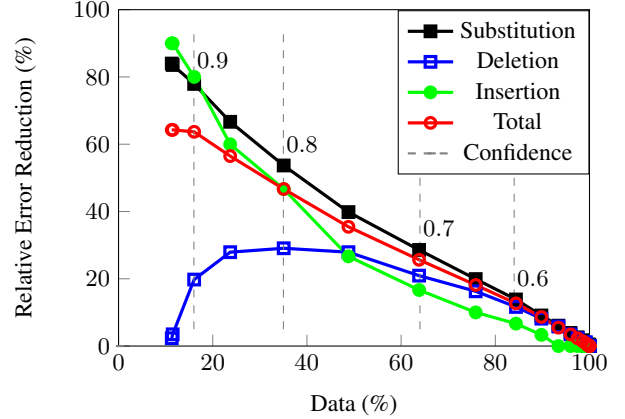


Fig. 2. Example of relative reduction in substitution, deletion, insertion and total error in confidence-based data selection

Type	Band		Error (%)			
	Model	Data	Sub	Del	Ins	Tot
Sup	Narrow	Narrow	24.1	8.3	3.1	35.5
	Narrow	Wide	23.4	19.0	1.3	43.7
Unsup	Wide	Wide	15.6	25.5	0.9	42.0

Table 1. Impact of band mismatch on unsupervised training

data. The middle line shows its performance on a down-sampled set of wide-band data containing news and topical broadcasts. As can be seen this mismatch has largely affected deletions. The last line shows that automatic transcriptions produced in such mismatched configuration, when used to train a wide-band recogniser, lead to an additional increase in deletions. Thus schemes that can lower deletions in selected automatic transcriptions would be of interest in such mismatched configurations.

If confidence estimates are supplemented with deletion estimates the standard confidence-based selection scheme needs to be altered. A range of schemes can be devised that weigh confidences and deletions in a number of different ways. One approach would be to discount confidences in equation (7) by scaled deletion estimates

$$\hat{c}_t = c_t - \theta_d d_t \quad (8)$$

$$\hat{c}_1 = c_1 - \theta_d d_1 - \theta_s s \quad (9)$$

This scheme has two free parameters and degenerates to the standard confidence-based selection scheme when θ_d and θ_s are zero. An alternative approach would be to threshold both estimates to yield estimates of correct hypotheses

$$\text{Cor}(\mathbf{c}; \theta_c) = \sum_{t=1}^T \delta(c_t \geq \theta_c) \quad (10)$$

substitution and insertion errors

$$\text{Inc}(\mathbf{c}; \theta_c) = \sum_{t=1}^T \delta(c_t < \theta_c) \quad (11)$$

as well as deletions

$$\text{Del}(\mathbf{d}; \theta_s, \theta_d) = \delta(s \geq \theta_s) + \sum_{t=1}^T \delta(d_t \geq \theta_d) \quad (12)$$

where θ_c and θ_s, θ_d are thresholds for correct hypotheses and deletions, d_t is an estimate of deleting a word following the hypothesis at time t and s is an estimate of deleting a word prior to the first hypothesised word. The confidence and deletion based data selection can be performed by rank ordering utterances/recordings according to an estimate of WER given by

$$\text{WER}(\mathbf{c}, \mathbf{d}; \boldsymbol{\theta}) = \frac{\text{Inc}(\mathbf{c}; \theta_c) + \text{Del}(\mathbf{d}; \theta_s, \theta_d)}{\theta_p \cdot \text{Inc}(\mathbf{c}; \theta_c) + \text{Cor}(\mathbf{c}; \theta_c)} \quad (13)$$

where θ_p is a penalty term to reduce over-estimation of reference words in the denominator since $\text{Inc}(\mathbf{c}; \theta_c)$ counts both substitution and insertion errors. Omitting $\text{Inc}(\mathbf{c}; \theta_c)$ term in the numerator would enable to perform data selection based on the estimate of deletion error alone, however, some safe guarding mechanism would be required to ensure that the overall WER is low. The thresholds $\boldsymbol{\theta}$ can be optimised by minimising mean squared error between the true and predicted WER on some held-out data using, for instance, a grid search or a constrained derivative free method [16].

