

# Induced Inflection-Set Keyword Search in Speech

Oliver Adams,<sup>♡♣</sup> Matthew Wiesner,<sup>♣</sup> Jan Trmal,<sup>♣</sup>  
Garrett Nicolai,<sup>◇♣</sup> David Yarowsky<sup>♣</sup>

<sup>♣</sup>Center for Language and Speech Processing, Johns Hopkins University, USA

<sup>♡</sup>Atos zData, USA

<sup>◇</sup>Department of Linguistics, University of British Columbia, Canada

oliver.adams@gmail.com, gnicolai@mail.ubc.ca,

{mwiesner, yenda, yarowsky}@jhu.edu

## Abstract

We investigate the problem of searching for a lexeme-set in speech by searching for its inflectional variants. Experimental results indicate how lexeme-set search performance changes with the number of hypothesized inflections, while ablation experiments highlight the relative importance of different components in the lexeme-set search pipeline and the value of using curated inflectional paradigms. We provide a recipe and evaluation set for the community to use as an extrinsic measure of the performance of inflection generation approaches.

## 1 Introduction

Keyword search (KWS) is the task of finding certain words or expressions of interest in a body of speech. KWS is relevant to incident-response situations such as those modeled by LORELEI (Strassel and Tracey, 2016) and was a focus of the IARPA Babel Program.<sup>1</sup> In the event of a humanitarian crisis, processing speech to determine mentions of certain keywords can inform better decision making when time is critical.

KWS is typically framed as searching for instances of a keyword in lattices that result from speech recognition decoding, as this means search is not restricted to a potentially incorrect one-best transcription. However, existing work on KWS assumes the relevant form of a keyword has been correctly specified. Many concepts to be searched for in speech take different forms through inflection as a result of the language’s morphosyntax. In most cases, distinctions between such inflections (e.g. kill, kills, killing, killed) are irrelevant to the problem of searching for the underlying concept of interest.

Producing such inflection sets manually is arduous, even for native speakers, yet curators of

<sup>1</sup>[www.iarpa.gov/index.php/research-programs/babel](http://www.iarpa.gov/index.php/research-programs/babel)




Figure 1: An example Zulu keyword lemma (left) is inflected (middle) and then searched for in a corpus of speech (right). ✓ and ✗ indicate correct/incorrect inflections, and correct/incorrect findings of the inflection in the corpus.

keyword lists may have to construct them cross-lingually using bilingual dictionaries, which typically only contain canonical forms. Compounding this issue are the limitations of existing language technology for most of the world’s languages across the whole KWS pipeline, including inflection generation, the language model (LM), the pronunciation lexicon, and the acoustic model.

In this paper we explore the application of inflection generation to KWS by searching for instances of a lexeme (see Figure 1). To the best of our knowledge, this task has not been investigated before. Using Bengali and Turkish as evaluation languages, we scale the number of inflections generated per lexeme-set to examine how the trade-off between false positives and false negatives affects downstream KWS. We additionally perform experiments that assume varying quality of inflection generation. Our findings show that lexeme-set KWS yields promising results even when all inflections must be generated on the basis of a distantly supervised cross-lingual approach to training inflection tools, though we observe that having a curated set of inflectional paradigms is important for achieving good performance. These first results encourage future work for lexeme-set search in speech, and the use of KWS as an extrinsic eval-

uation of inflection generation tools.

To this end, we make available to the community a lexeme-set KWS pipeline with baseline models for inflection generation, grapheme-to-phoneme conversion (G2P), multilingual acoustic modeling, and cross-lingual inflection generation, and a KWS evaluation set built on a suitable intersection of UniMorph inflection sets (Sylak-Glassman et al., 2015; Kirov et al., 2018) and the Babel speech (Andresen et al., 2016; Bills et al., 2016). The combination of these components serves as a novel downstream evaluation of inflection generation approaches, as well the other components in the pipeline. We make this recipe and evaluation set freely available online.<sup>2</sup>

## 2 The lexeme-set KWS Pipeline

The pipeline starts with a lexeme of interest from the evaluation set (§2.1). Inflections of the lexeme are generated using some generation tool or manual resource (§2.2). These inflections are then converted to a phonemic representation (§2.3) before being added to the lexicon used in speech recognition. KWS is then performed (§2.4) by decoding the speech and the model is scored on the how well it finds instances of the lexeme.

