Robust bird species recognition: making it work for dawn chorus audio archives

Dan Stowell*†1 and Mark Plumbley[1]

[1]Centre for Digital Music, Queen Mary University of London (C4DM, QMUL) – Mile End Road, London, E1 4NS, United Kingdom

Abstract

The recent (2013) bird species recognition challenges organised by the SABIOD project attracted some strong performances from automatic classifiers applied to short audio excerpts from passive acoustic monitoring stations. Can such strong results be achieved for dawn chorus field recordings in audio archives? The question is important because archives (such as the British Library Sound Archive) hold thousands such recordings, covering many decades and many countries, but they are mostly unlabelled. Automatic labelling holds the potential to unlock their value to ecological studies.

Audio in such archives is quite different from passive acoustic monitoring data: importantly, the recording conditions vary randomly (and are usually unknown), making the scenario a "cross-condition" rather than "single-condition" train/test task. Dawn chorus recordings are generally long, and the annotations often indicate which birds are in a 20-minute recording but not within which 5-second segments they are active. Further, the amount of annotation available is very small.

We report on experiments to evaluate a variety of classifier configurations for automatic multilabel species annotation in dawn chorus archive recordings. The audio data is an order of magnitude larger than the SABIOD challenges, but the ground-truth data is an order of magnitude smaller. We report some surprising findings, including clear variation in the benefits of some analysis choices (audio features, pooling techniques noise-robustness techniques) as we move to handle the specific multi-condition case relevant for audio archives.

Keywords: bird, bird song, automatic classification, multilabel, noise, archives

---

*Speaker
†Corresponding author: dan.stowell@qmul.ac.uk