4. BIDIRECTIONAL RECURRENT NEURAL NETWORKS

In this work bidirectional recurrent neural networks (BiRNN) [21] are used for estimating confidences and deletions. Figure 3 (a) shows the simplest form of BiRNNs that predicts confidence at time t by means of two recurrent units modelling past $\vec{\mathbf{h}}_t$ and future $\overleftarrow{\mathbf{h}}_t$ information. Various options exist how the past or history information

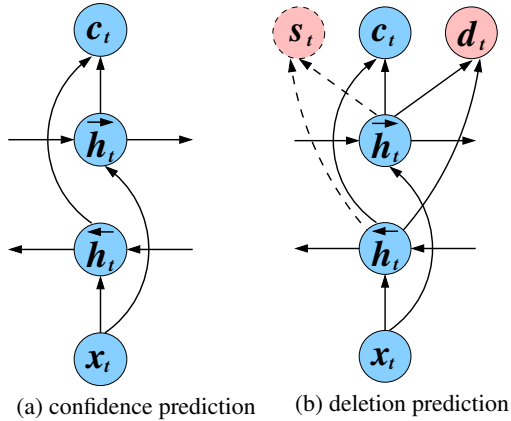


Fig. 3. Bidirectional recurrent neural networks

can be updated at the following time $t + 1$. In the simplest case

$$\vec{\mathbf{h}}_{t+1} = \sigma(\mathbf{W}^{(\vec{h})} \vec{\mathbf{h}}_t + \mathbf{W}^{(x)} \mathbf{x}_{t+1}) \quad (14)$$

where \mathbf{x}_{t+1} is an input feature vector, $\mathbf{W}^{(x)}$ is an input matrix, $\mathbf{W}^{(\vec{h})}$ is a history matrix and σ is a non-linearity such as sigmoid. The future information is typically modelled in the same fashion though this is not a requirement. The past and future information at any time t can be combined into a single context vector

$$\mathbf{h}_t = \left[\vec{\mathbf{h}}_t \quad \overleftarrow{\mathbf{h}}_t \right]^T \quad (15)$$

Given \mathbf{h}_t , the confidence c_t can be modelled by

$$c_t = \sigma(\mathbf{w}^{(c)T} \mathbf{h}_t + b^{(c)}) \quad (16)$$

where $\mathbf{w}^{(c)}$ is a parameter vector and $b^{(c)}$ is a bias, σ is any non-linearity mapping confidences into $[0, 1]$ range.

One possible way to extend this BiRNN from predicting just confidences to also predicting deletions is shown in Figure 3 (b). Here, at any time t an additional, deletion, prediction d_t is made that a deletion occurred between current and the following word

$$d_t = \sigma(\mathbf{w}^{(d)T} \mathbf{h}_t + b^{(d)}) \quad (17)$$

At time $t = 1$ an additional prediction is made to yield an estimate of deletion prior to the first hypothesised word

$$s = \sigma(\mathbf{w}^{(s)T} \mathbf{h}_1 + b^{(s)}) \quad (18)$$

Not that such definition cannot take into account multiple consecutive deletions. Training can be performed by minimising cross-entropies between true and predicted values of confidences and deletions

$$\mathcal{F}(\boldsymbol{\theta}) = \sum_{r=1}^R \left(\sum_{t=1}^{T_r} c_{r,t}^* \log(c_{r,t}) + d_{r,t}^* \log(d_{r,t}) \right) + s_r^* \log(s_r) \quad (19)$$

where the asterisk $*$ denotes true target values. For improved generalisation the objective function above can be regularised by adding a scaled L_2 norm of model parameters $\boldsymbol{\theta}$.

