



TB or not TB? Acoustic cough analysis for tuberculosis classification

Geoffrey Frost¹, Grant Theron², Thomas Niesler¹

¹Department of Electrical and Electronic Engineering, Stellenbosch University, South Africa

²SAMRC Centre for Tuberculosis Research, DSI/NRF Centre of Excellence for Biomedical Tuberculosis Research, Stellenbosch University, South Africa

{gfrost, gtheron, trn}@sun.ac.za

Abstract

In this work, we explore recurrent neural network architectures for tuberculosis (TB) cough classification. In contrast to previous unsuccessful attempts to implement deep architectures in this domain, we show that a basic bidirectional long short-term memory network (BiLSTM) can achieve improved performance. In addition, we show that by performing greedy feature selection in conjunction with a newly-proposed attention-based architecture that learns patient invariant features, substantially better generalisation can be achieved compared to a baseline and other considered architectures. Furthermore, this attention mechanism allows an inspection of the temporal regions of the audio signal considered to be important for classification to be performed. Finally, we develop a neural style transfer technique to infer idealised inputs which can subsequently be analysed. We find distinct differences between the idealised power spectra of TB and non-TB coughs, which provide clues about the origin of the features in the audio signal.

Index Terms: cough, tuberculosis (TB), BiLSTM, attention, style-transfer

1. Introduction

In 2021, 10 million people were reported to have developed tuberculosis (TB), of whom 1.5 million died. As a result, TB was the second most lethal infectious disease globally, ranking above HIV/AIDS and just below COVID-19 [1]. The majority of TB cases occur in developing nations where access to public health care is limited by complex socio-economic factors, making it difficult to identify and control the spread of the disease and resulting in patients receiving improper care [2].

Whilst published research covering cough classification is currently limited, a few studies have shown promising results when distinguishing between: wet and dry coughs [3, 4], pneumonia [5, 6], and more recently COVID-19 [7, 8, 9]. Because TB is predominately a respiratory disease, it results in patients developing a chronic cough. It has been shown in previous work that it is possible to distinguish between the coughs of TB patients and healthy controls by utilising simple statistical classifiers [10]. More recently, these methods have been evaluated on a dataset that aims to reflect *real-world* conditions, whereby coughers all suffer from some lung ailment that is in some cases TB [11]. Whilst frequency bands important for classification were identified [10], a thorough investigation into the acoustic patterns being learnt has not yet been conducted. Moreover, work considering TB cough classification has relied on linear models utilising fixed dimensional inputs which are typically frame-wise averages of acoustic features. Thus, temporal information present in a cough has so far been disregarded.

In this work we show that recurrent deep learning architectures can be used successfully for TB cough classification, and

improve upon existing methods. In addition, by incorporating an attention mechanism and a new loss term, combined with frugal feature selection, we show that model generalisation can be improved. Using the same attention mechanism, we are able to visualise the temporal regions of the feature space that are learnt to be important for cough classification. By considering idealised TB negative and TB positive coughs produced by a technique normally used for neural style transfer, we discuss the distinct characteristics of a TB cough captured by the neural network.

2. Data

We report classification results on a dataset comprising 74 individual patients and 1564 coughs. Previous work in TB cough classification has relied on relatively small datasets gathered in a single recording environment from a small number of patients. This is problematic when training deep-architectures due to their tendency to overfit, for example, to confounding socio-environmental factors which are especially important to disregard in a clinical setting [12]. Relying on recordings from a single environment restricts data diversity and consequently the final model's ability to generalise. In an attempt to address this, we combine the datasets used previously in [10] and in [11], referred to as the Brooklyn and Wallacedene datasets respectively. This is in an effort to yield a more environmentally diverse dataset. Brooklyn was collected in a noise-isolated facility from patients known to have TB and healthy controls, whilst Wallacedene was collected in a noisy environment, from patients who all suffer from either TB or some other lung ailment (confirmed later by sputum analysis). This combined dataset is summarised in Table 1.

