Bayesian Neural Networks for Acoustic Mosquito Detection

Ivan Kiskin (ikiskin@robots.ox.ac.uk), Adam D. Cobb, Steve Roberts

Abstract

- The **HumBug** project [1] aims to combat the spread of mosquito-borne diseases.
- We use a Bayesian convolutional neural network (**BCNN**) to successfully detect mosquitoes from their acoustics in challenging real-world conditions.
- Previous work has not quantified uncertainty, which we require for field deployment due to the variable nature of the target environment conditions.

HumBug pipeline

(a) Bednet with four smartphones positioned to trial the best location for recording mosquitoes.

(b) Bednet pockets to hold smartphones for recording.

Figure 1: Deployment in Tanzania (Oct 2020) to trial the effectiveness of acoustic mosquito detection with low-cost non-invasive measures.

- Smartphones are placed in bednet corners to allow autonomous data collection.
- Anonymised phone recordings synchronise to a central server and database.
- The recordings are screened with **Algorithm 1** to identify at-risk areas.

Algorithm 1 BCNN detection

1: Read 8 kHz mono wave file
2: Take sliding window log-mel transform (40×128 frames, each frame duration 64 ms)
3: Take N MC dropout samples from BCNN (2 conv/max-pool layers with 3×3 kernels)
4: Calculate mean probability $\bar{p}_C$, predictive entropy (**PE**), mutual information (**MI**)
5: Output the mean of $\bar{p}_C$, PE, MI per continuous section with $\bar{p}_{\text{mosquito}} > p_{\text{threshold}}$

Test performance

	Duration (h)	Class acc. (%)
Signal **A**	2.8	89.27 ± 0.07
Noise **A**	1.3	94.05 ± 0.11
Noise **B**	3.0	97.99 ± 0.05

Table 1: Out-of-sample performance estimated with $N = 10$ MC dropout samples (mean $\pm$ standard deviation). Signal **A**: lab and home mosquito recordings; Noise **A**: corresponding background. Noise **B**: bednet (**Figure 1a**) field background noise.

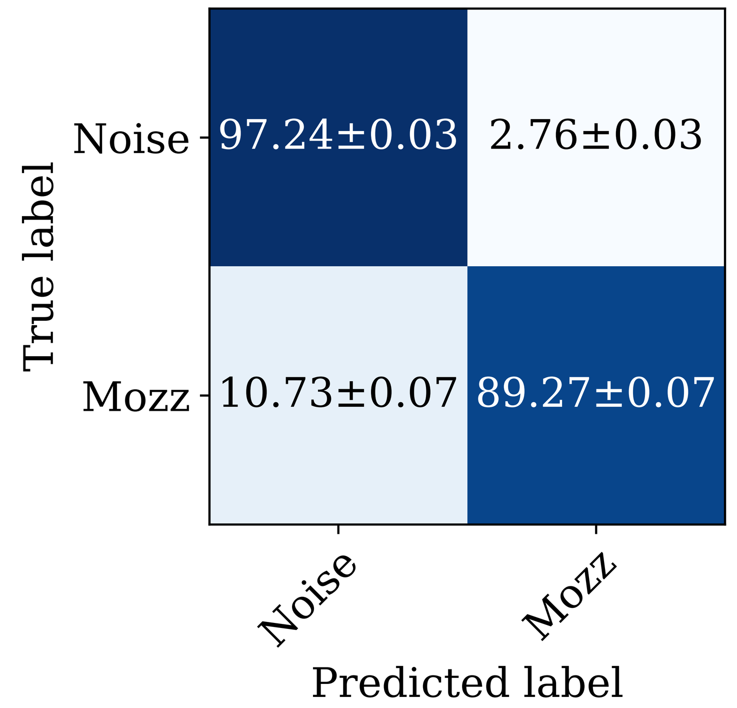

Figure 2: Out-of-sample performance on sources from **Table 1** combined.

Unlabelled data

- Import the audio and predicted labels to screen detections in Audacity [2].
- We can accept low entropy, high posterior predictive probability detections in the case where the audio background is noisy.
- Conversely, for good recording conditions we can accept higher entropy predictions to encourage higher recall of mosquito events.

