

F-POD Software Guide

Mar 2025

Background: PODs monitor the presence and activity of dolphins, porpoises and other toothed whales by the detection within the FPOD app of the 'trains' of echo-location clicks that they make. This gives them much lower false positive rates than can be achieved by identifying the characteristics of individual clicks, whether in wav files or other formats. There is more on the F-POD, the older C-POD, and relevant aspects of acoustics, in several PowerPoint presentations that can be accessed on the Chelonia website (www.chelonia.co.uk).

Feedback on this guide, and on any aspect of the instrument, is hugely valuable to us and to all users, so please don't hesitate to send comments, criticisms and questions to nick.trogenza@chelonia.co.uk

FPOD app: Installation and requirements:

The FPOD app 'FPOD.exe' runs under all versions of Windows. It requires a screen with a vertical resolution of at least 1080 pixels to display comfortably. On MACs FPOD.exe will run under "Parallels Desktop".

FPOD.exe is in *portable format* so that it can be run without being installed. It is checked for viruses etc. when uploaded to the Chelonia website. Download FPOD.exe from www.chelonia.co.uk. Place it in a folder wherever you like on your PC, or you can run it from a USB stick, an SD card or any other external drive. Deleting the file FPOD.exe is all that is required to remove it from your computer.

4K screen: strongly recommended, especially for validation work. Below 1024x768 the display is unsatisfactory.



Help: the red text in the pale green bar at the top of the screen provides a brief note on what is under the mouse pointer. If it is over the 'Delete file set' button on the Filters +files page this is shown:

This deletes the file selected for action and its associated set of files (FP1,FP3 and map files), even if they are not open. This is final. They do not go to the Recycle Bin!



Menu: this is shown or hidden by moving the mouse pointer up to the top of the FPOD app window and down again. Or clicking the large green label as above, or the small 'menu' button top left on the main screen.



The FPOD app and upgrades are free. To ensure you are notified of upgrades join the mailing list via https://www.chelonia.co.uk/contact_us.htm



Function keys are useful shortcuts. On some keyboards you can just press that key. On others you have to press the Fn key at the same time.

Favourite features

- [A trend estimation tool](#) is now included. This uses paired year ratio assessment (PYRA) and is useful for data sets of 3 years and longer.
- If you see something interesting in your data do CTRL + S and *save a 'comment' or settings file* with a really informative name like 'Dolphin social downsweeps' and you can then easily go back to view it on another day.
- You can play an audio representation of the rate of clicking shown on the screen. See the Display page info.
- Error estimation
- *Minute filters* can show you subsets of minutes, e.g. those in which you have marked a species sighting.
- *Validation sampling points* –statistically appropriate positions for sampling data can be set and visited very easily.
- *Third party utility* – this enables you to export a rich set of features of detections and re-import the results of your own, hopefully better classifier.

Quick start

Starting, deploying and stopping the F-POD are described in the *F-POD User Guide*.

Getting a file from the SD card:

- On retrieval of an SD card, insert it into the SD card reader or directly into your PC.
- Run FPOD.exe:
 - > Show the menu
 - > select the Read SD tab

Read SD | Files | Filters +files | Display | Trains | Settings | Navigation | About | Export

Create .FP1 or .CP1 raw data files on your PC from data on an SD card:


1: Site name - for file name no gaps TaterDhu

2: Metadata Deployment
(saved in the file) other information

Water depth m POD depth GMT Site code
LAT ... LONG ...

3: F-POD: read SD card FP1 file size limit 3GB join series Show date setting + C-POD reading

- > enter the location name in section 1. This name is *very important* and must be used consistently for that site.
- > add any deployment notes in section 2.
- > click 'Read F-POD SD card' in section 3.
- > select *any* file ending in .CHE on the SD card.

> select a directory for your files.  *Keep all the files for a project in **one** directory.* This really will save time and errors later.

All the .CHE files will now be sorted, read and converted to a number of FP1 files - their maximum size is set on the ReadSD page, so you can have huge single files if you wish:

3: F-POD: read SD card FP1 file size limit 3GB

- Do not put the POD number or date into the file name. These are added automatically in a tried and tested format that helps you handle and sort large numbers of files.
- Think carefully about how to name sites. For long term trend estimation a site that is more than 200m from another needs a different name.
- Do not put your POD files in many separate directories unless your project is very small and you have time to waste.

After you have read the SD card and had a look to see that the FP1 file matches what you expect you can erase the SD card and use it again. We recommend installing this SD card formatter: <https://www.sdcard.org/downloads/formatter/> from the SD Association. It is more reliable than others and only ever formats SD cards!

Getting an FP3 file from the FP1 file: 'Train detection'

Process the FP1 files through the KERNO-F classifier by going to the Trains page of the menu and clicking 'Detect click trains in FP1 files'

'Skip if FP3 exists' + 'search directories' lets you just pick the folder/directory where you have put the FP1.

This process will create an FP3 file for each FP1 file.

CLICK TRAIN DETECTION See instructional PowerPoints on FP3/CP3 files are the key output of PODs. FP3 files are created here from

skip if FP3 file exists search directories
 No dolphins No NBHF No sonars

Detect click trains in FP1 files - runs KERNO-F classifier

River dolphin detection is generally not as good as the marine dolphins on which the classifier was trained and they generally require train detection settings like these: (unless sonars are common)

CLICK TRAIN DETECTION

See instructional PowerPoints on FP3/CP3 files are the key output of PODs. FP3 files are created here from

skip if FP3 file exists search directories
 No dolphins No NBHF No sonars

Detect click trains in FP1 files - runs KERNO-F classifier

Noisy sites, which includes sandy rivers, may benefit from the exclusion of weaker clicks. This is available in the 'advanced settings'. Here clicks with an amplitude of less than 50 are excluded, and the detection rate for dolphins is generally improved.

Advanced KERNO-F settings		Limit clicks/m to	0
Duration N of cycles >	0	Force mins OFF	0
Amplitude >	50	Sensitivity level	0
No cetaceans at <	25 kHz		
Exclude frequencies	0 to 18 kHz	Click rate good >	0
Include only frequencies	22 to 221 kHz	NBHF boost	1
WUTS amplitude max	180	WUTS cluster max	3

<<< if you do have a cetacean clicking at <25kHz you can change that setting. Mostly best to keep the defaults!

A classification warning can be s
 FP3 file if the target is not set to

Save current train classifier settings to a .fpt file

Load saved train classifier settings from a .fpt file

If you are using these special settings, then save them as a 'train classifier settings file' using the button for that, and load them every time you process now files to maintain consistency.

What settings should I use when viewing and exporting data?

The standard settings for marine species are

ignore all filters (F4)

Train filters:
 marked trains included [more](#)
 high sp confidence only

Quality:

Hi Lo allQ
 Mod ?/echo

Species:

NBHF Other cet
 Sonar unclassified

Dolphins:

ignore all filters (F4)

Train filters:
 marked trains included [more](#)
 high sp confidence only

Quality:

Hi Lo allQ
 Mod ?/echo

Species:

NBHF Other cet
 Sonar unclassified

Porpoises:

(Note: there are some dolphins that make NBHF – narrow-band high-frequency) clicks.

Other filters can be applied to obtain a picture of the use of fast click trains etc. These filters are on Filters +files page:

Read SD	Files	Filters +files	Display	Trains	Settings	Navigation	About	Export	Results	Analysis																																																															
<div style="display: flex; justify-content: space-between;"> <div style="width: 45%;"> <p>Train filters</p> <table border="1"> <thead> <tr> <th></th> <th>Min</th> <th>Max</th> <th></th> <th>Min</th> <th>Max</th> </tr> </thead> <tbody> <tr> <td>median kHz</td> <td>0</td> <td>255</td> <td>av Bandwidth*</td> <td>0</td> <td>31</td> </tr> <tr> <td>N clicks in train</td> <td>5</td> <td>9999</td> <td>av Peak Position*</td> <td>0</td> <td>16</td> </tr> <tr> <td>Mean amplitude</td> <td>1</td> <td>9999</td> <td>exclude kHz:</td> <td>255 to</td> <td>255</td> </tr> <tr> <td>Mean N of cycles*</td> <td>0</td> <td>255</td> <td>WUTS risk*</td> <td>0</td> <td>7</td> </tr> <tr> <td>Clicks/s from</td> <td>1 to</td> <td>2500</td> <td></td> <td>Min</td> <td>Max</td> </tr> <tr> <td>exclude from</td> <td>2500 to</td> <td>2500</td> <td>Click rate confidence:</td> <td>0</td> <td>15</td> </tr> </tbody> </table> <p>Social calls included</p> <p>* F-POD only <input type="button" value="Clear all train filters"/></p> </div> <div style="width: 45%;"> <p>Click filters</p> <table border="1"> <thead> <tr> <th></th> <th>Min</th> <th>Max</th> </tr> </thead> <tbody> <tr> <td>kHz:</td> <td>0</td> <td>255</td> </tr> <tr> <td>exclude kHz</td> <td>255 to</td> <td>255</td> </tr> <tr> <td>click cycles:</td> <td>0</td> <td>255</td> </tr> <tr> <td>amplitude</td> <td>1</td> <td>9999</td> </tr> <tr> <td>NHBF index*</td> <td>0</td> <td>32</td> </tr> <tr> <td>IPI re expon</td> <td>27</td> <td>36</td> </tr> </tbody> </table> <p><input type="button" value="Filter C/FP1 for FISH tags"/> <input type="button" value="Clear all click filters"/></p> </div> <div style="width: 10%;"> <p>Minute filters F-POD exclude minutes with:</p> <p><input type="checkbox"/> boat sonars <input type="checkbox"/> WUTS risk <input type="checkbox"/> continuous noise <input type="checkbox"/> OpThreshold 0 to 8 <input type="checkbox"/> No Hi or Mod trains <input type="checkbox"/> angle > 180 <input type="checkbox"/> marked minutes <input type="checkbox"/> annotation n <> 0</p> <p><input type="button" value="save .fpf"/> <input type="button" value="load .fpf"/></p> </div> </div>												Min	Max		Min	Max	median kHz	0	255	av Bandwidth*	0	31	N clicks in train	5	9999	av Peak Position*	0	16	Mean amplitude	1	9999	exclude kHz:	255 to	255	Mean N of cycles*	0	255	WUTS risk*	0	7	Clicks/s from	1 to	2500		Min	Max	exclude from	2500 to	2500	Click rate confidence:	0	15		Min	Max	kHz:	0	255	exclude kHz	255 to	255	click cycles:	0	255	amplitude	1	9999	NHBF index*	0	32	IPI re expon	27	36
	Min	Max		Min	Max																																																																				
median kHz	0	255	av Bandwidth*	0	31																																																																				
N clicks in train	5	9999	av Peak Position*	0	16																																																																				
Mean amplitude	1	9999	exclude kHz:	255 to	255																																																																				
Mean N of cycles*	0	255	WUTS risk*	0	7																																																																				
Clicks/s from	1 to	2500		Min	Max																																																																				
exclude from	2500 to	2500	Click rate confidence:	0	15																																																																				
	Min	Max																																																																							
kHz:	0	255																																																																							
exclude kHz	255 to	255																																																																							
click cycles:	0	255																																																																							
amplitude	1	9999																																																																							
NHBF index*	0	32																																																																							
IPI re expon	27	36																																																																							

Other resources:

Join the POD users mailing list [here](#)

[This published paper](#) contains a detailed description of the process of validating detections and details of how the F-POD works, so it is a useful reference for any paper using POD data.