Table 1: *Dataset used for experimentation. TB and \overline{TB} indicate TB positive and negative respectively.*

	TB	\overline{TB}	Total
Patients	28	46	74
Total coughs	844	720	1564
Mean cough length (s)	0.60	0.64	0.62
Std dev cough length (s)	0.34	0.29	0.32

2.1. Cross-validation and testing

We divide the combined dataset into a training set (which is further subdivided for cross-validation) and a test set, containing 49 and 25 patients respectively. Importantly, both the Brooklyn and Wallacedene datasets are represented equally in all splits. Furthermore, splits are performed patient wise, ensuring that all coughs originating from the same patient are only present in one set, and we ensure a uniform distribution of TB positive and

negative (TB and $\overline{\text{TB}}$) patients across splits. The training set is further divided into 4 folds (each consisting of its own train and development set) for cross-validation using the same previously described procedure.

3. Models

We first train and evaluate several binary classifiers, including: logistic regression (baseline), a basic BiLSTM and a BiLSTM with attention. Next, we use the attention-based architecture to deduce important cough characteristics by generating idealised coughs for each class through a neural style-transfer technique. A sequential forward search (SFS) [13] is performed for both recurrent architectures to identify the most important frequencies for classification and to investigate its impact on model generalisation. This information is considered in conjunction with the temporal regions identified as important for classification by analysis of the attention weights.

3.1. Logistic regression

Previous work in TB cough classification has focused on simple linear models since it was observed that complex neural networks resulted in degraded performance. In both [10] and [11], logistic regression (LR) outperformed all other considered classifiers. As such, we use it as a baseline with which our architectures will be compared. LR is a simple approach that linearly models a probability $f(\mathbf{x})$ given a set of d predictors $\mathbf{x} \in \mathbb{R}^d$ using learnable parameters θ . This highlights an important limitation of LR: each predictor \mathbf{x} is a feature vector computed from a frame of audio. To obtain the probability that a cough is associated with TB, the average of the frame probabilities is computed. In doing so, any temporal information is lost.

3.2. BiLSTM

RNNs have successfully been used in several acoustic classification tasks. Acoustic feature vectors are processed sequentially, each updating the network’s internal hidden states which contain complex context-rich information and are available at the next time step. This allows the network to learn temporal relations important to the task at hand. In this work, we make use of a BiLSTM which extends the LSTM architecture [14] by processing the sequence in both forward and backward temporal directions.

A high-level diagram of the network is shown in Figure 1. The single BiLSTM layer has a 32-dimensional hidden state whereby the final outputs in both directions are concatenated to form \mathbf{q} . This embedding is then passed through a small feed-forward network with a 32-dimensional hidden layer and ReLU activations followed by an output layer. We include dropout before the first linear layer with a probability of 0.5. Lastly, to account for the unbalanced nature of our dataset, we use weighted cross-entropy as our loss function.

3.3. BiLSTM-Att

The development of the attention mechanism [15] has revolutionised deep learning research. With a focus on acoustic classification, attention-based architectures achieve near state-of-the-art results on tasks such as the Google speech commands dataset [16, 17]. In addition, the intuitive nature of the architecture allows for analysis of what the network is learning, reducing the black-box notion commonly associated with deep learning.

We develop an attention-based model by integrating an attention layer into the above BiLSTM architecture. Instead of

passing \mathbf{q} directly to the fully connected network as is done in the basic BiLSTM architecture, the attention mechanism uses \mathbf{q} as the query and outputs a weighted average (by the attention score) of all the BiLSTM outputs, thereby allowing the single output vector to capture information from the temporal regions most relevant for classification and suppress information from unimportant regions in time.