### 2.1 Evaluation Set

We evaluate systems based on their ability to complete the following task: given a lemma, find all occurrences of its inflections in speech. To create an evaluation set for this task, we use UniMorph data, which provides ground truth inflection sets for a substantial number of languages. We use as our evaluation set instances of words in the Babel 10h development set that also are inflections in the UniMorph data. We remove from this set a small number of inflections that occur in more than one paradigm, as well as those that don't occur in the Babel pronunciation lexicon. This means that we can use the Babel lexicon as an oracle pronunciation lexicon with respect to our constructed evaluation sets to compare against our other methods. The result is an evaluation set tailored to morphologically salient word forms, with 1250 Turkish paradigms and 59 Bengali paradigms. The set of evaluation languages that can be extended to other languages in the Babel set for which we have ground truth paradigms.

<sup>2</sup><https://github.com/oadams/inflection-kws>

### 2.2 Inflection Generation

Inflection generation is the task of producing an *inflection*, given a lemma and a *bundle* of morphosyntactic features. For example, `run + {PRES;3;SG}`  $\mapsto$  “runs”. The state of the art in inflection generation has arisen from the CoNLL–SIGMORPHON Shared Tasks (Cotterell et al., 2016, 2017, 2018; McCarthy et al., 2019), and typically consists of a modified sequence-to-sequence model with attention (Makarov and Clematide, 2018).

However, these systems are fully supervised, and hand-curated morphological dictionaries often do not exist. We instead turn to the methods of Nicolai and Yarowsky (2019), who use English annotation as distant supervision to induce target language morphology, using a widely-translated, verse-parallel text: the Bible. Starting from the inflection pairs extracted by their method, we ensemble generators trained using an RNN and DirecTL+ (Jiampojarn et al., 2010). For each lemma in the respective UniMorph, we generate hypotheses for each feature bundle, ensembling via a linear combination of confidence scores. This gives us a set of inflections for each of the lexemes in the evaluation set which can then be searched for in the speech.

### 2.3 Grapheme-to-Phoneme Conversion

To include hypothesized inflections in the KWS pipeline, orthographic forms of inflections must be mapped to a phonemic form consistent with the units used by the acoustic model (Maskey et al., 2004; Chen et al., 2016; Mortensen et al., 2018; Schultz et al., 2007; Kominék and Black, 2006; Deri and Knight, 2016; Trmal et al., 2017). We use a finite-state transducer model trained with Phonetisaurus<sup>3</sup> on 5,000 word forms in the target language.

### 2.4 Keyword Search

After generating inflections of lemmas in the evaluation set, these inflections are then included in the lexicon used in KWS. The KWS involves decoding the speech into lattices, and assessing lattice's inclusion of the keyword of interest. Our pipeline builds on the Kaldi OpenKWS system (Trmal et al., 2017), which uses the standard lattice indexing approach of (Can and Saraclar, 2011). We

<sup>3</sup>[github.com/AdolfVonKleist/Phonetisaurus](https://github.com/AdolfVonKleist/Phonetisaurus)

use augmented pronunciation lexicons for KWS, which has been shown to outperform proxy KWS, a popular alternative (Chen et al., 2013).

The novel problem of lexeme-set KWS is related to work on out-of-vocabulary KWS, which has been approached by handling sub-word units such as syllables and morphemes (Trmal et al., 2014; Narasimhan et al., 2014; van Heerden et al., 2017; He et al., 2016). In contrast to KWS with sub-word granularity, our approach is to generate likely full-word inflections given a lemma.

For language modeling, we used a 4-gram modified Kneser-Ney baseline (Kneser and Ney, 1995). We compare using as training data the in-domain Babel text to the Bible, a resource available for many languages, and which was the resource used for cross-lingual distant supervision for inflection generation described in Section 2.2. Hypothesized inflections not seen in the training data receive some probability mass in language model smoothing as is default in SRILM (Stolcke, 2002), the language modeling tool used.

Though a monolingual acoustic model could have been used, we chose to use a “universal” phoneset acoustic model which can effectively be deployed on languages not seen in training and is motivated by work in multilingual acoustic modeling (Schultz and Waibel, 2001; Le and Besacier, 2005; Stolcke et al., 2006; Vesely et al., 2012; Vu et al., 2012; Heigold et al., 2013; Scharenborg et al., 2017; Karafiát et al., 2018). We train an acoustic model on 300 hours of data from 25 languages using a common phonemic representation across languages. The training data includes 10 hours for each of 21 different languages from the IARPA Babel corpus, a 20 hour subset of the Wall Street Journal,<sup>4</sup> Hub4 Spanish Broadcast news,<sup>5</sup> and the Russian and French portions of the Voxforge<sup>6</sup> corpus.

