

A Corpus for Modeling Word Importance in Spoken Dialogue Transcripts

Sushant Kafle, Matt Huenerfauth

Rochester Institute of Technology
1 Lomb Memorial Drive, Rochester, NY
sushant@mail.rit.edu, matt.huenerfauth@rit.edu

Abstract

Motivated by a project to create a system for people who are deaf or hard-of-hearing that would use automatic speech recognition (ASR) to produce real-time text captions of spoken English during in-person meetings with hearing individuals, we have augmented a transcript of the Switchboard conversational dialogue corpus with an overlay of word-importance annotations, with a numeric score for each word, to indicate its importance to the meaning of each dialogue turn. Further, we demonstrate the utility of this corpus by training an automatic word importance labeling model; our best performing model has an F-score of 0.60 in an ordinal 6-class word-importance classification task with an agreement (concordance correlation coefficient) of 0.839 with the human annotators (agreement score between annotators is 0.89). Finally, we discuss our intended future applications of this resource, particularly for the task of evaluating ASR performance, i.e. creating metrics that predict ASR-output caption text usability for DHH users better than Word Error Rate (WER).

Keywords: Word Importance Annotation, Speech Recognition Evaluation, Word Importance Prediction

1. Introduction

There has been increasing interest among researchers of speech and language technology applications to identify the importance of individual words, for the overall meaning of the text. Depending on the context of how the importance of a word is defined, this task has found use in varieties of applications such as text summarization (Hong and Nenkova, 2014; Yih et al., 2007), text classification (Sheikh et al., 2016), or speech synthesis (Mishra et al., 2007).

Our laboratory is currently designing a system to benefit people who are deaf or hard-of-hearing (DHH) who are engaged in a live meeting with hearing colleagues. In many settings, sign language interpreting or professional captioning (where a human types the speech, displayed as text on a screen for the user), are unavailable, e.g. in impromptu conversations in the workplace. A system that uses automatic speech recognition (ASR) to generate captions in real-time could display this text on mobile devices for DHH users, but text output from ASR systems inevitably contains errors. Thus, we were motivated to understand which words in the text were most important to the overall meaning, to inform our evaluation of ASR accuracy for this task.

In this paper, we present a word-importance annotation of transcripts of the Switchboard corpus (Godfrey et al., 1992). While our overall goal is to produce measures of ASR accuracy for our caption application; to demonstrate the use of this corpus, in this paper, we present models that predict word-importance in spoken dialogue transcripts.

1.1. ASR Evaluation

ASR researchers generally report the performance of their systems using a metric called Word Error Rate (WER). The metric considers the number of errors in the output of the ASR system, normalized by the number of words human actually said in the audio recording. While WER has been the most commonly used intrinsic measure for the evaluation of ASR, there have been criticisms of WER (McCowan et al., 2004; Morris et al., 2004), and several re-

searchers have recommended alternative measures to better predict human task-performance in applications that depend on ASR (Garofolo et al., 1999; Mishra et al., 2011; Kafle and Huenerfauth, 2016).

Among these newly proposed metrics, a common theme has been: rather than simply counting the number of errors, it would be better to consider the importance of the individual words that are incorrect - suggesting that it would be better to more heavily penalize systems that make errors on words that are important (with the definition of importance based on the specific application or task). This approach of penalizing errors differentially has been shown to be useful in various application settings, e.g. in our research for DHH users, we have found that an evaluation metric designed for predicting the usability of an ASR-generated transcription as a caption text for these users could benefit from word importance information (Kafle and Huenerfauth, 2017). However, estimating the importance of a word has been challenging for our team thus far, because we have lacked corpora of conversational dialogue with word-importance annotation, for training a word-importance model.

1.2. Word Importance Estimation

Prior research on identifying and scoring important words in a text has largely focused on the task of keyword extraction, which involves identifying a set of descriptive words in a document that serves as a dense summary of the document. Several automatic keyword extraction techniques have been investigated over the years, including unsupervised methods using, e.g. Term Frequency x Inverse Document Frequency (TF-IDF) weighting (HaCohen-Kerner et al., 2005), word co-occurrence probability estimation (Matsuo and Ishizuka, 2004) – as well as supervised methods that leverage various linguistic features from text to achieve strong predictive performance (Liu et al., 2011; Liu et al., 2004; Hulth, 2003; Sheeba and Vivekanandan, 2012).