For this approach to be successful, the input features should include information about hypothesised word sequence that is correlated with different error types. A large body of work has been produced on hand-crafted features for these general machine learning approaches [24, 11, 28, 20]. For this preliminary exploration a simple and compact form was examined for each hypothesised word

$$\mathbf{x}_t = \begin{bmatrix} \lambda_t \\ c_t \\ \mathbf{e}(w_t) \\ \mathcal{O}(w_t|w_{t-1}, \dots) \\ \log(P(w_t|w_{t-1}, \dots)) \\ |w_t| \\ \delta_{t-1} \\ \delta_{t+1} \end{bmatrix} \quad (20)$$

This includes duration λ_t , confidence c_t as estimated by the decision tree, word w_t embedding $\mathbf{e}(w_t)$, n -gram language model statistics in the form of log-probability $\log(P(w_t|w_{t-1}, \dots))$ and the highest non-backoff n -gram order $\mathcal{O}(w_t|w_{t-1}, \dots)$, character length $|w_t|$ and any time gap between the current word and preceding δ_{t-1} and following δ_{t+1} word. Though this set is limited it provides a sufficient ground for exploring confidence and deletion prediction.

5. EXPERIMENTS

Experiments were conducted on three languages: Georgian, Tagalog and Swahili. These languages originate from different language families and possess limited resources available for the development of speech recognisers. All three languages were a subject of research under the IARPA BABEL program for developing agile and robust speech recognition technology that can be rapidly applied to any language. The latter two languages have also been included into the new IARPA initiative for machine translation for English retrieval and summarisation of text and speech (MATERIAL). For each language only limited quantities of transcribed narrow-band conversational telephone speech (CTS) data are available. This consists of

40-60 hours for training, 10 hours for development and 10-15 hours for evaluation. The Material program interested in both narrow-band CTS and a range of wide-band domains has also released 10 hours of transcribed wide-band news and topical broadcasts for performers to analyse how well they can handle large domain mismatches.

All speech recognisers made use of lattice-free maximum mutual information (LF-MMI) interleaved time-delay and long short-term memory neural network (TDNN-LSTM) [29, 30] acoustic models and n -gram language models [31]. Unless otherwise stated, all acoustic models were trained on narrow-band CTS data. Input features for acoustic models consisted of 24 dimensional filter bank coefficients, probability of voicing and pitch [32]. All language models were obtained by interpolating multiple source language models on transcriptions of development data. The source language models were built on training data transcriptions and various sources of web data harvested by Babel/Material program performers [33, 34]. The amount of web data varied from 140 million words for Georgian to 1 billion words for Tagalog and Swahili. Pronunciations for these words were obtained by a grapheme-to-phoneme model [35]. Table 2 shows performance of these recognisers on Babel development (`dev`) and evaluation (`evalpart1`) data. The

Language	Test set	Error (%)			
		Sub	Del	Ins	Tot
Georgian	<code>dev</code>	26.7	10.3	2.9	39.9
	<code>evalpart1</code>	22.9	8.0	2.9	33.8
Tagalog	<code>dev</code>	26.4	9.9	4.3	40.6
	<code>evalpart1</code>	25.6	9.4	4.2	39.2
Swahili	<code>dev</code>	24.1	8.3	3.1	35.5
	<code>evalpart1</code>	23.4	8.2	3.4	35.0

Table 2. Speech recognition performance on narrow-band development and evaluation data

total error rate, WER, in these limited resource conditions is high, ranging between 35 and 40%. For all languages the development data provides a good match to the evaluation data. In particular, the Georgian evaluation data appears to be significantly easier.