We design this architecture bearing in mind the fact that it will be used to aid in the understanding of the acoustic signature of a TB cough. Accordingly, a new loss term is introduced that encourages the embedding layer of the network (the output of the attention block) to generalise across patients of the same TB status. This is performed to inhibit our subsequent model analysis to be confounded by attributes learned irrelevant to TB cough classification, namely patient identity, an attribute present in cough [18]. This is accomplished by incorporating a GE2E loss term which was originally proposed to determine speaker embeddings by encouraging the network to keep embeddings close together when from the same target speaker, and further apart for different speakers [19]. We consider TB and $\overline{\text{TB}}$ coughs to represent two respective “speakers”. Hence the similarity between the embedding centroids of different patients with the same TB status is maximised, whilst minimising the similarity between embeddings of the TB and $\overline{\text{TB}}$ classes. The combined loss function used to train our network is given in Equation 1, where B is the batch size.

$$\mathcal{L} = -\frac{1}{B} \left(\sum_b \beta \cdot \mathbf{y}_b \cdot \log(\hat{\mathbf{y}}_b) + \alpha \sum_{j,i} L_{GE2E}(\mathbf{e}_{ij}) \right) \quad (1)$$

Here the first term is standard weighted cross-entropy where \mathbf{y}_b and $\hat{\mathbf{y}}_b$ are the vectors of ground truth and predicted probabilities respectively (where the dimension is the number of classes i.e. two) for a given cough in the batch and β is a vector of class weights (constant throughout training). The weight for the under-sampled class is set to 1, while for the over-sampled class it is the ratio of its occurrence in the training set to the total number of samples. In the second term, α is a regularisation parameter, L_{GE2E} is the function that computes the GE2E loss for a specific embedding in a given batch, and \mathbf{e}_{ij} is the embedding vector of the i^{th} cough from the j^{th} class.

4. Experimental procedure

With the exception of LR, which is trained using the standard `scikit-learn` recipe [20], models are trained for 15 epochs with a learning rate of 1×10^{-4} and batch size of 128. After training all 4 folds, the mean development AUC is computed for each epoch. The models from the epoch with the highest mean development AUC are selected, and at test time are ensemble. The decision threshold used for classification was $\bar{\gamma}$, the mean of the decision thresholds γ_n that result in the EER for each fold. Hence the decision threshold was chosen on the basis of the EER as in previous work [10, 11]. We note that it might be possible to improve performance if a strategy that chooses this threshold to optimise, for example, sensitivity and specificity, is adopted. However, we leave this investigation for future work.

4.1. Data Augmentation and feature extraction

We experimented with 3 data augmentation techniques: SpecAugment [21], random insertions and deletions, and speed-perturbation [22]. Initial experiments indicated that only speed perturbation was effective, and hence report only this form of




Figure 1: Structure of the basic BiLSTM (Network 1) and its attention variant (Network 2), with shared components indicated.

augmentation. We use warping factors of $\{0.9, 1.0, 1.1\}$ resulting in a 3-fold increase in the size of the dataset. We considered three types of acoustic feature: mel-spectrograms, linear filter-bank energies, and MFCCs (with appended velocity and acceleration, as well as cepstral mean and variance normalisation). The former was found to perform best for all architectures, with the ideal number of filter banks being 180, 128, and 80 for LR, the BiLSTM and BiLSTM-Att model respectively. Analysis was based on 2048-sample frames, with successive frames overlapping by 1536 samples (i.e. a frame-skip of 512). Variations in frame length and frame-skip were not considered in this work. All recordings were down-sampled to 44.1kHz before feature extraction.

4.2. Idealised coughs through neural style transfer


In an attempt to understand what the network is learning in order to distinguish between TB and TB coughs, we employ a technique similar to that used in neural style transfer [23] to synthesise an *idealised* cough for each class. This is accomplished by first defining a 2D parameter matrix $\hat{x} \in \mathbb{R}^{n \times d}$ (initialised to zeros) that will represent the input to the network, where d is the size of the acoustic feature vector seen by the network during training, and n is the number of frames. In this case we select $n = 80$. Next, we perform training as before, but instead of fixing the network input and output and optimizing the weights, we fix the weights and train \hat{x} . This allows the discovery of the input \hat{x} that best leads to the output class y for the trained weights, i.e. an idealised input cough feature representation for the output class in question.

5. Experimental results and discussion


We present classification performance for the various classifiers investigated and discuss our findings. In addition, we analyse the idealised cough mel-spectrograms produced when applying our adaptation of neural style transfer and observe the attention weights to infer the spectral and temporal regions that the classifier finds most useful for classification.