[Topic notes and glossary](#) ... at the end of this document

This [Youtube channel](#) has videos on numerous F-POD topics

PowerPoint presentations on these topics are available on the Chelonia website including these links

[PODs and acoustics](#) links to a set of resource videos / documents

[Validation of F-POD data](#)

[F-POD version history](#)

The FPOD app

What it does:

FPOD.exe **converts** the data stored on an SD card to two files:

- The **FP1** file that holds all the clicks in raw data.
- The **FP3** file that holds only clicks in click trains found by the KERNO-F classifier + information on the trains.
- ... both file types hold all the information on the angle of the POD to vertical and the temperature for each minute.

The classification process typically takes 1 minute per day of data.

The FPOD app contains many **display tools and filters** that are helpful in understanding the soundscape and cetacean click trains.

The FPOD app has flexible **export tools** to pass the data in various formats to spreadsheets, 'R' etc ...

Finally, the app enables you to change the clock time and settings of any F-POD.

The FPOD app display, filters and exports C-POD file data and is more powerful than the CPOD app, but that app is still needed for the processing of CP1 files to CP3 files using the KERNO classifier.

Files

On the SD card:

- A series of '.....CHE' data files, one for each day (from recent F-PODs) up to 4GB. If no errors are reported and the FP1 file looks as expected and is backed up then you do a 'quick format' of the card, and it is read for re-use.
If a data series is already on an SD card the POD will create a new file series with different series number. Both series will be read, and given distinct names.
- F_SETS.txt – this gives the F-POD new operational settings that are retained even with no batteries. The F-POD always runs with the last settings it was given, so it is normally run with blank SD cards.
- F_TIME.txt – this is a text file you can create via the Settings page of the menu to reset the F-POD to your local time.

On the PC:

- *.FP1 – data files generated when the SD card is read. They contain the clicks logged, the temperature and the angle off vertical upwards for each minute.
- *.FP3 - data files are generated when the FP1 file is processed via the button on the Trains page. They contain only the clicks that have been identified by the KERNO-F classifier as belonging to a train, each with a set of train descriptors. The FP3 file also holds some new minute descriptors on landmark sequences and ambient noise.
- a 'File set' **refers** to an FP1 file and the corresponding FP3 file. You can open multiple copies of each to view different parameters of each.
- *.f1m, *.f3m – these are map files that allow the app to move rapidly to any time, even in large files. These files must be in the same directory as the FP1 or FP3 file to be useful. If you delete them, No worries! they will be recreated if needed, so if you are in doubt you can discard them.
- *.txt – many export functions create text files that can be imported very easily into databases, spreadsheets and word-processors using copy and paste.
- *.arr – array files. These are large files generated from sets of FP3 files that have been named (or renamed) with their Y, X position on a rectangular grid, so that detections across the whole array can be displayed at speed.

Array files are opened by the 'Start array viewer' button on the 'view +' page

Cropping data files

Normally you need to remove the ends of the files – the times when the POD was not in the water, or the boat sonar was still being heard. This should be done once by the 'data manager' and the cropped FP1 file is then put in a folder called 'Useful files' which becomes the definitive data store for your project. All FP1 files from all sites and all years go into this one folder! Yes, it's scary! ... and some people hate the idea, but it works best.

To crop a file:

- Open a file set or just the FP1 file.
- Set a time selection enclosing the part you wish to retain – use the right-click pop-up menu.
- Click 'Create new cropped copies of open files'. The original files are not altered.
- Note: the FP3 file holds a record of the total number of clicks in the FP1 file, but these will now refer to the original FP1 file. If you re-process the cropped FP1, then this value is corrected, and the file warnings are updated.

Time selections are also a way of limiting the time range of a file that is available for analysis or export. They can be saved into the file header, and used or changed subsequently. Or time selections can be set as a temporary measure to enable you to quickly analyse that subset of data.

Pop-up Menu and Editing

Right-clicking in the Display area brings up this menu:

Select this file for action doesn't make anything happen! It just identifies which file will be the subject of various processes.

Put raw data into clipboard gives you all the detailed information on each click.

Explanations of what these are can be found here (add link)

Set selection ... These are the tools for setting a time selection that then limits what is shown, analysed or cropped to

Set train species These are editing tools. They work on either the train of the next click after the time at the pointer, or – if you select the 'NO' option that pops up – they work on all the trains shown on the screen. In that case it is important not to have any train that is split by the start or end of the screen. You can move the file a little way to the right or left using those arrow keys on the keyboard. The right key shows what is to the right of the pointer etc.

Send train to trash The train will disappear and can only be recovered by creating the FP3 file afresh.

Save comment file is very useful. Give the file a meaningful name and add .fpf to the name and it can be used to find your way back to the exact view you had.

Save time point / Go to A quick way of marking a point you want to go back to – like a bookmark

Go to next landmark sequence these are interesting series of click trains made by porpoises as they approach a landmark in a rather uniform landscape. It only works if some are found after searching for them via the Trains page.

Put name of source .wav This only applies to 'virtual FPOD files' synthesised from .wav files using a utility that is on the website. It only works if the sampling rate was 500k or more.

```
Select this file for action
Adjust Y scales
Put file name into results memo
  Put time into results memo
  Put raw data into clipboard
  Put wave data into clipboard
Set selection start
Set selection end
Clear selection
Set train species to NBHF
  Set train species to 'other cet'
  Set train species to sonar
  Set species to unclassified
Send train to trash
Mark this train
  Unmark this train
  Give train a type number
Save comment file
Save time point to file
Go to saved time point
  Go to next Landmark Sequence
Close this File
close
Put name of source .wav file on Clipboard
```

The FPOD app – Data display

Up to 6 data files can viewed together, numbered from 1 to 6 as you go up from the bottom, and they will always be shown with the times in synchrony along the horizontal axis. You can have 6 views of the same file, or the same view of 6 different files, etc opened via the *Files* page of the menu.

High resolution: the time scale is 100ms or less per pixel. Details of each click are shown, but there may be more than 1 click within the time span of a pixel. These displays show the structure of click trains and often allow you to see surface echoes, multipath propagation, evidence of the structure of the click beam etc. You can also view some of the fine time domain structure of each click, and more detail of some clicks.

Low resolution: the time scale is 1min or higher – up to 10 days. The data shown is *counts* of clicks or features of the clicks within that group of 1 or more minutes, and is shown in columns 3 pixels wide. These displays show the pattern of use of the location over time, and often show tidal and seasonal influences, storms, boat traffic etc.

Extra info can be show, or hidden (via the Display page of the menu)

The *angles* of the POD to vertical are shown as short lines at the top of the file 1 panel and give evidence of storms, currents, and movement of the POD. Whole histories of deployments that went AWOL can sometimes be inferred.

A *black line* on FP1 files shows the *average number of clicks logged per minute* in each time period.

A *green line* on File 1 shows the *operational threshold*. The amplitude threshold steps up when large numbers of clicks are logged over a few minutes. This can be switched off or adjusted via the Settings.

The *temperature* is displayed for the first FP1 and FP3 as a black and white line. To get a listing of the temperatures go to the Export page of the menu and use the ‘Summary of displayed minutes’ button. This also gives you details of the *battery voltages* and which battery stack was in use, and the *operational threshold* in use.

Battery voltage levels can be displayed for file1 only. It’s the option at the bottom of the drop-down list for file1.

Zooming



Select a new time scale by clicking one of these labels

Or:

Zoom in

- the keyboard down arrow: In the new view the time at the mouse pointer position will be moved to the centre of the screen.

or .. this is quicker for big changes:

- Left drag to right - a yellow line will appear showing the new range to be displayed. At the end of the drag the view zooms in to show it.

Zoom out


the keyboard up arrow. The new view has the time at the mouse pointer position at the centre of the screen.

Move earlier/later

the left arrow key shows data earlier than the pointer position

the right arrow key shows data after the pointer position

Data display – low resolution



The **hints** shown in red in the green bar at the top

This deletes the file selected for action and its associated set of files (FP1,FP3 and map files), even if they are not open. This is final. They do not go to the Recycle Bin!


explain the software item under the mouse pointer.

Some useful features:

- Waveform + spectrum will show you details from those few clicks on which the wavelength and duration of every cycle have been logged.
- Ignore all filters: extremely useful in high res mode. You can filter for whatever you like and then use this to see the full context. F4 is the shortcut.
- 'marked trains included' Trains can be marked by putting the mouse pointer just to the left of the first click in the train and selecting from the right-click pop-up menu. The Filters +files page has more on this. Viewing and analysing 'marked trains only' is very useful for studying behaviours. Click the 'marked trains included' text to step through the options.
- high species confidence only: The default is checked = on = true. Turning it off give detection rates a few percent higher, but mostly it's best to use the default setting – partly because you don't have to remember to turn it off!
- Species encounter classifier: The C-POD had two: GENENC and Hel1. If an open file was processed using one of these it will have the original KERNO classification plus a secondary classification using the named classifier that

offers one or more species classes. If any of these are selected they over-ride the KERNO classifications. Third party classifiers of F-POD data can also add species classifications.


