

The FPOD app - Version history and Notes

2/03/2025

To avoid unintelligible long version numbers the FPOD app has version numbers for 7 distinct functional areas:

1. Changes to the [KERNO-F version](#). These affect only the contents of the FP3 files produced by the KERNO-F classifier. The current version, since 2022, is V1.0 and this will not be updated until major improvements are possible, and version 1.0 onwards will remain available to make consistent analysis easy.
2. Changes to the [Export version](#) may also affect data exported. Some changes will affect data exported but none affect the data in the FP1 or FP3 file, so re-processing of files is not required to bring updates into effect.
3. Changes to the [Display version](#) affect only what you see on screen. There will be many changes intended to improve the display, ease of use, various bugs etc. Those do not affect the numbers you export and there is no need to re-process FP1 files or keep different versions of exported data as they will not be different.

Changes to the filtering defaults do affect what is exported and are in red below.

Numerous trivial changes are not documented...

Small and large versions have been introduced to fit better on some screens, 27/09/2022
4. [Warnings version 1](#) gave far too many warnings and was consequently difficult to use. Version 2 - the current version - can be quickly generated from FP3 files without re-processing of the FP1 files, and without changing any other features of the FP3 data.
5. [Technical version](#). Where users encounter problems in specific contexts (e.g. SD card corruption) code changes that are required to overcome them are incorporated in the next released version.
6. [Settings version](#). These are only the setting that go on the POD via an F_SETS.txt file. Various pre-defined settings can be applied to an FPOD and are accessed via the Settings page of the menu.
7. [Social Call version](#). This is the beginning of a development process, starting with smoothing train inter-click-intervals. Click rate slope values are now accessible via the 'Dev' export item.
8. [FP3 version](#). The FP3 file contains only those clicks that are found to form part of a train by the KERNO-F classifier. Originally (FP3version 0) each click had two 16byte records in the FP3 file, but only one in the FP1 file. FP3version 1 has three 16byte records for each click and the version number will indicate what data is in this extra record which does make version 1 files 50% bigger than FP3version zero files.

Here is a record of the version changes in each function, with current version at top of each list:

1. [KERNO-F version](#)

version	change	date
1.00	There is, for research purposes, now an option to block the feedback stage that runs after KERNO-F and can modify the results in any minute based on the results in the 5 minutes before and after it..	02/03/2025
1.00	Very large sonar clusters could crash KERNO-F. This limitation has been removed. It does not change the results obtained from any data file, so the version number is still 1.0.	27/09/2022
1.00	No changes Data processed with earlier version should be re-processed.	23/03/2022

2. [Export version](#)

3.2	Third Party support has arrived! Powerful exporting and re-importing versions are now accessible via a button, bottom right, on the Export page. This is to enable machine learning / AI developments. The distribution of ICI gradients across click rates can now be rapidly exported via the Export page. Files must be updated first - see Social Call version below.	26/10/2024
3.01	Text headers for species ID in various exports	13/09/2024
3.0	Includes export of smoothed ICI values, and file paths	05/08/2024
2.25	Inter-click-interval export now puts all ICIs in each train on its own row	20/04/2024
2.24	Exporting train ICI inflections has been revised. The earlier version did a form of autocorrelation and may have been more useful ... do let us know if you have used either	28/02/2024
2.23	Missing export columns fixed	25/01/2024