5.1. Classification

We present development set performance for each considered architecture in Table 2 and the associated test set performance in Table 3. A substantial increase in classification performance over the LR baseline is observed with regards to the basic BiLSTM model with all metrics either matched or improved upon, most notably the specificity. The test set AUC for both the basic BiLSTM and its attention variant are comparable, but there is a large discrepancy in the remaining metrics. This indicates that the decision threshold was not optimal and more robust alterna-



(a) Idealised mel-spectrograms for TB negative (left) and positive (center) coughs, with bins identified by SFS shown in yellow (right).



(b) Mean power of each idealised mel-spectrogram for the frequency bins deemed most important by SFS.

Figure 2: Idealised mel-spectrograms and mean spectral power.

tives to using the EER should be explored in future work.

When inspecting the effect of applying SFS on the two deep architectures, interesting observations can be made. We note a substantial reduction in the standard deviation of the EER-based thresholds determined for the BiLSTM-Att architecture (0.175 before and 0.070 after SFS) whilst the opposite is observed for the BiLSTM without attention. Despite achieving the highest test AUC, an increase in decision threshold standard deviation was observed (from 0.108 to 0.155). An increase in the variability of the decision threshold between folds indicates poorer generalisation. Conversely, with the BiLSTM-Att architecture, better generalisation across the folds is observed which is evident in the reduced standard deviation of the decision threshold. This is especially important with the implementation of a TB screening tool in mind, where model generalisation will be key.

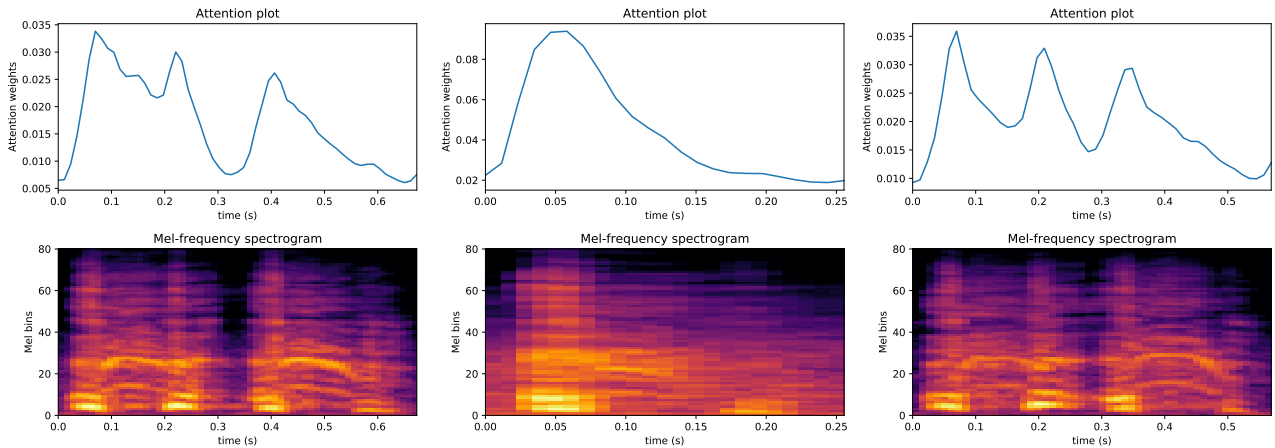


Figure 3: Various cough mel-spectrograms, all of which originate from the same patient, and their respective attention weights shown above them.

Table 2: Mean and standard deviation of the area under the ROC curve (AUC) and the EER decision thresholds ($\bar{\gamma}$) observed during 4-fold cross validation.

Model	$\bar{\gamma}$	AUC
LR (baseline) [10, 11]	0.272 ± 0.103	0.701 ± 0.127
BiLSTM	0.534 ± 0.108	0.777 ± 0.094
BiLSTM (SFS)	0.603 ± 0.155	0.919 ± 0.081
BiLSTM-Att	0.460 ± 0.175	0.873 ± 0.054
BiLSTM-Att (SFS)	0.568 ± 0.070	0.900 ± 0.092

Table 3: Test set performance for the models listed in Table 2, evaluated through various metrics: sensitivity, specificity, accuracy and area under the curve.