### 3 KWS Evaluation Metrics

We evaluate KWS performance on a per lexeme-set basis, rewarding the system when it finds any form of an evaluation lexeme, regardless of how it is inflected, while also penalizing failure to find any inflection.

As an evaluation metric we use term weighted value (TWV), a standard metric in KWS devel-

oped for the NIST 2006 Spoken Term Detection evaluation (Fiscus et al., 2007), which rewards joint maximization of recall with minimization of false positives. TWV relies on a threshold parameter to determine what minimum level of confidence is required by the system in order to assert keyword findings. There are several variations of term weighted value (TWV) that are different in the way the threshold is handled: Actual (ATWV), Optimum (OTWV), and Supreme (STWV).

ATWV is the TWV of the system given some global threshold (provided by the system) of confidence common to all keywords, and is the most common metric used to compare systems.

OTWV determines a per-keyword (in our case, per lexeme-set) threshold. For our purposes this is the most informative metric because it gives a better sense of how the ATWV would be if system effectively normalized confidences across lexemes. Improvements to TWV may also potentially be made beyond what is represented by the OTWV. Some inflections are more likely than others, yet the thresholds for OTWV are made at a per lexeme-set basis, not a per-inflection basis. Improving how the system weights the likelihood of different inflections (either during inflection generation or in the LM probabilities) would likely substantially improve ATWV.

STWV is a recall-oriented version of TWV that disregards the confidence of the terms and does not penalize false positives. It is thus similar to lattice recall and serves as a useful metric in system analysis for determining whether low ATWV/OTWV is due to large number of false positives or issues in effective speech word lattice decoding.

## 4 Experiments

We conduct experiments to see how performance of KWS relates to the number of inflections hypothesized by the cross-lingual distantly-supervised method described in Section 2.2 (henceforth RNN+DTL), before comparing it to several alternative benchmark methods.

As evaluation languages we used Bengali and Turkish, a subset of languages for which we have Bibles and that also occur in UniMorph. We observed similar trends and relative performance of methods for both languages so in the subsequent results we present the arithmetic mean of the results of Bengali and Turkish.

<sup>4</sup>LDC94S13B

<sup>5</sup>LDC98T29

<sup>6</sup><http://voxforge.org>




Figure 2: Variations on the term weighted value (TWV) metric for different numbers of generated inflections per morphosyntactic bundle (by RNN+DTL). Overgeneration of inflections improves recall, as captured by STWV, but leads to too many false positives when  $k \geq 80$ .

#### 4.1 The Number of Generated Inflections

To gauge how over-generation of inflections affects KWS performance we scaled  $k$ , the number of inflections generated per morphosyntactic bundle. Figure 2 illustrates how TWV varies with respect to  $k$  for inflections generated by RNN+DTL in a KWS system that uses the in-domain Babel LM. At values of  $k$  beyond 40 the ATWV and OTWV began to decrease, as the number of false positives was too high. The recall-oriented STWV continued to increase, peaking at 0.764 when  $k = 160$ . It is noteworthy that OTWV only began to decrease at such a value of  $k$ . For Turkish nouns, with 23 bundles per lexeme-set, a value of  $k = 40$  corresponds to 920 inflections, the vast majority of which are invalid inflections. This indicates that there is room for a substantial amount of inflection overgeneration in KWS, since the speech recognition can provide acoustic evidence against incorrect inflection candidates.

#### 4.2 Comparison of Inflection Approaches

To get a comparative sense of the KWS performance of RNN+DTL at the best value of  $k$ , we compare it with three other approaches: Oracle, UniMorph, and Lemmas, as shown in Table 1.