While this conceptualization of word importance as a keyword-extraction problem has led to positive results in

the field of text summarization (Litvak and Last, 2008; Wan et al., 2007; Hong and Nenkova, 2014), this approach may not generalize to other applications. For instance, given the sometimes meandering nature of topic transition in spontaneous speech dialogue (Sheeba and Vivekanandan, 2012), applications that process transcripts of such dialogue may benefit from a model of word importance that is more local, i.e. based on the importance of a word at sentential, utterance, or local dialogue level, rather than at a document-level. Furthermore, the dyadic nature of dialogue, with interleaved contributions from multiple speakers, may require special consideration when evaluating word importance. In this paper, we present a corpus with annotation of word importance that could be used to support research into these complex issues.

2. Defining Word Importance

In eye-tracking studies of reading behavior, researchers have found that readers rarely glance at every word in a text sequentially: Instead, they sometimes regress (glance back at previous words), re-fixate on a word, or skip words entirely (Rayner, 1998). This research supports the premise that some words are of higher importance than others, for readers. Analyses of eye-tracking recordings have revealed a relationship between these eye-movement behaviors and various linguistic features, e.g. word length or word predictability. In general, readers’ gaze often skips over words that are shorter or more predictable (Rayner et al., 2011).

While eye-tracking suggests some features that may relate to readers’ judgments of word importance, at least as expressed through their choice of eye fixations, we needed to develop a specific definition of word importance in order to develop annotation guidelines for our study. Rather than ask annotators to consider specific features, e.g. word length, which may pre-suppose a particular model, we instead took a functional perspective, with our application domain in mind. That is, we define word importance for spontaneous spoken conversation as the degree to which a reader of a transcript of the dialogue would be unable to understand the overall meaning of a conversational utterance (a single turn of dialogue) if that word had been “dropped” or omitted from the transcript. This definition underlies our annotation scheme (in 3 1) and suits our target application, i.e. evaluating ASR for real-time captioning of meetings.

In addition, for our annotation project, we defined word-importance as a single-dimensional property, which could be expressed on a continuous scale from 0.0 (not important at all to the meaning of the utterance) to 1.0 (very important). Figure 1 illustrates how numerical importance scores can be assigned to words in a sentence – in fact, this figure displays actual scores assigned by a human annotator working on our project. Of course, asking human annotators to assign specific numerical scores to quantify the importance of a word is not straightforward. In later sections, we discuss how we attempt to overcome the subjective nature of this task, to promote consistency between annotators, as we developed this annotated resource (see Section 3 1). Section 4 characterizes the level of agreement between our annotators on this task.

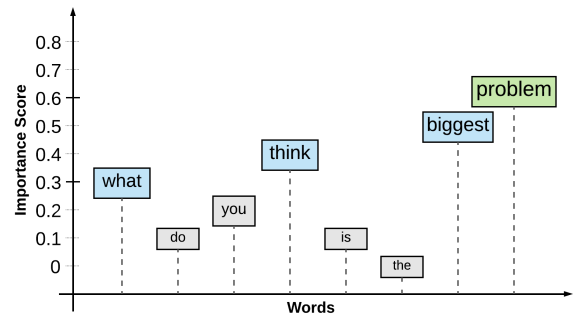


Figure 1: Visualization of importance scores assigned to words in a sentence by a human annotator on our project, with the height and font-size of words indicating their importance score (and redundant color coding: green for high-importance words with score above 0.6, blue for words with score between 0.3 and 0.6, and grey otherwise).

3. Corpus Annotation

The Switchboard corpus consists of audio recordings of approximately 260 hours of speech consisting of about 2,400 two-sided telephone conversations among 543 speakers (302 male, 241 female) from across the United States (Godfrey et al., 1992). In January 2003, the Institute for Signal and Information Processing (ISIP) released written transcripts for the entire corpus, which consists of nearly 400,000 conversational turns. The ISIP transcripts include a complete lexicon list and automatic word alignment timing corresponding to the original audio files¹.