Two approaches for confidence estimation are examined in this work: baseline decision trees (DT) and bidirectional recurrent neural networks (BiRNN) in Figure 3 (a). The continuous bag of words representation was chosen to train 50-dimensional word embeddings on all available text data using `fastText` library [36]. These embeddings provided the majority of input features in equation (20). The BiRNN made use of single layer 64-dimensional LSTM units to model the past and future information. The development sets were used to train BiRNNs parameters using `TensorFlow` library [37]. Table 3 contrasts performances of the original confidence scores as given by the recogniser (CN), after decision tree mapping (DT) and after BiRNN estimation. As the proportion of positive examples is not high the ROC curve was used to compute areas under the curve (AUC). The CN and DT columns show that decision trees have significantly reduced cross-entropies between true and estimated confidences thus leading to lower normalised cross-entropy (NCE) values. However, the AUC values have not changed due to monotonicity of the learned mappings incapable of changing the rank ordering. The last column shows that BiRNNs examined in this work provide small but consistent gains in the NCE and AUC values.

For confidence and deletion prediction the topology of BiRNN is adjusted according to Figure 3 (b). Apart from doubling dimensionality of LSTM units no other changes were made. For consis-

Language	Metric	Confidence Estimation		
		CN	DT	BiRNN
Georgian	NCE	-0.376	0.256	0.282
	AUC	0.831	0.831	0.847
Tagalog	NCE	-0.428	0.219	0.249
	AUC	0.810	0.810	0.825
Swahili	NCE	-0.297	0.372	0.392
	AUC	0.832	0.832	0.846

Table 3. Normalised cross-entropy and area under the ROC curve performance on narrow-band evaluation data

tency with the previous results Table 4 shows areas under the ROC curve. The first block repeats the results in Table 3. The first line in

Predictor	Prediction	Language		
		Swahili	Tagalog	Georgian
Cor/Inc	Cor/Inc	0.846	0.825	0.847
+Del	Cor/Inc	0.846	0.826	0.847
	Next Del	0.742	0.740	0.726
	Start Del	0.746	0.646	0.644

Table 4. Area under the ROC curve performance of BiRNNs for correct/incorrect, sentence start and next word deletion prediction on narrow-band evaluation data

the second block shows that the AUC performance remained largely unaffected from the introduction of two additional tasks. The next line shows that predicting whether a deletion occurred before the following word is challenging with AUC values going down by as much as 0.1. The last line shows that, apart from Swahili, predicting that a deletion occurred prior to the first hypothesised word is even more challenging. Both are believed to be the consequence of small number of positive examples and limited features used in this work.

The next experiment examined whether the AUC values in Table 4 are high enough to enable untranscribed data selection with desirable error characteristics. Of particular interest in this work are transcriptions with low numbers of deletions to counter-balance the impact of domain mismatch between the narrow-band CTS Babel and wide-band news and topical broadcast Material data. For this investigation only Swahili was used. Two data selection schemes described in Section 3 were investigated. The discounting coefficients in equations (8)–(9) were estimated using a grid search. For simplicity a single discounting coefficient was estimated $\theta_d = \theta_s = 5$. For the thresholding scheme in equations (10)–(12) the optimal thresholds were found by minimising mean squared error between the true WER and its estimate in equation (13) using grid search. This yielded thresholds for correct/incorrect hypotheses $\theta_c = 0.528$, deletion of the following word $\theta_d = 0.941$ and deletion prior to the first hypothesised word $\theta_s = 0.043$. Figure 4 shows how the number of deletions varies with the amount of data selected. Both schemes achieve lower deletion numbers for all quantities of the selected data with the thresholding scheme showing lowest numbers. Figure 5 shows the impact these schemes have on WER. The thresholding scheme largely follows the profile of the decision tree approach whereas the discounting scheme leads to higher WERs. As the parameters in these schemes were estimated on narrow-band data it is interesting to examine whether the estimated values would generalise to other domains. The equivalent of Figure 4 for the Material wide-band data is shown in Figure 6. This figure shows that the




Fig. 4. Impact of data selection schemes on deletion rate on narrow-band evaluation data

thresholding scheme generalised better to the wide-band data.