Model	Sens	Spec	Acc	AUC
LR (baseline) [10, 11]	0.889	0.625	0.720	0.769
BiLSTM	0.889	0.750	0.800	0.821
BiLSTM (SFS)	0.667	0.750	0.720	0.862
BiLSTM-Att	0.778	0.625	0.680	0.822
BiLSTM-Att (SFS)	0.778	0.813	0.800	0.850

5.2. Analysis and interpretation

Figure 2a depicts the idealised coughs synthesised by the BiLSTM-Att network using the neural style transfer method described in Section 4.2. Clear differences between idealised $\overline{\text{TB}}$ and TB cough mel-spectrograms are observed. This is further illustrated by comparing the mean power of these idealised coughs as a function of the frequencies determined to be most important by SFS, as shown in Figure 2b. For the idealised TB cough, we observe generally higher power at lower frequencies ($< 500\text{Hz}$) and the mid-band range of $1.8\text{kHz} - 3.3\text{kHz}$ whereas the $\overline{\text{TB}}$ cough has higher power between $1.2\text{kHz} - 1.8\text{kHz}$ and frequencies greater than 3.7kHz , which include frequencies far outside the typical range of human speech ($> 8\text{kHz}$). In Figure 3 we plot the attention weights as a function of time for three cough mel-spectrograms. We observe large importance being placed on regions where the signal has a high power and a large bandwidth, which coincide with the

initial bursts of energy for each coughing episode. Whilst only three examples are shown, these observations were made in general. This high energy portion of the coughing sound originates from the lung itself, in particular, the bronchi [24]. It therefore appears that, whilst TB can manifest in all regions of the respiratory tract, the model is relying on some change in the sound produced inside the lungs of TB and $\overline{\text{TB}}$ patients. Further research is necessary to deduce what the physiological causes of this difference in the audio signals could be.

6. Conclusion

In this work, we explored the use of recurrent networks for TB cough classification and use these trained networks to identify and interpret important cough characteristics in both frequency and time. A BiLSTM architecture is shown to improve on previous research, achieving a sensitivity and specificity of 0.89 and 0.75 respectively. This indicates that deeper architectures are viable for TB cough classification, and can improve upon previous state-of-the-art for TB screening. Furthermore, we show that by incorporating frugal feature selection our proposed attention-based architecture exhibits substantially better generalisation across folds than the other considered architectures. This is an important observation for future work, in which datasets will include many more recording domains and associated variability. Utilizing an attention architecture, the importance of certain temporal regions in the cough signal could be visualised. It was observed that the initial voiced regions of cough were the most important for classification. Moreover, by employing a neural style transfer technique, idealised TB negative and positive coughs were synthesised. Subsequent inspection revealed stark differences between energy content in specific frequency bands. In addition to providing new insights into the aspects of a tuberculosis cough that are important for classification, this provides evidence that the TB signal being learnt does indeed originate in the lungs. In future work, we look forward to evaluating our architectures on larger datasets currently being collected [25].

7. Acknowledgements

The authors gratefully acknowledge funding from the EDCTP2 programme supported by the European Union (grant RIA2020I-3305, CAGE-TB).