2.22	'Nall' dropped from exports - it's too difficult to use accurately, but can still be obtained by exporting N of clicks from the FP1 file, and that allows use of relevant filters. MinuteModallCIs added - this is a tool for exploring aspects of social communication. TrainCInflexions is a tool for exploring differences in click rate modulation behaviour between species.	30/12/2023
2.21	Sonar filter triggers can now be exported	10/12/2023
2.20	The warnings from version 2 are much more accurate and useful. On the Trains page you can quickly update the warnings in all your existing FP3 files using the 'search directories' option. This gives a listing of the risks found.	04/11/2023
2.14	Time Lost export in Detections & Environment now works on FP3 files, but they must be reprocessed to show this value. Click detail export: extrapolated amplitudes added.	15/09/2023
2.13	The trend analysis tool which implements PYRA - Paired Year Ratio Assessment - for the determination of trends is now working, and accessed via the Analysis page. TimeLost reporting in Detections and Environment fixed. Files need to be reprocessed through train detection to show these results. Errors in export of later F-POD angles fixed (0-180 in place of 0-255)	06/08/2023
2.06	'Detections and Environment' was exporting click rates and now exports click counts	17/03/2023
2.05	Exporting raw data on trains now includes species numbers allocated by the user Header N of minutes in incomplete periods was 1 less than in (correct) exported data	10/03/2023
2.04	Detections and Environment now exports average operational threshold values	23/01/23
2.03	Fix for: Analysis page results were slightly excessive numbers since technical version 1.04 released 4/6/22. Exported results were not affected.	8/1/2023
2.02	Train and Click details: to handle very large data sets a sample can now be exported from a fraction of minutes. Time export options improved.	1/1/2023
2.01	Site codes can now be 7 characters long	22/12/2022
2.00	'Detections and Environment' now exports separate columns of click counts and DPM for NBHF, dolphins and sonars, as .csv or tab separated text. Enjoy!	13/11/2022
1.14	Noise level export removed	11/11/2022
1.13	Site codes added to reduce volume of large data exports etc	08/11/2022
1.12	.csv file export added as an option	06/11/2022
1.11	Train details export to text file now puts a description of the data columns in the Results	31/07/2022
1.10	Improved reporting of Time lost in Detections and Environment export	07/07/2022
1.09	Adds clipped file names option	01/07/2022
1.08	Operational setting export now only exports from file type selected	27/06/2022
1.07	Encounter details for third party species classifiers implemented. N of clicks can now be exported from FP1 files.	14/06/2022
1.06	Train ICI inflexions distribution is now exportable.	04/06/2022
1.05	Angles and temperatures from F-PODs corrected. Some very early hardware versions may require division by 5.	20/05/2022
1.04	Saving export settings omitted some (little used) filters	08/05/2022
1.03	Train details export missing some files fixed	25/04/2022
1.02	Click counts header row now shows file type	08/04/2022
1.01	Exporting lists of file settings etc now includes the KERNO-F version used	25/03/2022

3. Display version

2.14	Fixed CP3 data was display blank for species and train quality values, and filtering by avNcyc (not available from CP3 files) excluded all trains.	12/02/2025
2.13	Click rate displays are now colour coded by the train rate confidence, red=high, gray = low	17/12/2024
2.12	Marking and changing the species class of trains can now apply to all those on show in any time resolution	10/11/2024
2.11	The PRF = Clicks/sec display now colour codes the rate graph by the slope, so rising click rates are green, falling are red and static is grey.	26/09/2024
2.10	FP3 files are upgraded to version 1 which includes values of smoothed inter-click-intervals	05/08/2024
2.08	Sitecodes are now 11 characters long. Pervious sites codes will no longer read. Apologies if you have set them...	25/01/2024
2.03	Adds audio representation of click rates . Various minor fixes	25/10/2023
2.0	The app is now rescaled to the screen it starts on.	29/06/2023
1.11	Graphing of IPIs within click restored. C/FP1 redraw suppression improved	05/04/2023
1.10	Diel graph bugs: scaling of PRF and click kHz, diel display from 6 files	19/02/2023
1.09	CPOD filtering by rate quality of trains fixed	22/12/2022
1.08	Battery stack numbering was offset by 1 for some PIC code versions	23/08/2022
1.07	Analysis was blank on some older C-POD files corrected.	02/08/2022
1.06	Angles and temperatures from F-PODs and NBHFi low res display corrected.	20/05/2022
1.05	Now fits screen with 1080 pixels vertically	02/05/2022
1.04	Many cosmetic changes	30/04/2022