In our project, a pair of annotators have assigned word-importance scores to these transcripts. As of September 2017, they have annotated over 25,000 tokens, with the overlap of approximately 3,100 tokens. With this paper, we announce the release² of these annotations as a set of supplementary files, aligned to the ISIP transcripts. Our annotation work continues, and we aim to annotate all of the Switchboard corpus and with a larger group of annotators.

3.1. Annotation Scheme

To reduce the cognitive load on annotators and to promote consistency, we created the following annotation scheme:

Range and Constraints. Each word is assigned a numeric score between $[0, 1]$, where 1 indicates a high importance score; the numeric score has the precision of 0.05. Importance scores are not meant to indicate an absolute proportion of the utterance’s meaning represented by each word, i.e. the scores do not have to sum to 1.

Methodology. Given an utterance (a speaker’s single turn in the conversation), the annotator first considers the overall meaning conveyed by the utterance, with the help of the previous conversation history (if available). The annotator then scores each word based on its (direct or indirect) contribution to the utterance’s meaning, using the rubric described in the Interpretation and Scoring section below.

¹<https://www.isip.piconepress.com/projects/switchboard/>

²<http://latlab.ist.rit.edu/lrec2018>

Range	Description
[0 - 0.3)	Words that are of least importance - these words can be easily omitted from the text without much consequence.
[0.3 - 0.6)	Words that are fairly important - omitting these words will take away some important details from the utterance.
[0.6 - 1]	Words that are of high importance - omitting these words will change the message of the utterance quite significantly.

Table 1: Guidance for the annotators to promote consistency and uniformity in the use of numerical scores.

Rating Scheme. To help annotators calibrate their scores, Table 1 provides some recommendations for how to select word-importance scores in various numerical ranges.

Interpretation and Scoring. Annotators should consider how their understanding of the utterance would be affected if this word had been “dropped,” i.e. replaced with a blank space (“_____”). Since these are conversations between pairs of speakers, annotators should consider how much the other person in the conversation would have difficulty understanding the speaker’s message if that word had been omitted, i.e. if they had not heard that word intelligibly.

4. Inter-annotator Agreement

There were 3,100 tokens in our “overlap” set, i.e. the subset of transcripts independently labeled by both annotators. This set was used as the basis for calculating inter-annotator agreement. Since scores were nearly continuous (ranges [0,1] with a precision of 0.05), we computed the Concordance Correlation Coefficient (ρ_c), also known as Lin’s concordance correlation coefficient, as our primary metric for measuring the agreement between the annotators. This metric indicates how well a new test or measurement (X) reproduces a gold standard or measure (Y). Considering the annotations from one annotator as a gold standard, we can generalize this measure to compute the agreement between two annotators. Like other correlation coefficients, ρ_c also ranges from -1 to 1; 1 being the score of perfect agreement. Concordance between the two measures can be characterized by the expected value of their squared difference as:

$$E[(Y - X)^2] = (\mu_y - \mu_x)^2 + \sigma_x^2 + \sigma_y^2 - 2\rho\sigma_x\sigma_y \quad (1)$$

where, ρ is the correlation coefficient, μ_x and μ_y are the means of the population of the variables X and Y , and σ_x and σ_y are their standard deviation. The expectation score coefficient (between -1 and 1) is calculated as follows:

$$\rho_c = \frac{2\rho S_x S_y}{(\bar{Y} - \bar{X})^2 + S_x^2 + S_y^2} \quad (2)$$

where, ρ_c is the correlation coefficient, \bar{X} and \bar{Y} are the mean of X and Y , and S_x and S_y are standard deviations.