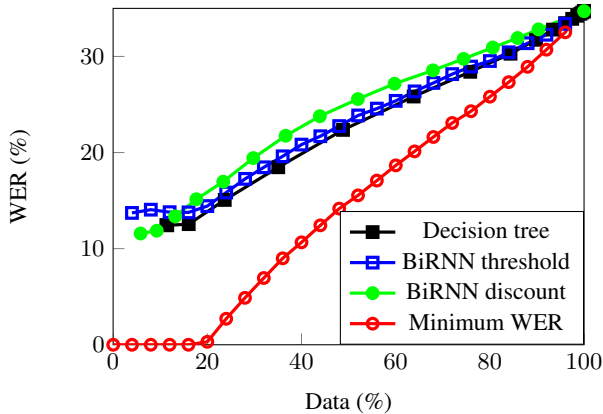


Fig. 5. Impact of data selection approaches on total word error rate on narrow-band evaluation

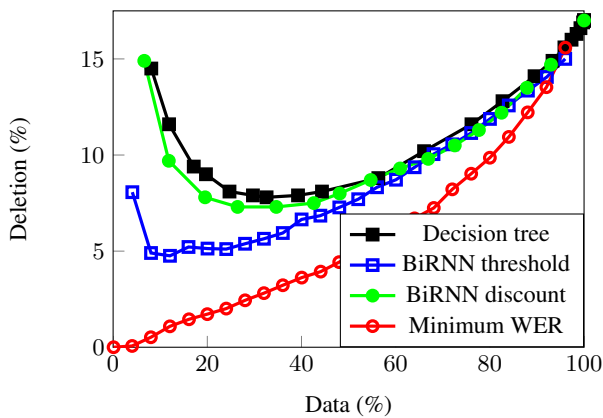


Fig. 6. Impact of data selection schemes on deletion error rate on Material wide-band data

The final experiment examined whether any of these schemes

could lead to WER improvements. As there is only general knowledge of the Material wide-band data domain composition (news and topical broadcast) a guess had to be made which of many possible Swahili wide-band audio files available on the internet to use. A subset of untranscribed Voice of America (Africa) data totalling 500 hours was selected and transcribed by the narrow-band recogniser from Table 2. A quarter of that data was chosen by each scheme to train a wide-band acoustic model. According to Figures 4 and 6 this should yield transcriptions with lower numbers of deletions. The WER performance of these models is compared in Table 5. The first

Selection	Error (%)			
	Sub	Del	Ins	Tot
narrow-band	23.4	19.0	1.3	43.7
confidence	15.6	25.5	0.9	42.0
threshold	16.5	25.2	1.1	42.8
discount	16.9	23.7	1.1	41.7

Table 5. Word error rate performance of data selection schemes on Material wide-band data

line is the narrow-band recogniser. The next three lines shows the standard confidence-based scheme and BiRNN-based thresholding and discounting schemes. The standard confidence-based selection yields gains over the narrow-band recogniser although this comes with a significant increase in the number of deletion errors such that they become the dominant source of errors. The thresholding scheme shows worse generalisation from the VOA data compared to the standard confidence-based approach. This may be attributed to the mismatch between the VOA and Material data. The final line shows that the discounting scheme shows better generalisation. This may be attributed to a lesser degree of divergence between the curves of the decision tree and the discounting scheme in Figures 4-6 than it is for the thresholding scheme. Both schemes show better deletion error characteristics. Overall these results show promise in confidence and deletion-based data selection schemes and highlight challenges faced in building general purpose wide-band recognisers without any available transcribed data.

6. CONCLUSIONS

This paper examined the problem of assigning accurate confidence scores to hypothesised words and predicting deletions produced by speech recognisers. The former, confidence estimation, problem has been researched before and a number of approaches have been proposed. This paper has examined bidirectional recurrent neural networks (BiRNN) - a powerful approach capable of incorporating information from both past and future words for predicting confidence in any given word. The latter, deletion prediction, problem has been researched less even though deletions are an important type of error that is hard to deal with in downstream and upstream tasks. This paper proposed a modification to the standard topology of BiRNNs to enable deletion prediction. The performance of BiRNNs was examined in challenging limited resource conditions across 3 different languages. The confidence prediction was found to yield better performance metrics than deletion prediction possibly due to the limited number of positive examples and features used in this exploration. The combination of the two predictions has been examined for unsupervised data selection in a highly mismatched domain. A simple discounting approach was found to provide small gains over the standard average frame-weighted confidence-based approach that does not take deletions into account.