1.03	File parameter displayed bug fixed.	24/04/2022
1.02	The default filter setting now includes no WUTS risk filtering. Previously this filtering was set at value 1 and was flagged up in the list of filters shown in various places. It typically 'costs' about 1 - 3% of detections. There is no change to the KERNO-F classifier, only the filtering of FP3 files. You do not need to take any action unless you exported data with earlier versions and want to compare it precisely with this or later versions.	12/04/2022
1.01	Error reading unexpected floating point marker from a text box corrected.	29/03/2022

4. Warnings version

2.00	The warnings from version 2 are much more accurate and useful. On the Trains page you can quickly update the warnings in all your existing FP3 files using the 'search directories' option. This gives a listing of the risks found. It has used information from a recent study of F-POD data published in PLOS ONE, and includes a 'sum of risks' value that is very useful. Values from any set of files can then be listed via the Filters+files page.	04/11/2023
1.03	Warnings were not written into the FP3 file in some cases. To correct this requires re-analysing the data through the KERNO_F classifier. That has not changed so the results will not be different (apart from the warnings).	07/07/2022
1.02	Target and actual modal NBHF kHz now reported in export of classification warnings	18/04/2022
1.01	Non-standard settings for KERNO-F v1 now reported in export of classification warnings	09/04/2022
1.00	No changes	

5. Technical version

2.6	1. Adds a 'redaction' utility that deletes the data from sections of files that are defined in a specific format in text lines pasted into the Results memo. 2. The settings used when running the KERNO-F classifier are now logged in the FP3 file and can be exported 3. Fixed: an error that allowed the marking or re-classifying of trains in Hi Res screens to go beyond the screen end within the last minute on display.	01/03/2025
2.5	Handles readings SD card in some cases where early files in a series have been lost.	17/05/2025
2.4	FP3 version 1 introduced. Each click record has an additional 16byte record, so these are about 50% bigger than version 0. It will hold information on social features of clicks and trains. FP3 files are updated when opened if 'use and show smoothed ICIs' on the Display page is checked.	05/08/2024
2.2	It is now possible to exclude 1 minute in 2,3..10 in running the KERNO classifier. This is provided to allow comparison between data logged with differing duty cycles. The feedback process has been adjusted for alternate minutes logged. The effect of longer duty cycles has not been evaluated. There is no effect on KERNO-F output from continuous data.	02/02/2024
2.12	The landmark sequence search now uses your own filters and exports initial and final kHz	06/12/2023
2.11	The SkipToScreen function now skips to a screen that always has the required number of clicks, and displays a short overlap with the previous screen to avoid missing features.	02/15/2023
2.10	F-PODs with acoustic release functionality in the FPGA now generate reports to enable Chelonia to optimise the release codes. These are text files produced when the SD card is read. Please email the text file to team@chelonia.co.uk	26/11/2023
2.00	The trend analysis tool is now working - gives trends on multi-year data sets	19/11/2023
1.10	Search directories for exporting now allows selection of file types	04/11/2023
1.09	Species numbers 1..6 can be added to trains	10/03/2023
1.07	Adds functionality for identifying.wav file source for virtual F-POD files generated by WavToPODfile.exe	20/12/2022
1.06	Added: Spaces in file names can be replaced by gaps	16/09/2022
1.05	Fixed: On some PCs FP1 files were truncated during train detection.	09/06/2022
1.04	Corrected: the first minute of files after file0 in a series was given the same time as the last minute of the previous file. Faster display of low resolution data by omitting calculation of non-displaying parameters.	04/06/2022
1.03	Affects C-PODs only: Reporting of differences between actual and registered gain is no longer automatically reported but this utility is still active in CPOD.exe	07/04/2022
1.02	2 Minor abnormalities in raw data affecting less than 1 click record per million fixed.	