We obtained an agreement score (ρ_c) of 0.89 between our annotators, which we interpret as an acceptable level of

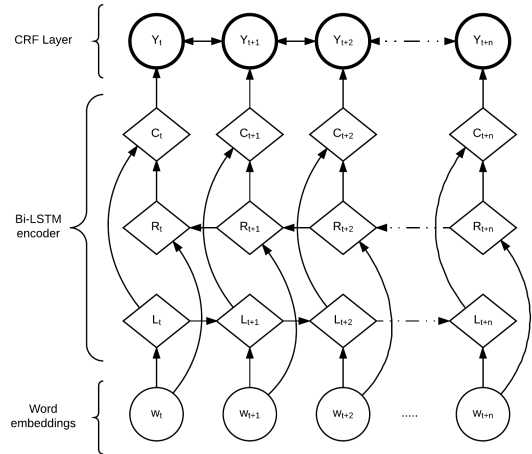


Figure 2: =

[General unfolded network structure of our model, adapted from (Lample et al., 2016). The bottom layer represents word-embedding inputs, passed to bi-directional LSTM layers above. Each LSTM takes as input the hidden state from the previous time step and word embeddings from the current step, and outputs a new hidden state. C_i concatenates hidden representations from LSTMs (L_i and R_i) to represent the word at time i in its context.]

agreement, given the subjective nature of the task of quantifying word importance in spoken dialogue transcripts.

5. Automatic Prediction

To demonstrate the use of this corpus, we trained a prediction model, by adopting the neural architecture described in (Lample et al., 2016) consisting of bidirectional LSTM encoders with a sequential Conditional Random Field (CRF) layer on top. Our input word tokens were first mapped to a sequence of pre-trained distributed embeddings (Pennington et al., 2014) and then combined with the learned character-based word representations to get the final word representation. As shown in Figure 2, the bidirectional LSTM encoders are used to create a context-aware representation of each word. The hidden representations from each LSTM were concatenated to obtain a final representation, conditioned on the whole sentence. The CRF layer uses this representation to look for the most optimal state (Y) sequence through all the possible state configurations.

The neural framework was implemented using Tensorflow, and the code is publicly available³. The word embeddings were initialized with publicly available pre-trained glove vectors (Pennington et al., 2014). The embeddings for characters were set to length 100 and were initialized randomly. The LSTM layer size was set to 300 in each direction for word- and 100 for character-level components. Parameters were optimized using the Adam (Kingma and Ba, 2015) optimizer, with the learning rate initialized at 0.001 with a decay rate of 0.9, and sentences were grouped into batches

³<https://github.com/SushantKafle/spechtext-wimp-labeler>




Figure 3: Confusion matrices for each model for classification into 6 classes: $c_1 = [0, 0.1)$, $c_2 = [0.1, 0.3)$, and so forth.

of size 20. We applied a dropout with a probability of 0.5 during training on word embeddings.

We investigated two variations of this model: (i) a bidirectional LSTM model with sequential CRF layer on top (LSTM-CRF) treating the problem as a discrete classification task, (ii) a new bidirectional LSTM model with a sigmoid layer on top (LSTM-SIG) for a continuous prediction. The LSTM-CRF models the prediction task as a classification problem, using a fixed number of non-ordinal class labels. In contrast, the LSTM-SIG model provides a continuous prediction, using a sigmoid nonlinearity to bound the prediction scores between 0 and 1. Using a square loss, we train this model to directly learn to predict the annotation scores, similar to a regression task.

5.1. Evaluation and Discussion

Partitioning our corpus as 80% training, 10% development, and 10% test sets, we evaluated our model using two measures: (i) total root mean square error (RMS) - the deviation of the model predictions from the human-annotations and, (ii) F_1 measure in a classification task - the ability of the model to predict human-annotations categorized into a group of classes. To evaluate performance in terms of classification, we discretized annotation scores into 6 classes: $[0, 0.1)$, $[0.1, 0.3)$, $[0.3, 0.5)$, $[0.5, 0.7)$, $[0.7, 0.9)$, $[0.9, 1]$. Table 2 summarizes the performance of our models on the test set, presenting average scores for 5 different configurations, to compensate for outlier results due to randomness in model initialization. While the LSTM-CRF had a better (higher) F-score on the classification task, its RMS score was worse (higher) than the LSTM-SIG model, which may be due to the limitation of the model as discussed in Section 5.

Model	RMS	F_1 (macro)
LSTM-CRF	0.154	0.60
LSTM-SIG	0.120	0.519

Table 2: Model performance in terms of RMS deviation and macro-averaged F_1 score, with best results in **bold** font.