- The small white lines at the top show the angles to vertical of the POD. You can see the deployment.
- The black and white line shows the temperature rise on immersion – this is near Odessa on the Black Sea.
- The black line shows the average number of clicks per minute.
- The green line shows the operational threshold of the POD moving in real time in response to the rate at which clicks were logged.



Skip to > 2

- This sets the minimum number of clicks on a high res screen, so the time on display moves forwards until a screen is found with at least this number. Skipping in an FP3 file to, say 512 clicks, at a resolution of 5ms will typically show intense feeding activity or social activity.



Data display – High resolution



What is shown:

You can select any of these parameters to display from each of the open files:

Frequency	<p>Colors of frequencies (kHz): (text is left aligned)</p> <p>Colour coded: 20 30 40 50 60 70 80 90 100 110 120 130 140 150 160 170 180 190 200 210 220</p>
Amplitude	<p>On an 8bit scale, 0..255, but this is compressed. Clipped clicks have extrapolated amplitudes.</p> <p>The default setting shows: <input checked="" type="radio"/> extrapolated <input type="radio"/> Pascals <input type="radio"/> linear <input type="radio"/> raw values</p>
Clicks/sec	<p>The reciprocal of the prior inter-click interval. The colour indicates the 'click-rate confidence' for the train– this is independent of the species classification confidence. Confidence level 0 has a significant error rate and above that errors are few and get less as the level rises.</p> <p>The lowest level is 0 and these train are shown in grey, higher levels are co loured yellow</p>

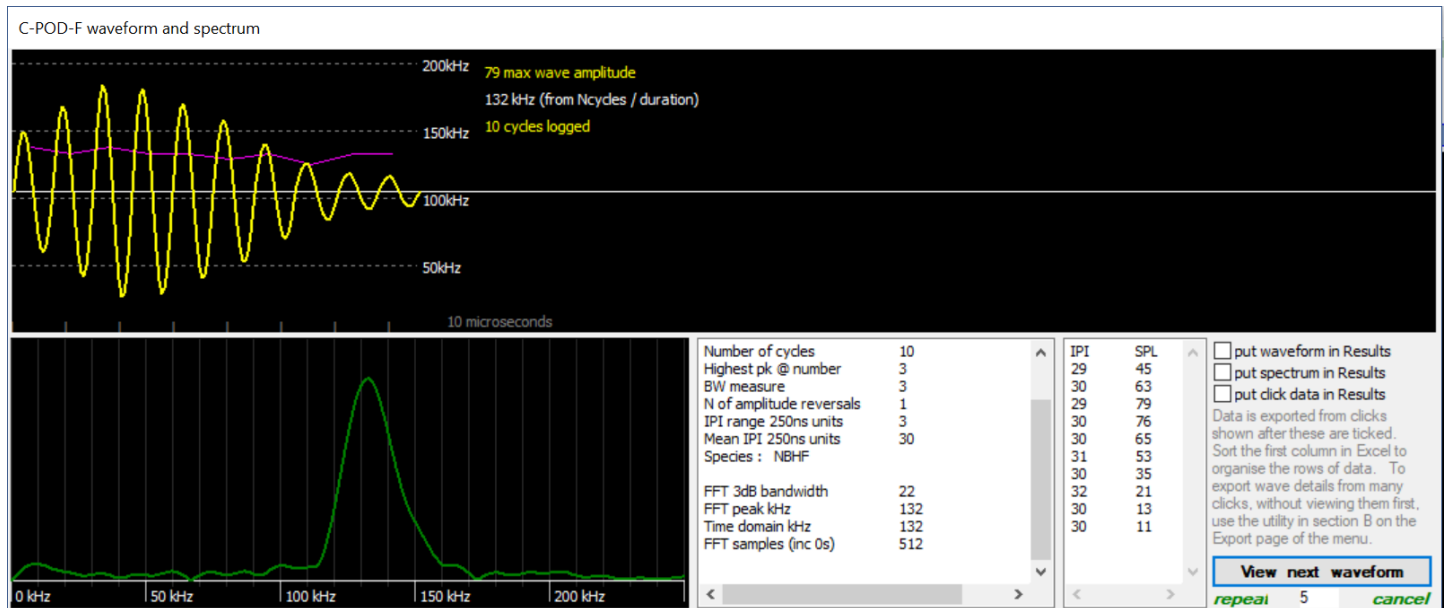
	<p>then red. You can filter levels on the Files +filters page in the Trains section here:</p> <p style="text-align: right;">Click rate confidence: Min Max 0 15</p>
<p>Train Q class</p> 	<p>The Q class is the ‘quality’ of the train = the confidence of it’s coming from a true train source, as opposed to a being a chance sequence of similar clicks at similar intervals. Normally only Q High and Moderate are treated as reliable, but it may be of interest at times to look at the lower Q classes. The lowest, called doubtful or ‘?’ includes trains that really belong to a higher Q class but appear to be composed of echoes of a higher Q train. They are ‘down-classed’ to avoid double-counting of trains with echoes.</p>
<p>Train Species class</p> 	<p>NBHF = narrow band high frequency – this species group includes all porpoises, some dolphins and pygmy and dwarf sperm whales.</p> <p>‘Other cet’ = all other cetaceans – dolphins, beaked whales, pilot whales etc.</p> <p>‘Sonar’ includes all sources that seem to be man-made.</p> <p>‘Unclassed’ these are trains with features that overlap two classes or lack any strongly discriminatory features.</p>
NBHF index	A metric representing how closely a click matches a typical porpoise clicks. It is based on frequency, duration (= n of cycles), bandwidth, and the position of the loudest cycle within the click.
N of cycles	This is used in place of ‘duration’. It is the number of waves, = cycles, in the click.
Inter-click-interval	Where a click has been inferred (i.e. it wasn’t picked up as part of the train) the interval used Takes into account the fact that the ICI is really two such intervals.
Bandwidth	An arbitrary scale derived from the variation in cycle periods and amplitudes within the click
Amplitude profile	The amplitude of a few cycle peaks around the loudest cycle.
Frequency profile	The frequency from the wavelength a few cycles around the loudest cycle.
Cluster size	The number of clicks identified as being part of the cluster of multipath replicates
Train WUTS risk	An estimate, of low accuracy, of the risk of this being a click train produced by a Weak Unknown Train Source.
Battery levels	Each battery stack is logged every minute, plus which stack is in use.

More details on these are in the document ‘FPOD click and train descriptions’.

Waveform display

Put a tick in the 'waveform + spectrum' box on the left side of the main display. You will then see a more detailed view of some clicks that have extra data saved. They are selected by a real-time train detection algorithm that is a simple but inaccurate method of selecting clicks for this special treatment. It does provide 'voucher' specimens of some clicks in many cetacean trains.

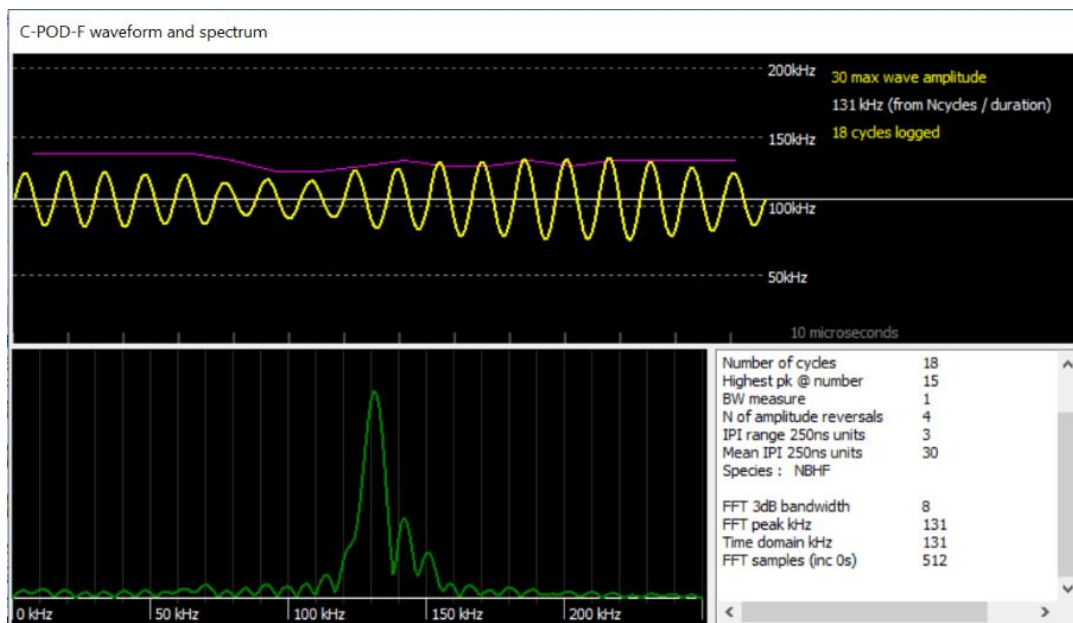
Here is the display of a click from a harbour porpoise:



The data stored was the amplitude and timing (to 250ns resolution) of the peak of each wave shown. The waveform, in yellow, has been constructed by fitting sine waves to the 10 points stored. This synthetic waveform has been sampled at 1m samples per second to provide the input data for the Fast Fourier Transform (FFT) that gives us the **frequency spectrum**, shown in green. The data has been expanded to 512 samples by zero-padding at each end.

The pink line shows the **cycle-by-cycle frequency**, which, in this case, shows a shallow downward trend. Actually it seems that this is common feature of porpoise clicks, which has previously been overlooked because the FFT spectrum gives no information about the sequence of frequencies within a sound. (The FFT does give phase information on each frequency but this is usually discarded because it is sensitive to the exact positioning of the input data set in the stream of values. Move a few samples earlier or later and the phase values jump around horribly so few analysts have found this information useful.)

The text boxes give largely self-explanatory information. It is notable that, for NBHF clicks, the time domain frequency estimation, which uses the length of the click and the number of cycles within it, gives the same frequency as the peak of the FFT frequency spectrum. The bandwidth (BW) measure here is an arbitrary construct but is more powerful than a conventional bandwidth measurement in its role in the classifiers used in the train detection process.



