

Video Script

Kaleidoscope Pro 5 Advanced Classifiers

Welcome to Kaleidoscope Pro from Wildlife Acoustics.

This video will show how to create and tune an Advanced Classifier to find specific vocalizations in new input files.

The Advanced Classifier will be much more discriminating than a Simple Classifier. An Advanced Classifier is used to reduce false positive rates by providing additional information to the analysis model to optimize the separation of statistically similar vocalizations.

If you want to find a specific individual species or vocalization in a batch of files, an Advanced Classifier may be the right tool for the job. If you want to do a general survey of many species, a basic cluster analysis or Simple Classifier will probably be a better solution for that job.

I've got two sets of audio recordings. I'm going to use the first set of data to build and train the Advanced Classifier. Then I'll use the Advanced Classifier to sort the second set of data. I'll choose the first data set for the Input directory. I'll assign an Output directory. I'll go to the Cluster Analysis tab and choose the first option to scan the input files for basic cluster analysis. I'm working with all default settings in this window. The default FFT window setting is optimized for most common birds. If I was looking for lower frequency sounds such as owls, I might set the FFT window to a larger size.

I'll run the cluster analysis batch process. Kaleidoscope Pro examines the input files for detected signals. The detected signals or vocalizations are then sorted based on statistical similarity. Clusters are created and sorted based on the similarities of the detected signals. The first listed cluster represents the most common set of similar detected signals.

Once the first cluster analysis batch process is completed, the Viewer and Results window open.

The goal is to train the Advanced Classifier to be able to accurately find the desired species or vocalization, and reduce false positive or false negative identifications. This is done by creating manual identifications for vocalizations within clusters, and then using those manual IDs to create and train the Advanced Classifier.

The next step is to copy the cluster names to the Manual ID column in the Results window. This provides general clusters destinations for the detected input signals from the new set of data.

I am now going to manually label the best examples of the vocalizations I do want to isolate.

I'll check the first vocalization and audition the audio. I happen to know this vocalization is a spring

Video Script

peeper. I'm looking for birds and the spring peeper is a frog. The spring peeper sings constantly through the night and if I check the next several clusters I find that they are all spring peepers. The vocalizations have been sorted into separate clusters because although they are the same species, there are enough differences in the recordings to group them separately. These clusters do not appear to contain any vocalizations that I want to isolate, so I'll leave their cluster names alone. These clusters will be ready to catch spring peepers in the new batch of input data.

Ok, this vocalization is not a peeper and if I audition it, it sounds like a titmouse. I do want the Advanced Classifier to be able to find titmice, so I'll create a button label with that name. I press the button and that applies a manual verification to the vocalization and I see the verification listed in the Results window. The Auto Next File button is engaged so when I verify the vocalization the next vocalization on the list is cued up. I see the pattern is the same so I'll add the manual titmouse verification to the next vocalization. I'll continue to verify titmice until I start to see vocalizations that are either poor examples or not a titmouse at all. Vocalizations are listed in the order of the most similar examples of the statistical pattern. So as I get further into the cluster the examples start to become less obvious.

What I'm doing by adding these manual verifications is training Kaleidoscope Pro to recognize the verified signals. If I see a poor example of the vocalization, or perhaps the detected signal contains a titmouse but also contains other birds or sounds, I will apply a blank label to that detected signal. Now this detected signal has no manual ID at all. That means it will not be used to train the Advanced Classifier. I only want training data that are good examples of what I do want to find, and what I don't want to find.

I'll advance through the next clusters. This cluster looks like it could be a robin. I'm not looking for robins so I'll leave the original cluster name alone. When new data is analyzed with the Advanced Classifier, any robin vocalizations will get sorted into this cluster. Here is another cluster that looks like titmice. Just like with the peepers I may find multiple clusters that contain titmice. I want to train the Advanced Classifier to recognize variations of the titmouse so I'll continue to label the best examples of that species.

Here are more peepers. These frogs are everywhere!

Now I get to a vocalization that is not a peeper or a titmouse. I'll audition and I think it's a black capped chickadee. I'll create another custom button label and start to verify the chickadee vocalizations. At this point I am building an Advanced Classifier that is trained to find both chickadees and titmice.

Ok, now right in the middle of this chickadee cluster I find something that is not a chickadee. It's another peeper! The peeper has enough in common with the chickadee that it got included in this cluster. I don't want the Advanced Classifier to think this vocalization is a chickadee, so I'll leave the original cluster name. Now if Kaleidoscope Pro is finding chickadees and it comes across a peeper that is similar to the chickadee, it will be able to discriminate between the two and the peeper will get separated and sorted into its own cluster.

Video Script

At this point I've created new cluster names for two species. If I wanted to identify other species besides the chickadees and titmice, I may be better off building separate Advanced Classifiers for those additional species.

I can check through the remaining clusters but I think I've got what I need for now.

I'll go to the Results window File menu and choose Save. The Results window represents an underlying cluster.csv file that was created in the Outputs directory during the initial cluster analysis batch process. Saving updates the .csv file with my manual verifications.

Now it's time to build the Advanced Classifier. Under the Batch tab I keep the Input Directory set to the original batch of recordings. I'll select a new Output Directory. I go back to the Cluster Analysis tab and choose the second option. The second clustering option will re-scan the original recordings and make new clusters based on the manual verifications in the underlying edited .csv file. I'll run the second cluster analysis batch process.

Only the vocalizations that I've labelled will be used as the statistical models for the new clusters. Once the batch process is complete the Viewer and Results window open. I can do a quick check and what I'm seeing here are named clusters based on my manual IDs and the original training data as detected signals.

When Kaleidoscope Pro runs a Cluster Analysis batch process it creates the cluster.csv file and it also creates a cluster.kcs file. The .kcs file is similar to the .csv file in that it contains information about the clusters. The purpose of the .kcs file is that it can be used as a classifier for new data. This second step of cluster analysis has created a cluster.kcs file that has been trained via the manual IDs. The new cluster .csv and cluster.kcs files have been created in the second Outputs directory. The new cluster.kcs file is the Advanced Classifier.

If I want to further tune the Advanced Classifier I can repeat the previous step of manually checking and labeling the detected signals in the clusters. I can then repeat the step of using the newly edited .csv file to compare against the original training data. That will produce a new .kcs file that is the tuned Advanced Classifier. When I'm happy with the results I get from the original training data I'm ready to use the Advanced Classifier to sort new input data.

Now for my Input directory I'll choose the second data set. I'll select a new Output directory. Under the Cluster Analysis tab I'll choose the third clustering option. This third option uses an existing cluster.kcs file to analyze and sort new input data. What's going to happen here is the new set of data will be analyzed using the trained cluster.kcs file to sort and identify the new vocalizations.

Once this last batch process is complete the Viewer and Results window open again. The vocalizations in the Results window have been detected from the new batch of files. They have been analyzed, named, and sorted based on the information in the Advanced Classifier.

Here are chickadees. Here are titmice. I don't see any false positive identifications because the trained

Video Script

Advanced Classifier can now discriminate between similar but different vocalizations. And lastly, here are the general clusters with their original names, which I can just disregard.

Kaleidoscope Pro has done a first-rate job of finding the specific vocalizations I'm looking for in the new input recordings.

Thank you for watching.