Confusion matrices in Figure 3 provide a more detailed view of the classification performance of each model. Since the LSTM-SIG was trained to optimize the accuracy of its continuous predictions, rather than its discrete assignment of instances to classes, it is not surprising to see a “wider diagonal” in the confusion matrix in Figure 3(b), which indicates that the LSTM-SIG model was more likely to misclassify words using ordinally adjacent classes. The figure illustrates that both models were worse at classifying words with importance scores in the middle range $[0.3, 0.7)$.

Treating our human-annotations as ground truth, we also computed the concordance correlation coefficient to measure the agreement between the human annotation and each model. The average correlation between the human annotator and the LSTM-CRF model was higher ($\rho_c = 0.839$), as compared to the LSTM-SIG model ($\rho_c = 0.826$).

6. Conclusions and Future Work

We have presented a new collection of annotation of transcripts of the Switchboard conversational speech corpus, produced through human annotation of the importance of individual words to the meaning of each utterance. We have demonstrated the use of this data by training word-importance prediction models, with the best model achieving an F_1 score of 0.60 and model-human agreement correlation of 0.839. In future work, we will collect additional human annotations for additional sections of the corpus. This research is part of a project on the use of ASR to provide real-time captions of speech for DHH individuals during meetings, and we plan to incorporate these word-importance models into new word-importance-weighted metrics of ASR accuracy, to better predict the usability of ASR-produced captions for these users.

7. Acknowledgement

This material was based on work supported by the National Technical Institute for the Deaf (NTID). We thank Tomomi Takeuchi and Michael Berezny for their contributions.