Menu

The 'menu' pages are shown by clicking on the large green text label, top right, or moving the mouse pointer to the top of the screen and down again, or clicking the tiny 'menu' button top left.

The *hints* shown in red in the green bar at the top explain most of the items when you hover over them, so this just picks out some other points from each page.

Read SD page ~~~~~ 1: *Location*: it is really important to give locations consistent unique names!

Page 2 describes

If your F-POD is 'acoustic release ready' it may generate text files with a name 'SendTo Chelonia...' Please do send them as we can then assess the risk of your acoustic release firing accidentally and tune it appropriately.

Files page

Tip: Get used to looking at whatever mixture of FP1, FP3, filters and display parameters helps you see what you want to see.

This is very useful – it provides a quick way to open whatever set of files you need to see the parameters you are interested in:

Open set: FP1 x1 + FP3 3 select either an FP1 or FP3 and this set will be opened

These buttons allow you to open sets of files, or even 6 different FP3 files so that you can then use the graphs function to compare them in various details – diel patterns, frequencies etc:

Open a file set

Open up to 6 unrelated files

clear time selection
save selection end times
delete saved selection end times

these refer to selections that you set via the right-click pop-up menu when viewing a file. They enable you to analyse parts of a file or create a cropped copy of it.

Filters +files page

Tip: List file times etc. is very useful in any project!

Tip: Use the Delete file set button to delete files – it gets rid of associated map files as well as the FP1 and FP3

All the text here has explanatory hints.

The filters are additive – to be shown a click or train must pass all the filters set, unless ‘ignore all filters’ is in force.

This section allows you to list all the settings of a batch of files to quickly check the uniformity of your project data collection:

Files: View metadata and classifier warnings:

List

<input type="radio"/> Classification Warnings	from
<input checked="" type="radio"/> Operational settings	<input checked="" type="radio"/> open F files
<input type="radio"/> Waveform capture settings	<input type="radio"/> open C files
<input type="radio"/> Instrument properties	<input type="radio"/> other F files
<input type="radio"/> File notes	<input type="radio"/> other C files
<input type="radio"/> Acoustic Release settings	
<input type="radio"/> Classifier settings	

List:

Classification warnings are now – in version 2 – valuable and you should look at them.

Set validation sampling points: The document ‘Validation of F-POD data’ explains this.

Deleting and renaming files is best done using these utilities:

File changes:

<input type="button" value="Delete file set"/>		
<input type="button" value="Append Files"/>	<input type="button" value="Crop sync"/>	<input type="button" value="Change start time"/>
<input type="button" value="Rename file set"/>	<input type="button" value="Crop"/>	<input type="button" value="Add m to start"/> <input type="text" value="1"/>
<input type="button" value="Rename 1 file"/>	<input type="button" value="FP3 to version1"/>	<input type="button" value="Add yr to start"/> <input type="text" value="1"/>

... because they rename or delete the associated map files.

Append files allows you to join files together correctly.

Problem files:

These utilities are rarely needed ... check ‘refresh maps’ and reopen any file where the FP1 and FP3 appear to be out of sync.

File repair: refresh maps

skip

errors:

report

hide

‘Hide errors’ is safe because any significant errors that require attention will still be shown.

'Redact open FP1s' is provided to meet the requirements of some military authorities. To use it do this:

1. Paste the list of times to be redacted into the Results memo:

LocationText

10.03.24 1734Z-1812Z

29.05.24 1755Z-1837Z

16.05.24 0353Z-0533Z

21.05.24 0200Z-0319Z

2. Open the relevant FP1 files, in time order so the earliest is File 1 (sort the files by name, so the earliest will then be the first) and use 'Open up to 6 unrelated files' and they will be opened in the correct order)
3. Click 'Redact open FP1s' Redacted copies will be made with RDCT at the end of the name.
4. Review the report in the Results memo.
5. If that is OK open any original FP1 and the RDCT file. View this in Low resolution and check that the process has been done correctly.
6. If the redaction is correct, delete the original FP1.

Please report any problems to team@chelonia.co.uk

Display page ~~~~~Tip: changing the display time width can be helpful.

The *hints* shown in red in the green bar at the top,

e.g.

This deletes the file selected for action and its associated set of files (FP1,FP3 and map files), even if they are not open. This is final. They do not go to the Recycle Bin!

explain the software item under the mouse pointer.

Several options here are provided to optimise displays for screengrabs to be used in documents. **Ctrl G creates an automated screengrab** of the data display area, with a white background.

Display line width: thicker lines can be useful when sharing screens.

Colour clicks by train number: in hi-res click amplitude displays clicks are normally coloured by their frequency. This option makes it easier to disentangle overlapping trains.

Show fast click rates/10 as a blue line: just to enable you to get a quick idea if the train is what you are looking for.

Put graph ... in results: For studies with many files exporting data to a database and analysing it there can be much faster.

Array viewer This can show an animation of detections made in a grid of loggers ... more details needed here. Contact us if you want to use this.

Times In the FPOD app times are stored as the number of minutes since the start of 1900 (as calculated wrongly by Microsoft ... so it disagrees with the correct value in R! by 1 or 2 days). This gives you access to those minute numbers.

File viewer *Show block n of any file*: this just lists the values in blocks of 512 bytes from any file. Can quickly show that a file is mainly blank/empty.

Click rate replay:

File name prefix:

Vaquita social call

.wav file duration up to 20 s

pitch 20

volume 20

add artificial click sounds

Reduce gaps to 1/8

This generates an audio file in WAVE format (filename.wav) in which the changing click rates on show are faithfully represented as changing tones. It's particularly interesting for social communications. You will need to try adjusting the pitch and duration to get a result that is meaningful to us.

auto-correlation

cross-correlation

For RESIDENCY try 1min time units

For TIDAL or DIEL patterns try 1h, 6h, or tidal time units, auto-correlation

Auto-correlation is useful for assessing the residency time of cetaceans, and for determining whether cyclic patterns are tide to days or tides. For residency time units of 1 or 5min are useful, for tides v. days time units of 6h or ¼ tide cycle are best.


Cross-correlation is useful for study movement between sites.

GRAPHS

The Graph button at the top of the screen graphs only the data displayed, and shows whatever you set here. By contrast the Analysis page includes the whole time range available, and has no graph options.

The available graphs here are these : * = most often used

- kHz*
- Clicks/s (PRF)*
- Duration (N of cycles)
- SPL (sound pressure level)
- Wavenumber of peak (Pk@)
- Bandwidth
- Terminal kHz
- Cluster size
- kHz profile
- NBHF index*
- Diel: N of clicks*
- Diel: kHz*
- Diel: Train clicks/s*
- Diel: Duration (N of cycles)
- Diel: SPL (sound pressure level)
- Diel: Av Train Duration
- N of clicks in Train
- KHZ.IPI prior
- KHZ.IPI minus1
- KHZ.IPI pk
- KHZ.IPI plus1
- KHZ.IPI plus2
- IPI trends
- auto-correlation PRF



The option to omit C/FP1 files (shortcut Shift F11) is provided to save time when the results from these files are not interesting.

Here's data from the coast of Georgia in the Black Sea, filtered to show Harbour Porpoises:


n = 19million = number of clicks in NBHF trains displayed

in the Clicks/s graph the horizontal, X, axis is on a log scale.


Here the data has been filtered to show 'other cet' i.e. dolphin species only.

Terminal kHz is the frequency of the last cycle in the click, and is a less precise (more compressed) metric, especially at high frequencies, resulting in graphs like this.

... could be useful for beaked whales.



While KERNO-F is in progress you see a display like this. If there are no detections it stays black but the minutes analysed, top left, shows it is still working.



Normally you just click '*Detect click trains in FP1 files -runs KERNO-F classifier*' and let your digital slave do billions of calculations for you. A separate document on special cases, like analysing pinger studies, will be provided ...

Generate Version 2 file warnings: This updates the file warnings – this version gives really useful warnings that can be viewed via the Filters +files page.

Export Fish tag data: PODs do store fish tag data that can be used to identify the tag number but this process is not yet in a usable form.

Show third party access form: third parties may be able to develop classifiers for species discrimination in specific locations or data sets, and this provides the interface for that. A rich set of features will be exportable and the results of your classification can be injected back into the FP3 file to give access to all the visualization, filtering and exporting.

Settings page

Settings are sent to the F-POD on the SD card in small text files 'F_SETS.txt' that are generated on this page of the menu. Having read the settings file the POD will retain those values indefinitely until it finds another settings file on an SD card that is inserted. It is a good idea to remove the SD card and delete its settings files after it has been inserted, and this is essential for the time setting file.

Changing operational settings of an F-POD

Settings for any F-POD:
Select battery type: Alkaline batteries
: in LF-POD
Continuous logging? continuous
Filter out boat sonars
Use automatic amplitude threshold control
Start on: 20 01 mm 1 dd 1 @ 00:00
[Click to hide advanced settings](#)
[Click to show acoustic release settings](#)

F-PODs have default settings that are the same as those shown unless the settings have been changed as below.
'F's run with their last settings until changed as below

Switch ANGLE for the settings file:
to: 77deg from: 0deg
100 0

Settings file description: Normal settings

Create an F_SETS.txt file

Select battery type: click the label to go through the options. This sets a minimum voltage for a stack of batteries, at which the POD will shift to the next stack on its list. The list starts with the weakest usable stack, so weak batteries are exhausted first in order to minimise discarding of part-used batteries.

Currently the POD housing has two battery stacks. The LF-POD has only one, so it can run them down lower until it fails.

You can override the defaults in the advanced settings by selecting 'user set switch voltage' and setting a new value in the 'Battery min voltage' box in advanced settings. The units here are 0.1V

Continuous logging?: Intermittent logging can be used to extend the running time. Power consumption in OFF minutes is low but the temperature and angle are still logged.

Filter out boat sonars: This activates an algorithm that runs in each minute and tries to detect up to two different boat sonars from the similarity of the IPIs around the peaks of successive clicks. If found, then a small range of click frequencies is excluded for the rest of that minute. The frequency range 117-160 kHz is excluded so that the sonar filter will never exclude NBHF clicks, but relatively few sonars are in that range – the commonest frequencies are 50, 80, 100 and 200 kHz sonars. If one of two sonars are found this is recorded in the data and in high resolution displays of click frequency the time and frequency range of the exclusion band activated are shown as a dark grey box.