6. Settings version

2.02	Adds an option for being asleep in any selection of calendar months - useful for polar deployments	20/04/2024
2.01	Adds an option for running 1 minute ON in 3minutes (1on, 2sleeping). Lower switch-over levels now set for Lithium batteries.	13/01/2024
2.00	All settings for use in the field now have a lower value for 'quiet minutes' to reduce hunting. This will have almost no effect on detections. The severe noise settings now exclude many clicks below 20kHz. This might reduce very distant dolphin detections and slightly increase all dolphin detections.	19/11/2023

1.08	A River Dolphin version has been added. This handles the issue of nearly continuous noise from transport of fine sand.	18/08/2023
1.07	Intermittent logging fixed. Text for high pass filter in F_SETS corrected	17/08/2022
1.06	Early F-PODs did not start with some previous versions	28/06/2022
1.05	Lithium battery switch-over level reduced to 15.5v from 17.5. Adds 3weeks running time.	17/06/2022
1.04	By default the real-time sonar filters are now OFF to avoid biasing dolphin click spectra.	04/06/2022
1.03	A redundant line has been removed to avoid a 'parsing error' report in Report.txt. You do not need to take any action.	16/04/2022
1.02	LF-POD settings for battery level at shut-down added. If you use LF-PODs with settings from this version you get slightly longer running times.	06/04/2022
1.01	Angle settings for switching on and off corrected. If you want to use this function you should use settings from this version or later.	05/04/2022
1.00	The default settings have changed from early FPODs. The change avoids some unstable switching between the logging thresholds that was seen in some quiet environments, and it reduces data volumes. This is valuable for all PODs, so you should change settings to this version when you can, but if you don't your data will still be good, and comparable.	23/03/2022

7. Social Call version

version	change	date
0.4	An update of FP3 version 1 files, using the button on the Trains page, now adds ICI slope values to each click after the 2 nd in a train. The distribution of slope values varies between species and activities. The slope values are $2000 * \text{Log}_{10}(\text{ICI} / \text{PriorICI})$	26/10/2024
0.3	FP3 files now have a version 1 that is generated automatically if 'use smoothed ICIs' on the Display page is ticked. You can then view and export smoothed values for ICIs that remove some noise from the ICI sequence of each train.	04/08/2024
0	This is a simple exploratory version that looks for minutes with a strong predominance of short ICIs in a narrow band with few longer ICIs. Minutes are given an annotation code of 175 and can be filtered using this.	31/12/2023

8. FP3 version

version	change	date
1	Three 16byte records for each click. A smoothed value for the inter-click interval preceding each clicks, except the first, in each train is stored in two bytes and is displayed when ICI or PRF (inter-click interval or click (=pulse) repetition frequency = clicks/s) is selected. To get the actual ICI export the raw click data via the right-click pop-up menu.	Not yet released
0	The original, standard version with two 16byte records for each click	

Some Notes on Filters:

There are filters and settings that affect **only the operation of the KERNO-F classifier**. They change the content of the FP3 file it produces, and do nothing later:

These three are worth changing in a few places, and the change is recorded in the file so you can see that fact in the Classification warnings subsequently if you are not sure whether you remembered to set it

No dolphins No NBHF No sonars

This one is also recorded in the file in the same way:

NBHF definition settings: defaults
 NBHF clicks, 'target' kHz 120

All these are not recorded in the file and changing the advanced settings is generally a bad idea, because a huge amount of testing will be needed + the appropriate data sets to test on.

NBHF lower range	18
NBHF upper range	15
NBHF clicks, peak at cycle number:	4
NBHF clicks, 'target' N of cycles	8

Advanced KERNO-F settings		Limit clicks/m to 0
Duration N of cycles >	0	<input type="checkbox"/> block feedback
Amplitude >	10	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

<< this is tiger country.
 Enter at your peril.
 Help is not on hand...
 But if you do have a cetacean clicking at <25kHz you can change that one...

There are filters that *operate only on the completed files*, FP1 or FP3, and do not change the file content:

The filters in force are shown in red in the box at the bottom of the screen e.g.

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

They are set here:

Train filters:
 marked trains included more
 high sp confidence only
 Quality 'Q': KERNO/F
 Hi Lo allQ
 Mod ?/echo
 Species: KERNO/F classifier
 NBHF Other cet
 Sonar unclassified

Display versions before 1.02 also had **WUTSrisk: 0-1** The default has been changed to no WUTS filtering because few files have any WUTS but it did cost a few percent of click trains. This is set on the Filters +files page. Apologies for not getting to that configuration at the start!