Oracle includes exactly the set of inflections that occur in the evaluation set. UniMorph includes all the inflections that occur in the UniMorph data, which differs from Oracle in that it contains true inflections that don't happen to occur in the Babel speech. We included this to as-

Inflections	LM	ATWV	OTWV	STWV
Oracle	Babel	0.315	0.463	0.866
UniMorph	Babel	0.392	0.513	0.864
RNN+DTL	Babel	0.133	0.269	0.577
Lemmas	Babel	0.169	0.219	0.281
RNN+DTL-NS	Babel	0.304	0.443	0.815
RNN+DTL	Bible	0.046	0.206	0.561

Table 1: Term weighted value (TWV) under varying conditions: Oracle inflections known to occur in the Babel speech; UniMorph inflections that additionally include true inflections not seen in speech; RNN+DTL-generated inflections via distant cross-lingual supervision; Lemmas-only search. Discounting spurious forms from RNN+DTL shows its high recall. Using an out-of-domain LM substantially decreases performance.

sess how true inflections of the lexeme that are not found in the speech affect performance. It helps substantially for ATWV and OTWV, but not for STWV. This somewhat counterintuitive result suggests that including more inflectional variants of a lexeme may not improve recall (i.e. improve STWV) but can decrease the number of false positives.

Lemmas searches only for citation-form lemmas. It has a relatively decent ATWV (even outperforming RNN+DTL, though not by OTWV) despite low recall (as indicated by STWV) because it has few false positives and also because most inflections sound similar to the lemmas via the addition of an affix. As a result, searching for the lemma often catches inflectional variants too.

We consider two further points of comparison. Firstly, RNN+DTL with only a Bible-trained LM, which underperforms other systems substantially except in lattice recall as indicated by STWV. Secondly, RNN+DTL-NS, which removes from RNN+DTL spurious inflections that weren't found in Oracle. Comparison of RNN+DTL-NS and RNN+DTL demonstrates that while the system has some robustness to overgeneration of inflected forms (§4.1), it is also the case spurious inflections not only increase false positives, but can actually hurt recall too.

These results indicate that correctly generating inflected forms and properly weighting the hypothesized inflections (either via the inflection generation module, or in the language model) is the most critical bottleneck in the pipeline. The high relative performance of UniMorph indicates

the value of making full use of available curated resources. Recent work has shown how effective inflection generation can be with limited resources in the target language (Cotterell et al., 2018). These results suggest that if such resources are not available, then in practice it is likely worth gathering training data with which train an inflection generator.

## 5 Conclusion

We have presented an evaluation of lexeme-set KWS. Our results shed light on the relative impact of undergenerating and overgenerating inflected forms on KWS, indicating that high recall can be achieved via an inflection method of cross-lingual distant supervision, but with the best all-round performance achieved by making use of Unimorph. We release our evaluation set along with scripts to reuse our pipeline so that the community can explore lexeme-set KWS as an extrinsic evaluation of inflection generation.

## Acknowledgements

We would like to thank all reviewers for their constructive feedback.