8. References

- [1] World Health Organization, "Global tuberculosis report 2021," 2021.
- [2] N. Foster, A. Vassall, S. Cleary, L. Cunnama, G. Churchyard, and E. Sinanovic, "The economic burden of TB diagnosis and treatment in South Africa," *Social science & medicine*, vol. 130, pp. 42–50, 2015.
- [3] Y. A. Amrulloh, D. A. Wati, F. Pratiwi, and R. Triasih, "A novel method for wet/dry cough classification in pediatric population," in *2016 IEEE Region 10 Symposium (TENSYP)*, 2016, pp. 125–129.
- [4] V. Swarnkar, U. R. Abeyratne, A. B. Chang, Y. A. Amrulloh, A. Setyati, and R. Triasih, "Automatic identification of wet and dry cough in pediatric patients with respiratory diseases," *Annals of biomedical engineering*, vol. 41, pp. 1016–1028, 2013.
- [5] Y. Amrulloh, U. Abeyratne, V. Swarnkar, and R. Triasih, "Cough sound analysis for pneumonia and asthma classification in pediatric population," in *2015 6th International Conference on Intelligent Systems, Modelling and Simulation (ISMS'15)*, 2015.
- [6] U. R. Abeyratne, V. Swarnkar, A. Setyati, and R. Triasih, "Cough sound analysis can rapidly diagnose childhood pneumonia," *Annals of biomedical engineering*, vol. 41, pp. 2448–2462, 2013.
- [7] M. Pahar, M. Klopper, R. Warren, and T. Niesler, "COVID-19 cough classification using machine learning and global smartphone recordings," *Computers in Biology and Medicine*, 2021, epub.
- [8] P. Mouawad, T. Dubnov, and S. Dubnov, "Robust detection of COVID-19 in cough sounds," *SN Computer Science*, vol. 2, pp. 1–13, 2021.
- [9] P. Bagad, A. Dalmia, J. Doshi, A. Nagrani, P. Bhamare, A. Mahale, S. Rane, N. Agarwal, and R. Panicker, "Cough against COVID: Evidence of COVID-19 signature in cough sounds," *arXiv preprint arXiv:2009.08790*, 2020.
- [10] G. Botha, G. Theron, R. Warren, M. Klopper, K. Dheda, P. Van Helden, and T. Niesler, "Detection of tuberculosis by automatic cough sound analysis," *Physiological Measurement*, vol. 39, 2018.
- [11] M. Pahar, M. Klopper, B. Reeve, G. Theron, R. Warren, and T. Niesler, "Automatic cough classification for tuberculosis screening in a real-world environment," *Physiological Measurement*, vol. 42, 2021.
- [12] C. J. Kelly, A. Karthikesalingam, M. Suleyman, G. Corrado, and D. King, "Key challenges for delivering clinical impact with artificial intelligence," *BMC medicine*, vol. 17, pp. 1–9, 2019.
- [13] F. J. Ferri, P. Pudil, M. Hatef, and J. Kittler, "Comparative study of techniques for large-scale feature selection," in *Machine Intelligence and Pattern Recognition*. Elsevier, 1994, vol. 16, pp. 403–413.
- [14] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, pp. 1735–1780, 1997.
- [15] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.
- [16] O. Rybakov, N. Kononenko, N. Subrahmanya, M. Visontai, and S. Laurenzo, "Streaming keyword spotting on mobile devices," *arXiv preprint arXiv:2005.06720*, 2020.
- [17] D. C. de Andrade, S. Leo, M. L. D. S. Viana, and C. Bernkopf, "A neural attention model for speech command recognition," in *INTERSPEECH*, 2020.
- [18] M. Zhang, Y. Chen, L. Li, and D. Wang, "Speaker recognition with cough, laugh and 'wei'," in *2017 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*, 2017.
- [19] L. Wan, Q. Wang, A. Papir, and I. L. Moreno, "Generalized end-to-end loss for speaker verification," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018.
- [20] F. Pedregosa *et al.*, "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
- [21] D. S. Park, W. Chan, Y. Zhang, C.-C. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le, "SpecAugment: A simple data augmentation method for automatic speech recognition," in *INTERSPEECH*, 2019.
- [22] T. Ko, V. Peddinti, D. Povey, and S. Khudanpur, "Audio augmentation for speech recognition," in *INTERSPEECH*, 2015.
- [23] L. A. Gatys, A. S. Ecker, and M. Bethge, "A neural algorithm of artistic style," *arXiv preprint arXiv:1508.06576*, 2015.
- [24] J. Korpáš, J. Sadloňová, and M. Vrabec, "Analysis of the cough sound: an overview," *Pulmonary pharmacology*, vol. 9, pp. 261–268, 1996.
- [25] "Cough Audio triaGE for TB (CAGE-TB)," <https://www.cagetb.org/>, accessed: 2022-06-19.