The value of this filter is that it greatly reduces the risk of boat sonars 'jamming' the detection process, forcing the amplitude threshold up, and filling the memory.

Use automatic amplitude threshold control: This raises the minimum amplitude of a click that will be logged if there are a number of successive noisy minutes and lowers it after some quieter minutes. These threshold levels are shown by a green line on low resolution displays and can be exported. You might wish to disable this for short specialised studies with high limits on numbers of clicks logged, but it is a good idea to normally have it in use as it removes the 'maxing out' problem of C-PODs with its knock-on effect of impaired train detection.

Start on: this allows you to delay the start to some date. You could use this to 'daisy-chain' multiple PODs in a remote location to obtain full coverage of one or more years.

Click to show advanced settings:

Advanced settings:		Filter	80kHz	Base deadband	5	Wav min amp	10
Min nCyc	5	Max 1 Bias	17	Use ADF filter below	120	Wav min step	7
NoisyMins	3	Max 2 Bias	21	Use BW filter below	0	Wav train SNR	15
QuietMins	4	Max 3 Bias	30	BW tolerance	3	Wav ICI	62 to 65
NoisyLevel	20	Max 4 Bias	50	Envelope filter below	110	N in train	3 to 15
QuietLevel	8	Stmax	220	Sonar trigger Ncyc:	14	Wavs 0.25GB limit	4
AmpStep	5	Stmin	18	Fixed exclusion minIPI	255	Wav / m	1536
Steepness	4	Min Pmax	12	Fixed exclusion maxIPI	0	<input type="checkbox"/> Wav record all clx	
Limit on clicks/m:	<input checked="" type="checkbox"/> use reverb filter			exclude 50kHz			
16k, 1 block	Thermal adjustment <input checked="" type="checkbox"/>			Battery switch @	48		

Set for:

minimal noise

normal operation

severe noise

NBHF only

tank tests

river dolphins

Normal settings

The basic advice here is: Don't change any of these! The defaults are included in the F_SETS file. Some pre-set options are included:

'tank tests' cancels various filters that aim to exclude weak broadband noise, the threshold adaptation, the sonar filters, input analogue filters and the angle switching. Your tank tests may still give misleading predictions of performance at sea for several reasons, e.g. generating a 130 kHz click with a square envelope and 8 cycles may be translated by the behaviour of the signal generating system into a click with the loudest cycle at the start. That would be correctly rejected by the KERNO-F classifier as not being an NBHF click, although the spectrum of such a click would probably pass FFT tests of spectrum frequency and bandwidth of an NBHF click.

'severe noise' is suitable for rough shallow water, or high levels of shrimp noise. Ideally you should compare two F-PODs with different settings in such locations, but this requires cetaceans to be common.

'river dolphins' aims to reduce the impact of noise from moving sand. It is not needed in muddy rivers.

Fixed exclusion minIPI	72
Fixed exclusion maxIPI	88
50kHz excluded	

Where 50kHz filters are very prevalent clicking the 'exclude 50kHz' button gives these settings which remove the need for the sonar to detector to find the sonar in each minute. The sonar detector is then able to find and exclude up to two other sonar frequency bands. You can set any pre-excluded band of this sort using the 'fixed exclusion' settings. These are defined as inter-peak-intervals = wavelength in 250ns units = 4000 / your target frequency. Allow plus and minus 10% of the frequency for the min and max.

Other controls:

Wav record all clicks: overrides the normal process in which likely cetacean clicks, assessed by being in what looks like a train, are the only ones that have full waveforms stored. This does not affect detection performance but the volume of data will increase substantially so it is only suitable for short deployments.

The other controls are best left alone because it is easy to end up with files recorded with different settings and then you need to know how much this may have affected the results, and that can only be determined empirically in sea trials. Most of the controls have 'hints' visible at the top of the screen.

Many sea trials were run to arrive at the default settings. Do not follow advice in any published paper to vary these settings without consulting Chelonia first, as inappropriate variations in C-POD settings or classifiers have been proposed in published papers. These have been proposed on the basis of obtaining improved detection rates in samples of data that are far too small to give good 'generalization performance' and the consequent errors can be large.

Switch angle:

This is used to allow you to start an F-POD many days before deployment and keep it horizontal and asleep so that it uses little power until it is deployed. If there will be no long delay then do not use this but force the POD to be ON at all angles. This is because angle settings can lead to data loss if it is deployed at an OFF angle or gets into one by mistake e.g. on a rig that falls on its side etc. Deep F-PODs or F-PODs for beaked whale monitoring may be deployed 'upside down'. The angle options do allow this but should be tested to ensure they are working correctly before committing the POD to the

Switch ANGLE
for the settings file:

to: 180deg from: 0deg

255 0

lid

ON OFF

deep.

The diagram represents a POD with its lid at the centre and its transducer housing at the outside.

Acoustic Release

... is not yet available. It screws on in place of the standard lid and allows a restraining line to be set

Acoustic release:	20 yy	MM	dd	hh
Do not release before:	99	1	1	1
Force release on:	99	1	1	1

free when a POD- specific acoustic code is received from a deck unit. It relies on the continuing activity of the POD itself, so it has a 'Force release on ...' date. A GPS to phone system tracking system is planned to allow recovery after such an event. It also has a power saving 'Do not release before' option.

RESET the F-POD CLOCK

The clock is set to GMT = UTC. You can reset it to your local time. Follow the instructions on the Settings page. An error of less than 100ms is easily achieved. Accurate setting is particularly valuable if you wish to compare data from different PODs in detail as events will then be visible on the same high resolution screen.

Navigation page ~~~~~

You can use these scaling factors to give each data panel graph exactly the range you need – most often F5 and F6 get you near by doubling or halving the current scale. The factors here give more precision if you need it.

The counts value also scales the height of the black line in FP1 displays that shows the number of clicks logged.

About page ~~~~~

The versions of various component of the app are shown on this page.

The KERNO-F version is 1.0 and will remain the same and not be subject to minor improvements. At some point a version 2 that is substantially more powerful will be introduced as an alternative option.

The changes in the various component versions are listed in the FPOD version history on the Chelonia website.

Delphi source code and coding notes: This gives the code used to unpack the raw data from the FP1 in this actual version of the software, plus notes on program.

Export page ~~~~~

Tip: the 'search directories' function is very useful

All items have info shown in red in the green bar at the top.

Read SD Files Filters +files Display Trains Settings Navigation About Export Results Analysis

Data Export Train filters: Q: Hi,Mod,Species: NBHF,'Dolphins',sonar, Good Sp ID only
 No click filters

use current time selection use end times saved in each file include annotations
 search directories for: FP1 FP3 CP1 CP3 Lat Long as decimal

A Data:
 Detections and environment
 DPM (detection positive minutes)
 DP10M (detection positive 10 minutes)
 DPH (detection positive hours)
 N of clicks
 Train duration total
 Inter-click intervals
 Train ICI inflections
 Minute Modal ICIs
 (dev item)

Period:
 1 second
 15 seconds
 1 min
 10 mins
 hours
 1/4 tides
 days
 all
 omit negative time periods

Output to:
 text file
 Results
 tab separator
 csv files
 separate text files

file name:
 include headers
 clipped file name
 site code only
 file name + path

Date format: d/m/y y-m-d neither MinN Julian day neither

Export from open file: 1 2 3 4 5 6

B F-POD Train details Full train details Wave details Click details. 1min in:

C F-POD Encounter details MinGap 10 MaxLength 10 minutes graph

D Summary of displayed minutes List file times List sonar filter triggers

'Detections and environment' combines the most useful data for many purposes. Includes temperature and angle values

Tide periods can be offset - see display page.

Inter-click intervals exports successive ICIs for all filtered trains = a large volume of data. Best to filter for high ICI quality (Filters +files page)

To export file warnings from many files go to the Trains page. For displayed files go to the Filters+file page

Detections and Environment: This adds information to each logging period about the angle, logging threshold, and some measures of the soundscape. It's a good plan to export this first and then drop other specific exports into your spreadsheet alongside it, so you keep the time units (the 'Period' 1,10, 60minutes etc).

Choice of operational statistic: There are four important principles here:

7. Larger units are preferred because they show less variability between time periods, sites, and deployments than smaller units. These are each larger than the one before: NofClicks/day; DPM/day (= detection positive minutes); DP10M/day (= detection positive 10 minute periods); DPH/day. Monthly values may be preferred.
8. But the unit should not exceed roughly 30% of the maximum possible or it will be significantly underestimating the highest rates. If you are near that, move to a smaller unit.
9. NofClicks conflates behaviour – like the use of very fast click trains in social communication with presence. Using *Train duration total* avoids that and is preferable.
10. Counting the number of trains is not a good idea. Trains in the POD data are mostly fragments of much longer real trains made by the animal and their duration and number reflects how fast the animal is scanning its environment i.e. it conflates rate of angular movements of the animal's head with presence.

Inter-click intervals gives you large volumes of detailed information on this if you wish to study it.

Train ICI inflections is a research tool to study the structure of click trains.

Minute Modal ICIs is a research tool to identify social communication by finding minutes with a sharp mode in the distribution of ICIs.

(Dev item) this currently exports the distribution of the rates of change of smoothed ICIs (the slope of clicks/s graph) at different ICIs. It is another research tool for the study of species differences in echo-location and social communication.

There are various other options on this page. In each case there is useful information on what they do in the red text at the bottom of the FPOD window when your pointer is over the option.

Access third party analysis tool: Species classification is a 'clustering' type of problem and well suited to various machine learning methods, while the train identification is a different type of signal processing problem for which few existing machine learning methods are useful. Here you can export click details and train details and some powerful KERN-F derived features. Then there is a simple interface to allow third party classifications to be written back into the FP3 file so that all the display, analysis, filtering and exporting functions are available on this new classification. The names for new and different classification groups can also be incorporated and used by the FPOD software.