7. REFERENCES

- [1] B. Li, T. N. Sainath, A. Narayanan, J. Caroselli, M. Bacchiani, A. Misra, I. Shafran, H. Sak, G. Punduk, K. Chin, K. C. Sim, R. J. Weiss, K. W. Wilson, E. Variiani, C. Kim, O. Siohan, M. Weintraub, E. McDermott, R. Rose, and M. Shannon, “Acoustic modeling for Google Home,” in *Interspeech*, 2017.
- [2] A. Ram, R. Prasad, C. Khatri, A. Venkatesh, R. Gabriel, Q. Liu, J. Nunn, B. Hedayatnia, M. Cheng, A. Nagar, E. King, K. Bland, A. Wartick, Y. Pan, H. Song, S. Jayadevan, G. Hwang, and A. Pettigree, “Conversational AI: The science behind the Alexa prize,” in *arXiv preprint arXiv:1801.03604*, 2018.
- [3] S. J. Young, “Still talking to machines (cognitively speaking),” in *Interspeech*, 2010.
- [4] G. Saon, T. Sercu, S. Rennie, and J. H.-K. Kuo, “The IBM 2016 English conversational telephone speech recognition system,” in *Interspeech*, 2016.
- [5] W. Xiong, J. Droppo, X. Huang, F. Seide, M. Seltzer, A. Stolcke, D. Yu, and G. Zweig, “The Microsoft 2016 conversational speech recognition system,” in *ICASSP*, 2017.
- [6] T. Alumäe, D. Karakos, W. Hartmann, R. Hsiao, L. Zhang, L. Nguyen, S. Tsakalidis, and R. Schwartz, “The 2016 BBN Georgian telephone speech keyword spotting system,” in *ICASSP*, 2017.
- [7] K. M. Knill, M. J. F. Gales, K. Kyriakopoulos, A. Malinin, A. Ragni, Y. Wang, and A. Caines, “Impact of ASR performance on free speaking language assessment,” in *Interspeech*, 2018.
- [8] H. Jiang, “Confidence measures for speech recognition,” *Speech Communication*, vol. 45, pp. 455–470, 2005.
- [9] G. Evermann and P. C. Woodland, “Large vocabulary decoding and confidence estimation using word posterior probabilities,” in *ICASSP*, 2000, vol. 3, pp. 1655–1658.
- [10] L. Gillick, Y. Ito, and J. Young, “A probabilistic approach to confidence estimation and evaluation,” in *ICASSP*, 1997.
- [11] M. Weintraub, F. Beaufays, Z. Rivlin, Y. Konig, and A. Stolcke, “Neural-network based measures of confidence for word recognition,” in *ICASSP*, 1997.
- [12] M. S. Seigel and P. C. Woodland, “Combining information sources for confidence estimation with CRF models,” in *Interspeech*, 2011.
- [13] C. Sutton and A. McCallum, “An introduction to conditional random fields,” *Foundations and Trends in Machine Learning*, vol. 4, no. 4, pp. 267–373, 2010.
- [14] K. Kalgaonkar, C. Liu, Y. Gong, and K. Yao, “Estimating confidence scores on ASR results using recurrent neural networks,” in *ICASSP*, 2015.
- [15] M. A. Del-Agua, A. Gimenez, A. Sanchis, J. Civera, and A. Juan, “Speaker-adapted confidence measures for ASR using deep bidirectional recurrent neural networks,” *IEEE/ACM TASLP*, vol. 26, no. 7, pp. 1198–1206, 2018.
- [16] A. Ragni and M. J. F. Gales, “Automatic speech recognition system development in the “wild,”” in *Interspeech*, 2018.
- [17] B. Strope, D. Beeferman, A. Gruenstein, and X. Lei, “Unsupervised testing strategies for ASR,” in *Interspeech*, 2011.
- [18] H. Y. Chan and P. C. Woodland, “Improving broadcast news transcription by lightly supervised discriminative training,” in *ICASSP*, 2004.