## References

- Jess Andresen, Aric Bills, Eyal Dubinski, Jonathan G. Fiscus, Breanna Gillies, Mary Harper, T. J. Hazen, Amy Jarrett, Bergul Roomi, Jessica Ray, Anton Rytting, Wade Shen, and Evelyne Tzoukermann. 2016. IARPA Babel Turkish language pack IARPA-babel105b-v0.5. LDC2016S10.
- Aric Bills, Anne David, Eyal Dubinski, Jonathan G. Fiscus, Breanna Gillies, Mary Harper, Amy Jarrett, Mara Encarnacin Prez Molina, Jessica Ray, Anton Rytting, Shelley Paget, Wade Shen, Ronnie Silber, Evelyne Tzoukermann, and Jamie Wong. 2016. IARPA Babel Bengali language pack IARPA-babel103b-v0.4b. LDC2016S08.
- Doğan Can and Murat Saraclar. 2011. Lattice indexing for spoken term detection. *IEEE Transactions on Audio, Speech, and Language Processing*, 19(8):2338–2347.
- Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. 2016. Acoustic data-driven pronunciation lexicon generation for logographic languages. In *Proc. ICASSP*.
- Guoguo Chen, Oguz Yilmaz, Jan Trmal, Daniel Povey, and Sanjeev Khudanpur. 2013. Using proxies for OOV keywords in the keyword search task. In *2013 IEEE Workshop on Automatic Speech Recognition and Understanding*, pages 416–421. IEEE.
- Ryan Cotterell, Christo Kirov, John Sylak-Glassman, Géraldine Walther, Ekaterina Vylomova, Arya D. McCarthy, Katharina Kann, Sebastian Mielke, Garrett Nicolai, Miikka Silfverberg, David Yarowsky, Jason Eisner, and Mans Hulden. 2018. [The CoNLL–SIGMORPHON 2018 shared task: Universal morphological reinflection](#). In *Proceedings of the CoNLL–SIGMORPHON 2018 Shared Task: Universal Morphological Reinflection*, pages 1–27, Brussels. Association for Computational Linguistics.
- Ryan Cotterell, Christo Kirov, John Sylak-Glassman, Géraldine Walther, Ekaterina Vylomova, Patrick Xia, Manaal Faruqui, Sandra Kübler, David Yarowsky, Jason Eisner, and Mans Hulden. 2017. [CoNLL–SIGMORPHON 2017 shared task: Universal morphological reinflection in 52 languages](#). In *Proceedings of the CoNLL SIGMORPHON 2017 Shared Task: Universal Morphological Reinflection*, pages 1–30, Vancouver. Association for Computational Linguistics.
- Ryan Cotterell, Christo Kirov, John Sylak-Glassman, David Yarowsky, Jason Eisner, and Mans Hulden. 2016. [The SIGMORPHON 2016 shared Task—Morphological reinflection](#). In *Proceedings of the 14th SIGMORPHON Workshop on Computational Research in Phonetics, Phonology, and Morphology*, pages 10–22, Berlin, Germany. Association for Computational Linguistics.
- Aliya Deri and Kevin Knight. 2016. [Grapheme-to-phoneme models for \(almost\) any language](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 399–408, Berlin, Germany. Association for Computational Linguistics.
- Jonathan G Fiscus, Jerome Ajot, John S Garofolo, and George Doddington. 2007. Results of the 2006 spoken term detection evaluation. In *Proc. SIGIR*, volume 7, pages 51–57.
- Y. He, P. Baumann, H. Fang, B. Hutchinson, A. Jaech, M. Ostendorf, E. Fosler-Lussier, and J. Pierrehumbert. 2016. Using pronunciation-based morphological subword units to improve oov handling in keyword search. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 24(1):79–92.
- Georg Heigold, Vincent Vanhoucke, Alan Senior, Patrick Nguyen, Marc’Aurelio Ranzato, Matthieu Devin, and Jeffrey Dean. 2013. Multilingual acoustic models using distributed deep neural networks. In *ICASSP*.
- Sittichai Jiampojarn, Colin Cherry, and Grzegorz Kondrak. 2010. [Integrating joint n-gram features into a discriminative training framework](#). In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 697–700, Los Angeles, California. Association for Computational Linguistics.