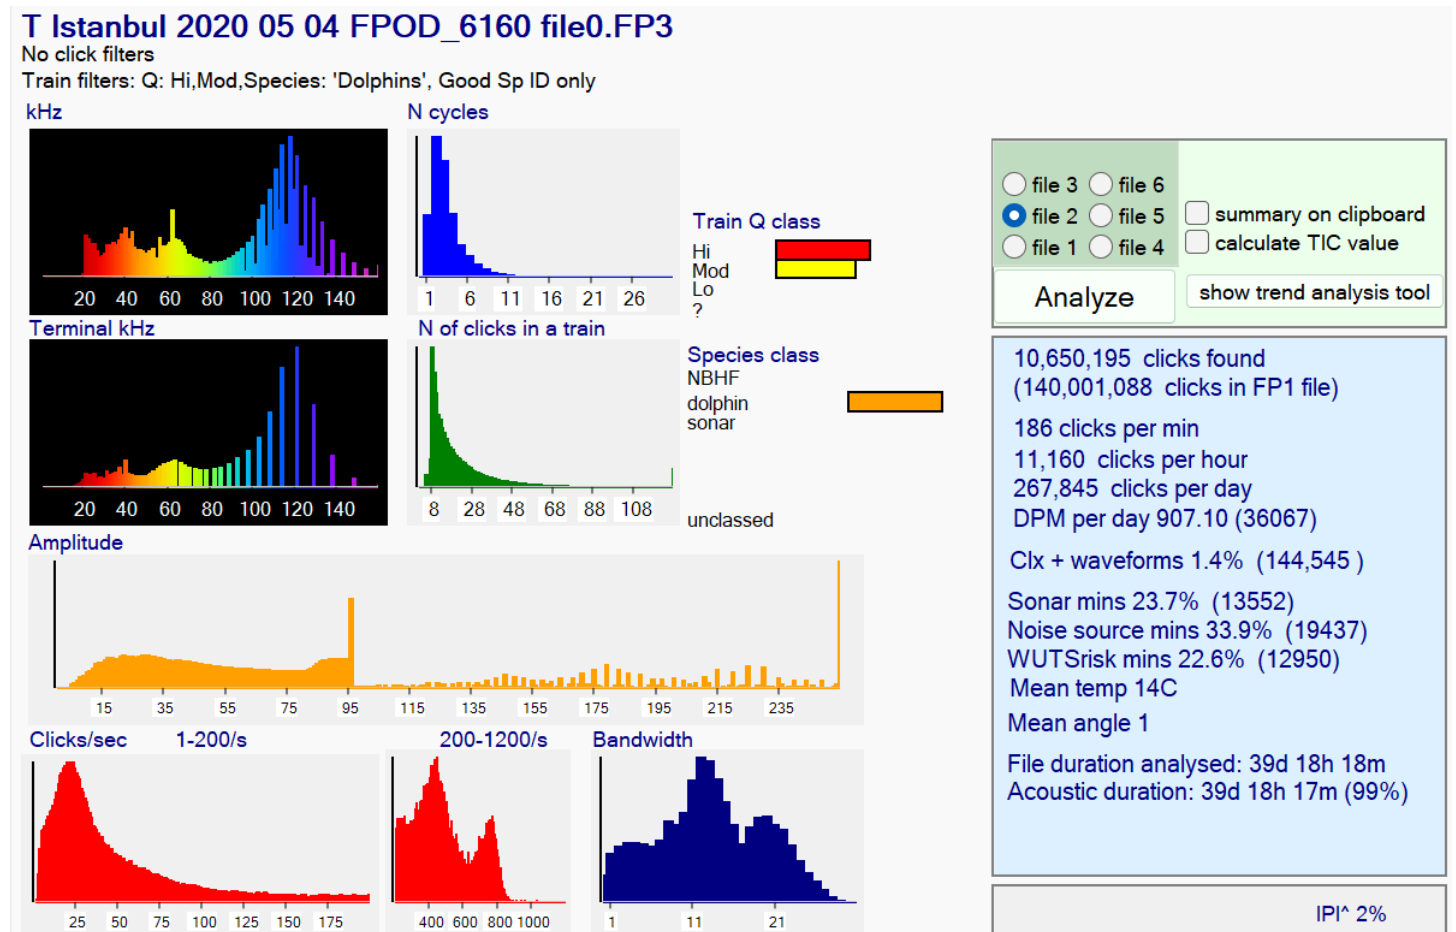
Results page ~~~~~

This is a main pathway for exporting data. Cut to clipboard is very useful because it clears the Results memo.

Add sightings to files: This is a powerful tool. You can, for example, find filters that identify the presence of sounds from a pinger or sonar, and use that to identify those minutes in an exported list of minutes; then format that as in the template spreadsheet '*FPOD sightings entry.xls*'; paste the relevant columns into the Results page and read them into the file. Then you can filter minutes by the presence or absence of this feature. Or the list might be of species sightings.

Analysis page ~~~~~

This gives you a quick view of the whole of a file, in contrast to the graphs you get when you use the 'graph' button at the top of the screen – the 'graph' button shows you only the data on screen.



If an FP1 file is analysed the last entry in the results list is how long it would take to fill a 32GB SD card at the rate found.

The Clicks/sec graph is split to show the range above 200/s upscaled by x10 as otherwise interesting variation here is lost.

The frequency graphs have a gappy appearance at high frequencies due to the resolution at which this data has been stored.

The IPI^ value is the proportion of clicks in which the sum of the two wavelengths before the loudest cycle re shorter than or equal to the next two. This is typical of NBHF clicks.

The TIC value is the 'train interval clustering index' – a measure that is higher for cetaceans than random events or noise.

Amplitude graphs can have a saw-tooth profile due to the effect of the time-domain filters that exclude shorter clicks of higher bandwidth in a series of steps. These filters, plus the very compact mode of data storage enable the F-POD to run for long periods and provide 'high value' data as inputs to a powerful train detection process that uses the time domain features of clicks.

The Trend Analysis Tool

The tool determines trends in the detection rates from PODs at one or more fixed sites.

It uses 'Paired Year Ratio Assessment' = PYRA.

This method is described in 'Estimating cetacean population trends from static acoustic monitoring data using Paired Year Ratio Assessment (PYRA)' Grist *et al*

<https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0264289>

The distinctive advantage of PYRA is that it avoids the need for estimation of both seasonal and diel patterns. Cetaceans often show big seasonal and diel patterns of activity and their estimation from gappy data is a major source of error in some other methods.

How it works:

Paired Year Ratio Assessment requires 2 or more years of data and handles gaps in the data by using only data from days that can be 'paired' or 'matched' with data from the same day in the following year. So the sum total of the detection metric in the second year is divided by the matching total in the first year to give a ratio that represents the growth or shrinkage – the 'change value' - of the data logged.

This two-year data window is moved forward one day at a time to produce a change value for each date. The date for this change value is the date between the two years. If the data window is moved further the number of matched days eventually falls as the second year becomes more incomplete. When the volume of detections becomes very low the change value becomes more and more wild! That's also true at any time within the body of the data when there are few detections, so to avoid wild results a small 'ballast' is added to both years. This has the effect of pushing the change value towards 1.0 = no growth.

The process ends when there is no longer any paired data in the Y2 period. It can also start when there is only 1 paired day, but these 'thin ends' are of little value. So change values based on less than a mean of 1 DPM (detection positive minute) per day are not graphed ... but all paired data is used in the assessment of overall trends.

A measure of uncertainty is obtained by random resampling of whole days, with replacement, within a 21-day window centred on the paired day.

What does it tell us?

Sampling the trend in cetacean activity at a number of sites is often a very much cheaper way of getting a handle on what is happening to the population of a cetacean than doing repeated line transect surveys of its entire distribution from boats, planes or drones. Much larger numbers of detections can be made within the lower cost, and they give much greater precision. So what's not to like?

The sites chosen will usually cover only a small part of the distribution of the species (the monitoring of the Vaquita is a notable exception). So redistribution of the animals is potentially a major issue. Has the species declined or has it moved away? Has it increased or become more concentrated in the area logged? Habitat type will often be a factor in redistribution, and analysis of trends at logging sites with differing characteristics may be relevant (the app allows you to select and view subsets of the sites and to export data from all sites individually). These questions need to be considered and related to other biological or physical data, so that these trend results form a part of that larger picture.

The bottom line is: the trend values are simply the trend seen in the set of sites, and their predictive value for the 'local population' and the whole population needs to be determined using other data which might include prior data from such trend monitoring elsewhere, data on fish stock changes, physical oceanographic data, etc.

Limitations

Gappy data: when logged periods don't align across years the un-paired days are not used. A lot of data can be 'lost' this way, so if the whole year cannot be logged it's good to log the same parts of successive years.

Site stability: sites should be fixed, and any change in position of more than 200m should be treated as a new site. This is because we know that detection rates of porpoises vary significantly at this scale, even when the habitat looks superficially fairly uniform. (A fully randomised spatial design is theoretically possible and would use only the overall statistics for each year but would require many more sites).

If a site changes – perhaps construction works starts nearby, or it is a maturing aquaculture site – then a local trend may be due to those changes, so it should be excluded from the overall trend analysis, but its own trend may still be of real interest.

Grist *et al.* (2022) discuss the limitations of PYRA.

Steps in using the trend analysis tool

In brief:

1. Give your data files site codes that identify each site correctly. If they have no site code the early part of the file name can be used and must uniquely identify each site in the study.
2. Crop all files to remove any time not in the water.
3. List all file times, (see 'Filters +files' page) then sort by dates of start and end to identify the dates that should be used in the trend analysis.
4. Enter the earliest and latest days, of any file, in section 1 and also any exclusion period if part of each year was not normally logged or was subject to something that might interfere with the normal distribution of the cetaceans:

5. Select a detection metric. DPM – detection positive minutes is the default. Larger units may 'saturate' and it's a good idea to use only units that do not go above about 30% of the maximum value they can have in the (which is 24 for DPH, 144 for DP10M, 1440 for DPM etc.) in the present data set or in the likely future of the project. Nclx may be conflate presence with feeding or social activity.
6. Click 'Confirm Time Frame and exclusions'
7. Set the filters for the input data e.g. Hi and Mod Quality, NBHF species, High species confidence. You could do assessments where you filtered out fast trains or slow trains etc. if you wished. If you are reading the data from a text file the filters will be those in force when the data was originally read.
8. In section 3 select where you will import data from.
9. Select 'Define sites by' unless you are reading from a text file.
10. Click Confirm filters.
11. Click 'Select files for data import'. The number of files and sites found will be shown. Later you can export a text file of the data and read that if you wish to save time.
12. Generally keep defaults for the resample runs, bin size and ballast more in Notes below.
13. Click 'save settings' at the top of the form and give it your project name.
14. Click Analyse. File reading and processing are followed by a 'View graphs...' popup message.
15. Now it gets interesting! But before you get into viewing the data click 'Put site list on clipboard'. The list will be on the clipboard so you can paste it directly into a spreadsheet and check that your files are correctly grouped into sites. If any have the wrong site code that can be changed via the File changes section on the Filters +files page – and you'll need to repeat the process above. If section 5 has disappeared, click 'Show PYRA controls'

Notes

The **resample runs** sets the number of times each paired day will be resampled – see ‘How it works’ above. If you set the resample runs to 1 there is actually no resampling - your original data is used. Increasing the number of resampling runs give a smoother distribution of change values. Larger numbers are a good idea.