- [19] J. Ma, S. Matsoukas, O. Kimball, and R. Schwartz, “Unsupervised training on large amounts of broadcast news data,” in *ICASSP*, 2006.
- [20] M. S. Seigel and P. C. Woodland, “Detecting deletions in ASR output,” in *ICASSP*, 2014.
- [21] M. Schuster and K. K. Paliwal, “Bidirectional recurrent neural networks,” *IEEE Tran Sig Proc*, 2010.
- [22] T. Kemp and T. Schaaf, “Estimating confidence using word lattices,” in *Eurospeech*, 1997.
- [23] F. Wessel, K. Macherey, and R. Schlüter, “Using word probabilities as confidence measures,” in *ICASSP*, 1998.
- [24] T. Schaaf and T. Kemp, “Confidence measures for spontaneous speech recognition,” in *ICASSP*, 1997.
- [25] M.-H. Siu, H. Gish, and F. Richardson, “Improved estimation, evaluation and applications of confidence measures for speech recognition,” in *Eurospeech*, 1997.
- [26] J. Davis and M. Goadrich, “The relationship between precision-recall and ROC curves,” in *ICML*, 2006.
- [27] S. J. Young, G. Evermann, M. J. F. Gales, T. Hain, D. Kershaw, X. Liu, G. Moore, J. Odell, D. Ollason, D. Povey, A. Ragni, V. Valtchev, P. C. Woodland, and C. Zhang, *The HTK Book (for HTK Version 3.5)*, University of Cambridge, <http://htk.eng.cam.ac.uk>, 2015.
- [28] J. Ma and S. Matsoukas, “Unsupervised training on a large amount of Arabic broadcast news data,” in *ICASSP*, 2007.
- [29] V. Peddinti, Y. Wang, D. Povey, and S. Khudanpur, “Low latency acoustic modeling using temporal convolution and LSTMs,” *Signal Processing Letters*, 2017.
- [30] D. Povey, A. Ghoshal, G. Boulianne, L. Burget, O. Glembek, N. Goel, M. Hannemann, P. Motlicek, Y. Qian, P. Schwarz, J. Silovsky, G. Stemmer, and K. Vesely, “The Kaldi speech recognition toolkit,” in *ASRU*, 2011.
- [31] A. Stolcke, “SRILM - an extensible language modeling toolkit,” in *ICSLP*, 2002.
- [32] P. Ghahremani, B. BabaAli, D. Povey, K. Riedhammer, J. Trmal, and S. Khudanpur, “A pitch extraction algorithm tuned for automatic speech recognition,” in *ICASSP*, 2014.
- [33] G. Mendels, E. Cooper, V. Soto, J. Hirschberg, M. Gales, K. Knill, A. Ragni, and H. Wang, “Improving speech recognition and keyword search for low resource languages using web data,” in *Interspeech*, 2015.
- [34] L. Zhang, D. Karakos, W. Hartmann, R. Hsiao, R. Schwartz, and S. Tsakalidis, “Enhancing low resource keyword spotting with automatically retrieved web documents,” in *Interspeech*, 2015.
- [35] M. Bisani and H. Ney, “Joint-sequence models for grapheme-to-phoneme conversion,” *Speech Communication*, 2008.
- [36] P. Bojanowski, E. Grave, A. Joulin, and T. Mikolov, “Enriching word vectors with subword information,” 2016.
- [37] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia,

R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mane, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viegas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng, “TensorFlow: large-scale machine learning on heterogeneous distributed systems,” in *arXiv preprint arXiv:1603.04467*, 2015.