



The 7th International DCLDE Workshop

[Detection, Classification, Localization, and Density Estimation]

July 13 - 16, 2015 • La Jolla, CA

Scoring Tool

The scoring tool is designed to compare detections with the ground truth files provided for the workshop. It accepts files in the workshop's CSV format:

For the high-frequency task, the result file should contain comma separated value (CSV) entries with each line as follows (see the DCLDE dataset [description](#) for further details on species abbreviations, etc.):

```
project, site, species-abbreviation, start-time, end-time
```

Time stamps are provided in ISO 8601 format: YYYY-MM-DDTHH:MM:SS with an optional decimal and fractional seconds following the seconds field. Example for Risso's dolphin detection at CINMS site B:

```
CINMS,C,Gg,2012-04-09T06:44:03.0,2012-04-09T07:16:40.0
```

Low-frequency-task results for blue and fin whales are similar:

```
D CPP, C, Bp, 2013-02-04T15:13:15.8, 2013-02-04T15:13:16.3
```

New in 2018: In the low-frequency-task ground truth files, an additional call quality field is included. This field can have one of two values, "good" or "poor", and is used for scoring purposes. This field will be ignored if provided in result files. By default, "poor" quality detections are not counted as false positives or false negatives.

Spaces between fields may be included or omitted.

Scoring is implemented with a set of Matlab scripts whose entry function is *dclde2018*. It accepts two arguments, groundtruth and detections. To compare multiple files, cell arrays of filenames may be passed as arguments.

Examples:

```
r = dclde2018('analyst.csv', 'mywonderfuldetector.csv');

r = dclde2015({'analyst-CINMS18C_1.csv', 'analyst-CINMS19C_2.csv'}, ...
             {'mydet-CINMS-18C_1.csv', 'mydet-CINMS-19C_2.csv'});
```

In both cases, a result structure will be returned. It contains the following members:

- precision – # correct detections / # detections
- recall - # detected ground-truthed calls / # ground-truthed calls
- falsePosl – An indicator array for each detection indicating whether or not it is a false positive (1 → false positive).

There are a number of coverage statistics. For ground-truthed signals, this means the amount of time (or percentage) of the signal that was covered by a detection. For detections, this is the amount of time (or percentage) of the detection that corresponded to a ground-truth call(s). These are captured in the following fields:

- truthCoverage – Time (Matlab serial dates) each ground truth signal is covered.
- truthCoveragePct – Percentage of each ground truth signal that corresponds to one or more detections.

Detection signals have similar fields: detectionCoverage and detectionCoveragePct.

Low ground-truth coverage indicates that portions of the signal are not being detected. Low detection coverage indicates that the detections are longer than the ground truth detections and that the detector might be overzealous (remembering that analysts make mistakes as well).

The final fields are:

- fragmentation – For each ground truth signal, the number of detections associated with it. Values greater than 1 indicate that the signal was broken into two or more possibly overlapping detections.
- Efragmentation – Empirical expected fragmentation based on the mean of all fragmentation values that do not correspond to missed ground-truth signals.
- associations – Indicator functions (stored as a Matlab sparse matrix) containing values of 1 any time a ground truth signal (rows) has been associated with a detection (columns).

Additional descriptions of coverage and fragmentation can be found in [Roch et al. \(2011\)](#).

When multiple sites are in the set of processed CSV files, use function `dclde2015multisite` which will partition detections by project and site and compare ground truth and detections from the same places. The same metrics are returned, but are returned as an array with one element per site and a final element that represents a combined score.

An example invocation:

```
% Find the filenames of all .csv files from directories mydetectiondir and dclde2015\hf.
% Both must be relative to the current directory and the dclde2015 directory
% must be in our path or the current directory.
gtdetections = findfiles('dclde2015\hf', '*.csv$', 'regexp');
detections = findfiles('mydetectiondir', '*.csv$', 'regexp')
% Score the files
r = dclde2015multisite(gtdetections, detections);
% Look at detections from first site CINMS/B
>> r(1)
ans =
    precision: 1
    recall: 0.3289
    falsePosl: [26x1 double]
    truthCoverage: [76x1 double]
    truthCoveragePct: [76x1 double]
    truthCoverageOverallPct: 0.0871
    detectionCoverage: [26x1 double]
    detectionCoveragePct: [26x1 double]
    detectionCoverageOverallPct: 0.8577
    fragmentation: [76x1 double]
    Efragmentation: 1.2800
    associations: [76x26 double]
    truth: [1x1 struct]
    detections: [1x1 struct]
    project: 'CINMS'
    site: 'B'
```

Both scoring functions support additional arguments that may be helpful to some users. These are specified as optional keyword value pairs that may follow the mandatory filename arguments:

- 'AllQualities', true|false (**New for 2018**) Default false. If false, detection or omission of “poor” quality ground truth detections is ignored. If true, results are scored against all ground truth detections, regardless of ground truth detection quality.
- 'IgnoreLabels', true|false – If true, labels are ignored when comparing detections to ground truth. This can be useful for testing detectors that do not detect specific calls/species.
- 'RemoveZeroDuration', true|false – If true, any detection with a duration of zero will be removed before processing starts. When false, warnings are issued for zero length detections. If zero length detections are permitted, it is possible to have ground truth detections that were matched, but have no coverage.

- 'SpecificLabels', mapObj – Only consider detections from specified species (and possibly specific calls from that species) when scoring. All other detections are not scored. See function description of read_annotations for details.

These functions have been tested with Matlab 2014b but will most likely work with Matlab 2008b or later. Note that you must set your path to access the directory containing the files.

New for 2018: Further information on detection quality parameter

Defining a detectability threshold for ground truth detections is tricky. We do not want to penalize detection of very faint but true calls, however many faint calls are only recognizable as calls to a human analyst based on context (e.g. call is close in time to other calls).

In this dataset, we have chosen to distinguish between “good” quality calls, which we would expect a detector to find, and “poor” quality calls, which we would not expect to be detected. By default, this year’s scoring tool does not count detection of “poor” quality calls toward performance metrics, but they are included in the ground truth files so that if they are detected, there will be no false positive penalty.

If the 'AllQualities' argument is used, the tool will score relative to all calls in the ground truth set, regardless of quality. This is likely to be a much tougher score. High performance with the 'AllQualities' flag enabled would suggest that the classifier is able to leverage call context as well as the acoustic characteristics of the call itself.