- Martin Karafiát, Murali Karthick Baskar, Shinji Watanabe, Takaaki Hori, Matthew Wiesner, and Jan “Honza” Černocký. 2018. [Analysis of multilingual sequence-to-sequence speech recognition systems](#). *arXiv:1811.03451*.
- Christo Kirov, Ryan Cotterell, John Sylak-Glassman, Géraldine Walther, Ekaterina Vylomova, Patrick Xia, Manaal Faruqi, Sebastian J. Mielke, Arya McCarthy, Sandra Kübler, David Yarowsky, Jason Eisner, and Mans Hulden. 2018. [UniMorph 2.0: Universal morphology](#). In *Proceedings of the 11th Language Resources and Evaluation Conference*, Miyazaki, Japan. European Language Resource Association.
- Reinhard Kneser and Hermann Ney. 1995. Improved backing-off for m-gram language modeling. In *1995 International Conference on Acoustics, Speech, and Signal Processing*, volume 1, pages 181–184. IEEE.
- John Kominek and Alan W Black. 2006. [Learning pronunciation dictionaries: Language complexity and word selection strategies](#). In *Proceedings of the Human Language Technology Conference of the NAACL, Main Conference*, pages 232–239, New York City, USA. Association for Computational Linguistics.
- Viet Bac Le and Laurent Besacier. 2005. First steps in fast acoustic modeling for a new target language: application to Vietnamese. In *ICASSP*.
- Peter Makarov and Simon Clematide. 2018. [UZH at CoNLL–SIGMORPHON 2018 shared task on universal morphological inflection](#). pages 69–75.
- Sameer Maskey, Alan Black, and Laura Tomokiya. 2004. Bootstrapping phonetic lexicons for new languages. In *Eighth International Conference on Spoken Language Processing*.
- Arya D. McCarthy, Ekaterina Vylomova, Shijie Wu, Chaitanya Malaviya, Lawrence Wolf-Sonkin, Garrett Nicolai, Miikka Silfverberg, Sebastian J. Mielke, Jeffrey Heinz, Ryan Cotterell, and Mans Hulden. 2019. [The SIGMORPHON 2019 shared task: Morphological analysis in context and cross-lingual transfer for inflection](#). In *Proceedings of the 16th Workshop on Computational Research in Phonetics, Phonology, and Morphology*, pages 229–244, Florence, Italy. Association for Computational Linguistics.
- David R. Mortensen, Siddharth Dalmia, and Patrick Littell. 2018. [EpiTran: Precision G2P for many languages](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC-2018)*, Miyazaki, Japan. European Languages Resources Association (ELRA).
- Karthik Narasimhan, Damianos Karakos, Richard Schwartz, Stavros Tsakalidis, and Regina Barzilay. 2014. [Morphological segmentation for keyword spotting](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 880–885, Doha, Qatar. Association for Computational Linguistics.
- Garrett Nicolai and David Yarowsky. 2019. [Learning morphosyntactic analyzers from the Bible via iterative annotation projection across 26 languages](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1765–1774, Florence, Italy. Association for Computational Linguistics.
- Odette Scharenborg, Francesco Ciannella, Shruti Palaskar, Alan Black, Florian Metze, Lucas Ondel, and Mark Hasegawa-Johnson. 2017. [Building an ASR system for a low-research language through the adaptation of a high-resource language ASR system: preliminary results](#). In *International Conference on Natural Language, Signal and Speech Processing (ICNLSSP)*.
- Tanja Schultz, Alan W Black, Sameer Badaskar, Matthew Hornyak, and John Kominek. 2007. Spice: Web-based tools for rapid language adaptation in speech processing systems. In *Eighth Annual Conference of the International Speech Communication Association*.
- Tanja Schultz and Alex Waibel. 2001. Experiments on cross-language acoustic Modeling. *EUROSPEECH’01*, pages 2721–2724.
- A. Stolcke, F. Grezl, Mei-Yuh Hwang, Xin Lei, N. Morgan, and D. Vergyri. 2006. [Cross-domain and cross-language portability of acoustic features estimated by multilayer perceptrons](#). In *ICASSP*.
- Andreas Stolcke. 2002. Srilm-an extensible language modeling toolkit. In *Seventh international conference on spoken language processing*.
- Stephanie Strassel and Jennifer Tracey. 2016. [LORELEI language packs: Data, tools, and resources for technology development in low resource languages](#). In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC 2016)*, pages 3273–3280, Portorož, Slovenia. European Language Resources Association (ELRA).
- John Sylak-Glassman, Christo Kirov, David Yarowsky, and Roger Que. 2015. [A language-independent feature schema for inflectional morphology](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 674–680, Beijing, China. Association for Computational Linguistics.
- Jan Trmal, Guoguo Chen, Dan Povey, Sanjeev Khudanpur, Pegah Ghahremani, Xiaohui Zhang, Vimal Manohar, Chunxi Liu, Aren Jansen, Dietrich Klakow, et al. 2014. A keyword search system using

open source software. In *2014 IEEE Spoken Language Technology Workshop (SLT)*, pages 530–535. IEEE.

Jan Trmal, Matthew Wiesner, Vijayaditya Peddinti, Xiaohui Zhang, Pegah Ghahremani, Yiming Wang, Vimal Manohar, Hainan Xu, Daniel Povey, and Sanjeev Khudanpur. 2017. The Kaldi OpenKWS System: Improving Low Resource Keyword Search. In *INTERSPEECH*, pages 3597–3601.

C. van Heerden, D. Karakos, K. Narasimhan, M. Davel, and R. Schwartz. 2017. Constructing sub-word units for spoken term detection. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5780–5784.

Karel Vesely, Martin Karafiát, Frantisek Grezl, Marcel Janda, and Ekaterina Egorova. 2012. The language-independent bottleneck features. In *Spoken Language Technology Workshop (SLT), 2012 IEEE*.

Ngoc Thang Vu, Florian Metze, and Tanja Schultz. 2012. Multilingual bottle-neck features and its application for under-resourced languages. In *The third International Workshop on Spoken Language Technologies for Under-resourced languages*.