

Kaleidoscope Pro 4.1 Auto ID for Bats

Ian Agranat, Wildlife Acoustics, Inc. Concord, Massachusetts, U.S.A.

Paper in progress...

Our previous work on classifiers for bats prior to 4.1 is described here:

<http://www.wildlifeacoustics.com/downloads/batid.pdf>

With version 4.1 of Kaleidoscope Pro, several improvements to the algorithms have been incorporated. I am working on a new paper to describe these techniques in detail. Until then, this document serves as a brief summary.

As before, the series of zero crossing dots (after applying a division ratio) can be represented as an observation sequence comprising two-dimensional vectors of the frequency and delta frequency of each of the frequency-divided dots.

Prior to 4.1, the sequences were divided into individual echolocation pulses. With 4.1, we consider both pulse-by-pulse level analysis as before, and the sequence level analysis comprising all of the dots across multiple pulses in a sequence. The sequence level analysis allows the models to represent features such as average time between calls and the variability of Fmin across pulses that were not previously captured when considering only pulse-by-pulse analysis. This improves classification accuracy for some species.

Prior to 4.1, an HMM was created for each species. With 4.1, there is one global HMM with states seeded with KMeans++ to maximize separation between states and trained on the entire training data set.

A GMM is also used to represent eleven common call parameters such as TBC, Fc, Sc, etc. The Fisher score of a given bat pulse or sequence of pulses relative to the GMM and HMM model parameters forms a feature vector used to compare pulses or sequence for similarity using the dot product. This measure of similarity between pulses or sequences is then used to cluster the training data.

For clustering, we use an algorithm that forms spherical micro-clusters of some maximum radius of similar pulses or sequences based on the dot product of Fisher scores.

When clusters form containing only pulses or sequences from one species, the classifier will assign any pulses or sequences falling into such clusters as belonging to the species class. However, given the statistical overlap of many species, there will also be many clusters containing pulses or sequences from multiple species. In this case, we additionally create pair-wise classifiers using the Winnow algorithm for all the training data in overlapping classes. This is essentially what we did in prior versions, except now we do not form pair-wise classifiers for non-overlapping species for efficiency. If the classifier finds a new pulse or sequence falling into a cluster representing multiple species, the pair-wise classifiers for

those cluster-specific overlapping species are then used to determine which species is more likely using a hinge-loss function.

For the pulse level, given the significant statistical overlap across species, an additional fuzzy logic layer is created by testing the distribution of pulse-level classifications or misclassifications for each species decision in each micro-cluster assignment. Thus, each pulse in a sequence can be classified and the fuzzy logic layer applied to aggregate the pulses into the most likely species for the sequence based on pulse-by-pulse analysis. We found this aggregated pulse-level classification tends to be slightly more accurate than the sequence-level classification. If the pulse-level classification matches the sequence-level classification, then we have a confident match. Otherwise, the “NoID” result is returned. Agreement in runner-up choices (e.g. 2nd and 3rd most likely pulse and sequence level classifications) are output as possible alternative identifications.

Overall, we find this new approach to be more accurate (though slightly less sensitive e.g. more conservative) than our previous classifiers on more identifiable call sequences. More importantly, the new classifiers are more likely to return “NoID” when the sequence is ambiguous e.g. contains multiple bats or deviates more significantly from the training data.

The pulse-level clustering is also used to form a visual call reference library from the larger clusters for each species. These represent the more commonly occurring pulse shapes in the training data for each species. The classifier sensitivity settings e.g. more sensitive, balanced, or more accurate, require zero, one, or two minimum pulses to be represented by these more common pulse types.

Performance of these classifiers based on a random 50/50 split of data into training and testing can be found at

<http://www.wildlifeacoustics.com/products/kaleidoscope-software-ultrasonic/classifiers>

If you would like to be notified when more details are published, please sign up on our mailing list to receive Kaleidoscope technical support bulletins.

<http://www.wildlifeacoustics.com/company/contact>

And as always, if you have any questions, please direct them to our technical support team at

support2017@wildlifeacoustics.com