8. Bibliographical References

- Garofolo, J. S., Auzanne, C. G. P., and Voorhees, E. M. (1999). The TREC spoken document retrieval track: A success story. In Ellen M. Voorhees et al., editors, *Proceedings of The Eighth Text REtrieval Conference, TREC 1999, Gaithersburg, Maryland, USA, November 17-19, 1999*, volume Special Publication 500-246. National Institute of Standards and Technology (NIST).
- Godfrey, J. J., Holliman, E. C., and McDaniel, J. (1992). Switchboard: Telephone speech corpus for research and development. In *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP-92)*, volume 1, pages 517–520. IEEE.
- HaCohen-Kerner, Y., Gross, Z., and Masa, A. (2005). Automatic extraction and learning of keyphrases from scientific articles. In Alexander F. Gelbukh, editor, *Computational Linguistics and Intelligent Text Processing, 6th International Conference, CICLing 2005, Mexico City, Mexico, February 13-19, 2005, Proceedings*, volume 3406 of *Lecture Notes in Computer Science*, pages 657–669. Springer.
- Hong, K. and Nenkova, A. (2014). Improving the estimation of word importance for news multi-document summarization. In Gosse Bouma et al., editors, *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics, EACL 2014, April 26-30, 2014, Gothenburg, Sweden*, pages 712–721. The Association for Computer Linguistics.
- Hulth, A. (2003). Improved automatic keyword extraction given more linguistic knowledge. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing, EMNLP 2003, Sapporo, Japan, July 11-12, 2003*.
- Kafle, S. and Huenerfauth, M. (2016). Effect of speech recognition errors in text understandability for people who are deaf or hard-of-hearing. In *Proceedings of the 7th Workshop on Speech and Language Processing for Assistive Technologies (SLPAT)*. Interspeech.
- Kafle, S. and Huenerfauth, M. (2017). Evaluating the usability of automatically generated captions for people who are deaf or hard of hearing. In Amy Hurst, et al., editors, *Proceedings of the 19th International ACM SIGACCESS Conference on Computers and Accessibility, ASSETS 2017, Baltimore, MD, USA, October 29 - November 01, 2017*, pages 165–174. ACM.
- Kingma, D. P. and Ba, J. (2015). Adam: A method for stochastic optimization. In Yoshua Bengio et al., editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Lample, G., Ballesteros, M., Subramanian, S., Kawakami, K., and Dyer, C. (2016). Neural architectures for named entity recognition. In Kevin Knight, et al., editors, *NAACL HLT 2016, The 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, San Diego California, USA, June 12-17, 2016*, pages 260–270. The Association for Computational Linguistics.
- Litvak, M. and Last, M. (2008). Graph-based keyword extraction for single-document summarization. In *Proceedings of the Workshop on Multi-source Multilingual Information Extraction and Summarization, MMIES '08*, pages 17–24, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Liu, B., Li, X., Lee, W. S., and Yu, P. S. (2004). Text classification by labeling words. In Deborah L. McGuinness et al., editors, *Proceedings of the Nineteenth National Conference on Artificial Intelligence, Sixteenth Conference on Innovative Applications of Artificial Intelligence, July 25-29, 2004, San Jose, California, USA*, pages 425–430. AAAI Press / The MIT Press.
- Liu, F., Liu, F., and Liu, Y. (2011). A supervised framework for keyword extraction from meeting transcripts. *IEEE Trans. Audio, Speech & Language Processing*, 19(3):538–548.
- Matsuo, Y. and Ishizuka, M. (2004). Keyword extraction from a single document using word co-occurrence statistical information. *International Journal on Artificial Intelligence Tools*, 13(1):157–169.
- McCowan, I. A., Moore, D., Dines, J., Gatica-Perez, D., Flynn, M., Wellner, P., and Boulard, H. (2004). On the use of information retrieval measures for speech recognition evaluation. Technical report, IDIAP.
- Mishra, T., Prud'hommeaux, E. T., and van Santen, J. P. H. (2007). Word accentuation prediction using a neural net classifier. In Petra Wagner, et al., editors, *Sixth ISCA Workshop on Speech Synthesis, Bonn, Germany, August 22-24, 2007*, pages 246–251. ISCA.
- Mishra, T., Ljolje, A., and Gilbert, M. (2011). Predicting human perceived accuracy of ASR systems. In *INTER-SPEECH 2011, 12th Annual Conference of the International Speech Communication Association, Florence, Italy, August 27-31, 2011*, pages 1945–1948. ISCA.
- Morris, A. C., Maier, V., and Green, P. D. (2004). From WER and RIL to MER and WIL: improved evaluation measures for connected speech recognition. In *INTER-SPEECH 2004 - ICSLP, 8th International Conference on Spoken Language Processing, Jeju Island, Korea, October 4-8, 2004*. ISCA.
- Pennington, J., Socher, R., and Manning, C. D. (2014). Glove: Global vectors for word representation. In Alessandro Moschitti, et al., editors, *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25-29, 2014, Doha, Qatar, A meeting of SIGDAT, a Special Interest Group of the ACL*, pages 1532–1543. ACL.
- Rayner, K., Slattery, T. J., Drieghe, D., and Liversedge, S. P. (2011). Eye movements and word skipping during reading: effects of word length and predictability. *Journal of Experimental Psychology: Human Perception and Performance*, 37(2):514.
- Rayner, K. (1998). Eye movements in reading and information processing: 20 years of research. *Psychological bulletin*, 124(3):372.
- Sheeba, J. and Vivekanandan, K. (2012). Improved keyword and keyphrase extraction from meeting transcripts. *International Journal of Computer Applications*, 52(13).
- Sheikh, I. A., Illina, I., Fohr, D., and Linares, G. (2016).

- Learning word importance with the neural bag-of-words model. In Phil Blunsom, et al., editors, *Proceedings of the 1st Workshop on Representation Learning for NLP, Rep4NLP@ACL 2016, Berlin, Germany, August 11, 2016*, pages 222–229. Association for Computational Linguistics.
- Wan, X., Yang, J., and Xiao, J. (2007). Towards an iterative reinforcement approach for simultaneous document summarization and keyword extraction. In John A. Carroll, et al., editors, *ACL 2007, Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics, June 23-30, 2007, Prague, Czech Republic*. The Association for Computational Linguistics.
- Yih, W., Goodman, J., Vanderwende, L., and Suzuki, H. (2007). Multi-document summarization by maximizing informative content-words. In Manuela M. Veloso, editor, *IJCAI 2007, Proceedings of the 20th International Joint Conference on Artificial Intelligence, Hyderabad, India, January 6-12, 2007*, pages 1776–1782.