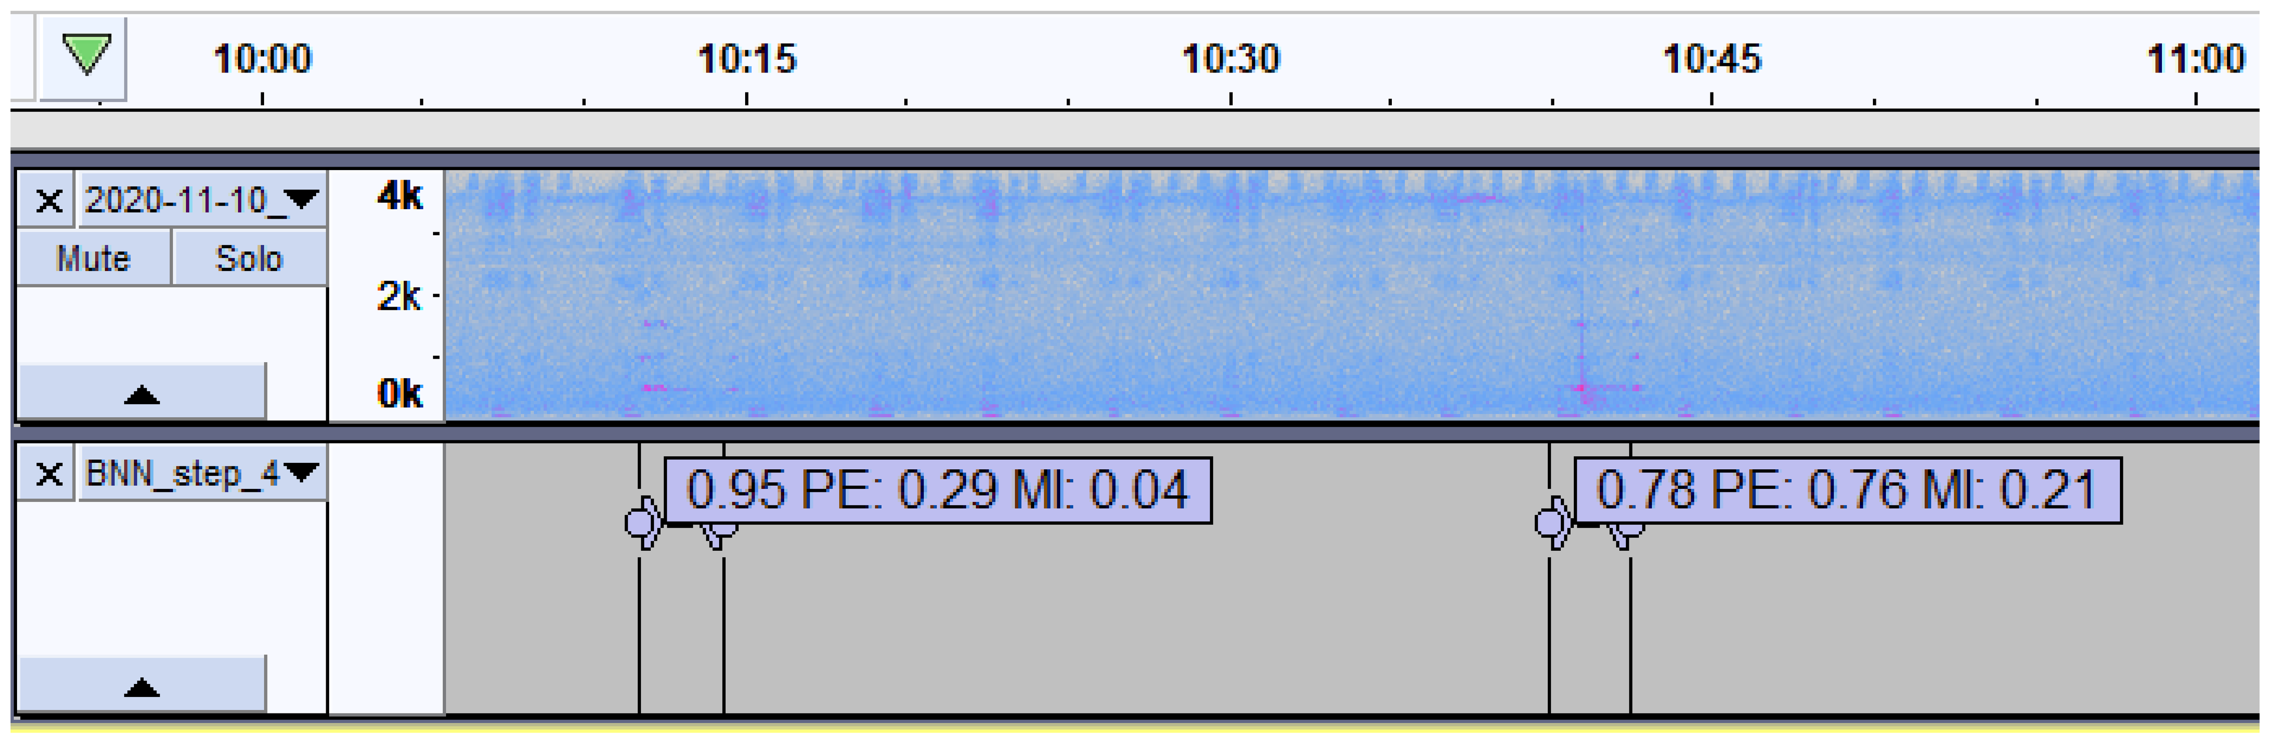

Figure 3: BCNN predictions on unlabelled field data (Nov 2020) in Audacity in the form: $\{\hat{p}_{\text{mozz}}, \text{PE}, \text{MI}\}$. Two windows with mosquito present were correctly identified in this section of audio, recorded in a similar arrangement to that shown in **Figure 1**.

References

[1] The HumBug project homepage: https://humbug.ox.ac.uk/

[2] Audacity. Open-source cross-platform audio editor. https://www.audacityteam.org/

Acknowledgements

- A thanks to all the team members at Ifakara Health Institute, the University of Oxford, and Kinshasa School of Public Health. For a full list please visit https://humbug.ox.ac.uk/the-team