Bin size: Each resampled PYRA change value will be added into a bin to produce percentile ranges and a histogram of results. The size of each bin can be set using the bin size value. Change values beyond the value of the most extreme bins will be added to that bin.

Ballast intensity

This is explained in ‘How it works’ above. The default is to add a ballast that is 5% of the overall sum of detections to both the year before and the year after the time point for the change value. It’s effect on the overall trend is very small, but it heavily damps values where very few detections are present, pushing the change value towards no change.

Data collected

The F-POD data on the SD card consists of a header that includes information about the POD, its settings, and the time of the start. On the PC all that data is retained, and there is more space in the header for other information, including the user entries on the Read SD page.

The main body of the data consists of minute timestamps, click records, and waveform records. Each takes up 16 bytes. The click record structure is (IPIs = inter-peak-interval in 250nanosecond units; amp = peak amplitude of a cycle) :

Byte number	Click record
0	time in 5us within minute, MSB (most significant byte, max value = 183) of loudest cycle.
1	time in 5 microsecond units
2	time in 5 microsecond units, LSB (least significant byte)
3	N of cycles in click. max=255
4	Pk@ 4bits; IPI range within click 4 bits (Pk@ is the wave number of the loudest wave)
5	IPIpreMax
6	IPI Max
7	IPI post1
8	IPI post2
9	amp of cycle before peak amp
10	amp of peak (=Pmax)
11	amp of cycle after peak (P+1)
12	IPI before first cycle
13	Duration, most significant 4bits : AmpReversalCount (in lower nibble)
14	Duration, least significant byte
15	IPI of last Cycle

The amplitudes are in a compressed 8 bit format. Clipped waves have their rise time, from zero-crossing to clipping stored and this is extrapolated to the true peak. For heavily clipped low frequencies this goes up to 12,800. The whole scheme can be exported via the Exports page.

The code for reading each 16 byte record can be accessed via the 'About' page of the menu of the FPOD app.

These data give a rich time domain description of each click that is used to give high discriminatory power to the train detection process.

Some clicks have up to 21 cycles of 'raw data' that is intended to provide more detail for various purposes. It consists of a value of 250 as a record type identifier, followed by a series of pairs of values giving the amplitude and period of 7 successive waves.

Minute timestamps carry information on temperature, the angle of the POD, the level of the threshold setting, the voltage of the battery stacks and which is in use.

Warning: POD dates

These give the same value as you get in Microsoft's Excel, Access etc, but:

These are based on the number of days since the start of 1900 and give a different value in 'R' which gives a 1 or 2 day difference. 'R' is correct!

Thanks to Leah Crowe for identifying this problem.

Topic Notes and glossary

Absorption	Sound is absorbed during its passage through seawater at a rate that is higher for higher frequencies of sound. The characteristics of broadband clicks consequently changes during transmission with the higher frequencies being reduced more than the lower frequencies. This difference increases with the distance from the source.
ADCP	Acoustic Doppler current profiler. These typically operate at 1Mhz or higher but often produce lower frequencies that may affect the distribution of cetaceans and be detected by C-PODs.
Amplitude	The height of the pressure wave of a sound. Also expressed as SPL = sound pressure level. This is used here as a simple instantaneous pressure scale value, but in some systems, it refers to a value in decibels of the RMS value of the pressures of the whole wave.
amplitude profile	The sequence of heights of the waves in a click. All clicks have three amplitudes logged – the loudest and the one before and after. Clicks with full waveform data have more. The amp profile of NBHF clicks and boat sonars is fairly flat. That of dolphin clicks and many non-cetacean sources is much further from flat.
AmpReversalCount	This is the number of times the cycle amplitude trend (whether a wave peak is higher or lower than the one before) within a click switches from rising to falling. It is assumed to be rising at the start. It is lower in cetacean clicks than noise and lowest in NBHF clicks.
Bandwidth	A measurement of the frequency width of the FFT spectrum. The width at half the maximum is often used = the 3dB bandwidth. Concatenation of reflections of a click, or of refracted replicates will make the click longer and this reduces the bandwidth. The F-POD exploits a different approach to bandwidth by logging a high precision measure of the variation in length of the waves around the loudest wave in the click. These values are least affected by ambient noise and are independent of the length of the click.
BBT	Broad band transients. Dolphins that whistle make short clicks. All short clicks have a Fourier spectrum with a broad bandwidth, so 'broad band' is predicated by 'transient'. NBHF clicks are longer. Longer clicks could be broadband or narrowband.
Beam structure	Clicks are emitted from the cetaceans' melon as a narrow beam of sound that is loudest near the center. In dolphins the mix of frequencies changes strongly across the beam, with low frequencies dominating the outer edges of the beam, and vice-versa. This can often be seen in POD data, with surface echoes revealing the frequency content of a different part of the beam from that forming the direct path to the POD. For NBHF clicks there is little change in frequency content across the beam.
Bedload	See sediment transport
click	See event. Click is normally taken to imply a short episode of increased amplitude. Whistles, by contrast, are not short and are not heard as a click but as a tone that may sweep up and down in pitch.
Click rate	The number of clicks per second. Also called PRF – pulse repetition frequency. It is the reciprocal of inter-click-interval.
clipping	The pressure of a loud clicks may reach the maximum output level of an amplifier or ADC (analogue to digital converter) before it has reached the peak of that cycle. It is then 'clipped'. The time taken between crossing zero and

	clipping can be used to extrapolate the peak value.
clusters	See multipath
coherence	See train coherence
crop	Make a new shorter file that is a copy of part of a file with time removed from one or both ends.
cycle	One wave of sound pressure going from zero up, then down below zero then back up to zero.
Date format	MinuteN refers to the minute number in FPOD time – it is the number of minutes since the start of the year 1900. Divide this by 1440 to get the time used in Excel and Access, which is a floating point number representing the number of days and fraction of the last day since the same start date.
dB	Decibels are intended to represent intensity (= power) of a sound and is 10 times the log to base 10 of the ratio of rms values. It is a long out-dated unit that is rightly deprecated by the SI system of units. It is not a physically precise unit of comparison because that depends on the acoustic impedance of the media being identical, it is troublesome to learn and it contributes to errors! The alternative is to report pressures directly as Pascals or rms pressures.
Detections and Environment	This adds information to each logging period about the angle, deadband, time lost (to 'maxing out' = reaching the limit for the minute) and some measures of the soundscape.
Doubtful	See Q class
DP10M	Detection Positive Ten Minute periods. These are 10minutes set by the clock, so their boundaries always correspond with the hour ends. This unit reduces the effect of variations between loggers but will saturate (reach the maximum possible) before DPM if the density of animals is high.
DPM	Detection Positive Minutes. This is a good general purpose measure of the level of acoustic activity. It's not the best for every purpose and more information is on the website.
dur / duration	Duration of clicks in FPOD data is often represented as the number of cycles in a click. The duration in 250 nanosecond unit is also stored and used to calculate the frequency of the click.
EDR	Effective detection radius. This is the radius of a circle that would contain as many animals as were detected in total i.e. the number missed inside this radius is the same as the number detected outside it. There should be a time period specified i.e. 'the 1 minute EDR' as it will vary with this time period. It is not the same as the mean distance at detection, which is typically much larger.
effective detection radius	See EDR
elephant outside the room	See the discussion of Notes and Warnings
envelope	An imaginary line stretched over the amplitude peaks of the waves in a click.
error estimation	See Warnings
event	'Clicks' are detected by PODs as short periods in which successive wavelengths are similar. These are tonal 'events', and there is no requirement that they must be louder than the sound stream before or after, although they usually are.
Export SPL as Pa	Ticking this box will show the amplitude of each click in Pascals by correcting for the frequency-dependent sensitivity of the system and the non-linearity of the amplitude scale. However, it makes it less clear when a frequency has

	reached and may have exceeded the maximum response for that frequency.
False negative	Clicks made by cetaceans that were missed by the detector. Because it includes clicks from cetaceans at all distances from the detector the idea is better represented by the 'detection function' that describes the proportion of false negatives in relation to the distance from the detector.
False positive	Clicks identified as a coming from a cetacean that did not. In the case of trains the concept is a bit different as the inclusion of a single click that did come from a cetacean means that it correctly represent the presence of a cetacean.
FFT	Fast Fourier Transform. This takes in a set of amplitudes, measured at uniform intervals, and generates half as many complex numbers. One of the pair of values in each complex number represents the intensity of one frequency within the set of amplitudes. The set typically has 256, or 512, or 1024 (or other power of 2) values.
Fish tag	See VEMCO fish tag
Frequency domain	The information within the frequency spectrum. This includes no information on the order in which frequencies appear in a sound.
frequency profile	The F-POD stores the wavelengths of the successive waves in a click. These give information about the frequencies in the click and their time sequence, and can be graphically displayed as the frequency that would have that wavelength in a pure sine wave. Frequency profiles cohere in trains. As with the amplitude profile it is flatter in clicks from NBHF sources and boat sonars than cetaceans and noise sources.
generalization performance	The success of a classifier when applied to data that was not part of its training set. Some papers have based advice on better POD settings on limited studies providing a training set only. This is risky, at best, predictably bad at worst.
guard band	A simple technique of detecting a narrow band signal is to compare the frequency of interest with a 'guard band' of higher or lower frequencies. The very first POD used 3 guard bands with used defined threshold ratios for each.
Hi, Mod, Lo, Doubtful	See Q class
hints	See the text in the green bar at the top of the FPOD app screen for hints on the function of whatever is under the mouse pointer.
ICI	Inter-click-interval. The time between the starts of successive clicks in a train. For cetacean click trains this ranges between 1 or a few seconds and less than 1ms
Include headers	Don't include them if you are exporting from multiple files to go into a database or spreadsheet.
IPI	Inter-peak-interval. The time interval between the peaks of the waves in a click. In the F-POD this is measured in 250ns units.
KERNO	Chelonia's automated classifier for data from C-PODs.
KERNO-F	Chelonia's automated classifier for data from the F-POD.
kHz	1000 cycles per second. Human hearing in children goes up to 20kHz which is the upper limit of the audio range, with higher frequencies being called ultrasound.
landmark sequence	A linear trend of falling inter-click-intervals that may be seen when porpoises approach an acoustic landmark from 250m or more away.
Lo	See Q class

MaxPk	The SPL of the loudest cycle in a click
Mod	See Q class
multipath	<p>Loud clicks are reflected structures within the cetacean and then, perhaps repeatedly, by the sea surface, the seabed, and possibly other structures. The sound beam is also bent this way and that by differences in sound velocity created by differences in salinity, temperature and pressure. These refract the beam and create a heterogeneous wave front, so hydrophones at different positions do not receive identical signals even if the beam was originally uniform in cross section, which it isn't. The surprising thing is that these pathway effects change over remarkably short time periods.</p> <p>The multipath cluster to some extent reflects the whole cross-section of the beam so it has a wide spread of frequencies when the source is a BBT click and a narrow spread when the source is an NBHF click.</p> <p>The coherence of cluster characteristics is useful in train evaluation.</p>
nAll	The number of all clicks in one minute in a CP1 or FP1 file.
NBHF	Narrow band high frequency clicks are made by porpoises and a few dolphin species. They are usually around 130kHz and relatively long – 8 to 12 cycles with a typical amplitude envelope. The species that make NBHF clicks do not whistle.
NBHF index = NBHFi	A numerical index, used in POD software, of the resemblance of a click to a typical NBHF click. This is derived for each click in the CP1 or FP1 file but is much more powerful (i.e. discriminatory) in FP1 files because they include the successive wavelengths and amplitudes around the loudest cycle of the click. Clicks 'source code ...' on the About page for full source code for this classifier.
Ncyc – number of cycles	The number of cycles in a click. Typical mean of NBHF clicks = 8, and is much lower for BBT clicks from dolphins.
Noise levels	F-POD data is not suitable for measurement of noise levels. This needs a broadband WAVE file recorder with an appropriate frequency range and anti-aliasing filtering.
Panguard	A widely used open source app for the analysis of wav file data to detect, classify and localise cetaceans.
Pascal Pa	The basic unit of pressure: 1 Pa is 1 Newton / sq. m. Hydrophones are pressure sensors that give a voltage that is proportional to the pressure. The weakest porpoise clicks detected by the POD have a peak pressure of around 1Pa.
PeakAt Pk@	The number of the cycle in a click that has the highest amplitude. Typically 1 for dolphin clicks and 4 or more for NBHF clicks, and very variable for boat sonar pulses.
Phase	Two waves are 'in phase' if their peaks, at some point in space, occur at the same time. If they are of equal amplitude but come from slightly different directions, they will at some places be in phase with increased amplitude, and at other places they may cancel out. Such constructive and destructive interference is common and is sometimes easily seen in POD data.
Play sound	The 'play' button translates inter-click intervals into tones that rise as the click rate (1/ICI) rises. It uses the amplitude and smooths the transition between ICIs
PriorIPI	The inter-peak-interval before the designated start of the click. In some dolphin clicks it may have been part of the click but not defined as such by the algorithm running in real-time on the POD that tries to define the start and end of clicks.
PYRA	Paired Year Ratio Assessment. This is a simple robust method for estimating trends in multi-year data sets. It is the basis of the trend estimation tool.

Q class	Hi, Mod, Lo, doubtful. These classes try to represent the probability that the train came from a true train source and was not a chance association of similar clicks from different sources. They do not map from C-PODs to F-PODs as in the latter doubtful trains are classified less often and trains identified as trains of echoes are put in this class.
refraction	The bending of sound by unevenness of the medium it is travelling through. Sea water is never uniform in salinity, temperature (and movement and pressure) so a sound does not travel in a perfectly straight line without distortion, but can diverge, converge and interfere constructively or destructively with itself.
replicates	The same click arriving later than the earliest version having travelled along a different path as a result of reflection and/or refraction.
rms	Root mean square. The intensity of a sound is represented by the rms of the successive amplitude values that make up its waveform.
Sediment transport	Seabed material brought into suspension by currents. It generates sediment transport noise. Particles around 0.25mm diameter generate sound around 130kHz. All sandy seabeds move sometimes. If they did not, they would be overlain by mud.
social calls	NBHF species probably all use distinctive patterns of ICI to create social calls. Other species, which includes most dolphins, use whistles as social calls, and also use bursts of clicks at a very high rate.
Sonar filter	See text
Sonar risk	The estimated probability of the presence of a boat sonar in some data, typically 1 minute of data. The detection of a sonar risk may bias the detection of cetacean.
Sound pressure level	Here this is used to refer to instantaneous pressures, e.g. of the peak of a wave or click.
Species class	The primary species classes given by the KERNO classifiers are NBHF, other cet (= BBT = other cetaceans), boat sonars and unclassified.
Spectrum	The distribution of energy across the frequencies in a sound. The FFT
SPL	See Sound pressure level.
Threshold	When the F-POD is running it responds to a series of minutes with high click counts by raising its detection threshold. This is not a simple threshold but is applied first to short clicks and only later to long clicks. The benefit is that large number of weak short clicks are of little value in detecting dolphins, but weak long clicks are more often useful in detecting NBHF species. This is sometimes referred to as the operational deadband setting. It has several steps. The starting point and criteria for changing up or down are under user control, but most changes will be worse than the default settings!
TIC – time interval clustering	A measure of how clustered detection are in C-POD data.
Time	Times are stored in C-POD files as the number of minutes since the start of the year 1900. This is the same start as is used in Access and Excel and many other apps. Microsoft Office stores the number of days, plus the fraction of a day. To create a date/time value in 'Office' format from a PODmins value: PODmins/1440 The POD stores times within the minute in 5 microsecond units (max value = 12million) that are reported in data readouts as their microseconds value. This can be added to the Office-style date/time value after dividing it by the number of a microseconds in a day (1440*60,000,000).

Time domain	The entire set of sequential values of sound pressure in a record, and the various descriptors that can be derived from these values (i.e. without using the FFT or similar to infer what frequencies are present). The time domain gives useful features such as the wavenumber of the loudest wave in a click, the relationship between the wavelengths and amplitudes of numbered waves, the mean wavelength. Some time domain features may represent the physics of the sound production mechanism e.g. the propagation of vibrations through the physical structure.
Time endpoints	Time points set, most easily, by right-clicking on the display of a file and selecting 'set selection start' or 'set selection end'.
Time selection	A part of a file defined by time endpoints, one of which may be the start or end of the file. They enable you to select part of file and analyse it or make a copy of it. The easiest way to set one is to use the right-click menu on the main display area to set a time selection start or end at the position of the mouse pointer. You can save this in the file header using the 'save selection end times' label, bottom left, on the Files page. If you only set a start or only an end, then the other end will be the start or end of the file.
tonal	Having a frequency peak, which may be broad or narrow. White noise, pink noise etc. show no strong peak. More tonal clicks have more cycles of more similar wavelengths and vice-versa.
train	A more or less regularly spaced series of more or less similar elements. Cetacean click trains show time spacing changes that sometimes exceed a doubling or halving of ICI. The click features also change through the train, especially in the case of dolphin clicks which show strong variation across the sound beam. The POD captures only a fragment of the train produced by the animal as its sound beam sweeps across the POD hydrophone.
train coherence	A measurement of the similarity of successive ICIs and successive click features. Trains logged often represent a series of points across the beam of sound produced by the animal, and the clicks show a stronger resemblance to each other than to any generalization of the features of the clicks of the species, or to the 'on axis click'.
trend	The trend analysis tool generates values for the trend at a lite or set of sites for whatever parameter – such as DPM – is selected. See PYRA
Use end times saved in each file	The auto-crop utility can set and save endpoints in each file to the first and last midnight when the POD is running, and this will use those (or any other saved endpoints, but only one pair can be saved) Only useful if PODs were set to be OFF before immersion.
validation	Any method of testing whether an identified train was what it was classified as being. This is a crucial process and often conducted in a highly subjective way. It is amenable to some verification, but much more work is needed to put this on firm basis.
VEMCO fish tag	Small data loggers that can be attached to fish and that repeatedly transmit sounds at 70 kHz that encode the date collected. PODs log these transmissions.
Warnings	Warnings of possible sources of error in files are generated at the end of running KERNO-F, or by running the process on FP3 files via a button on the Trains page – that will update them to the latest version. They are seriously useful and can be viewed for any open FP3 file or exported for a batch via the Filters +files page.
Wave	See cycle

Wave number	The waves within a click are numbered from 1.
Weak unknown train sources	= WUTS. These are trains of sounds that can resemble BBT or NBHF click trains, but do not come from cetaceans. Sometimes they have a very distinctive character, starting with very short ICIs that increase steadily. Their origin is unknown but is probably biological with more than one species producing these sounds. In some places they are a serious problem for automated classifiers.
whistles	Long tonal sweeps made by most dolphin species as communication calls.
WUTS	See <i>Weak Unknown Train Sources</i> . These trains are, at present, only seen in POD data because they are too hard to find without train detection. Their source is almost certainly a number of animal species that are at present unidentified. They are not shrimps which do not make trains of clicks, but stridulation by a crustacean is possible. Other possibilities include mollusc radulas scraping the hydrophone surface, and polychaete worms moving through sand.



We welcome all feedback, especially criticisms and suggestions for improvement. Please do not hesitate to contact nick.trogenza@chelonia.co.uk Thank you!