How does KERNO-F work?

KERNO-F uses the old-fashioned 'feature engineering' approach in a layered process:

1. Analyse raw data to give measures of the frequency, bandwidth and amplitude of clicks, and their resemblance to typical NBHF (porpoise) clicks.
2. Identify the exponentially decaying multipath clusters that are typical of cetacean clicks.
3. Identify features of sediment transport noise or other continuous noise sources in whole minutes of data.
4. Seek sequences of similar clicks at similar intervals (trains), allowing the click character and intervals to evolve over the course of the train. This uses multiple hypothesis testing.
5. Remove one from every pair of possible trains found that share one or more clicks.
6. Derive features for distinguishing trains that come from one of the sources listed above from those that arise by chance in background noise. KERNO-F uses 80+ features derived from the raw data in a process that is heavily biased against false positives.
7. Derive features for distinguishing 'species classes' or guilds of trains - 'NBHF', 'Other cetaceans', 'Boat sonars'. 70+ features are used.
8. Calculate classification distances between the guilds. These distances are vectors and are not reversible.
9. Use these distance measures for individual trains in combination with whole minute metrics to classify each train, if that can be done with adequate confidence (i.e. with a bias against false positives), or ascribe them to 'unclassified species'.
10. Populate a record of the detection metrics of each minute.
11. Write these trains into an FP3 file.
12. Use an 11-minute-long window of these detection metrics as the basis for feedback in which each species group may be boosted by positive feedback (i.e. unclassified trains are re-classified) or negative feedback.
13. Write these changes into the existing FP3 file.
14. Register in the file various potential sources of error in the whole file that may merit inspection:

- 0 Validate OtherCet - low total number
- 1 Validate NBHF - low total number
- 2 Validate OtherCet as they are <1% NBHF
- 3 Validate NBHF as they are < 1% OtherCet
- 4 Boat sonars more numerous than NBHF
- 5 Boat sonars more numerous than Other cet
- 6 Validate NBHF as modal kHz is more than 2kHz below target value
- 7 Validate NBHF as modal kHz is more than 4kHz above target value
- 8 Validate NBHF as modal NofCycles is abnormal
- 9 Validate OtherCet as modal click-rate is abnormal
- 10 Validate NBHF as modal click-rate is abnormal
- 11 Validate OtherCet as modal kHz is abnormal
- 12 Inspect source of abnormal modal kHz in this file
- 13 POD more than 30degrees from vertical in more than 5% of minutes
- 14 Inspect source of noise seen in >5% of minutes
- 15 High fraction of minutes with elevated WUTS risk
- 16 No warnings for this file

15. Finalise the FP3 file and deliver a numerical report on the whole process, which includes the following measures:

NBHF clicks Hi+ModQ (Hi+ModQ is the selection of classification confidence that is generally used)
N promoted NBHF via feedback
OtherCet clicks Hi+ModQ
N promoted OtherCet via feedback
Sonar clicks Hi+ModQ
Minutes with continuous noise source such as sediment transport noise
Minutes with boat sonars
'WUTs ratio' a measure of how many minutes have features suggestive of weak unknown train sources
Target kHz used to identify NBHF click trains
Median kHz of the median kHz in trains identified as NBHF

Frequently Asked Questions

Why are some good trains not classified?

This is often seen and arises from two factors: the trains, in their context, have some resemblance to false trains in some other context, and both would become false positives if the criteria were weaker; or, the algorithm could be better! So it 'doesn't make sense' in the context of the file it is in, but does make sense in the context of using the same process to analyse a very diverse set of acoustic data.

Why is it so slow?

Actually it's very fast! Typically it processes 10million clicks per minute on a modern PC. In 10k clicks - often found in 1 minute - there are more different trains of 5 or more clicks than seconds since the origin of the universe ... and there are a lot of minutes. The F-POD app is trying to find all the good trains. All known solutions to this problem are imperfect. The code is written in Delphi, which is one of the very fast languages, like C. In speed these are about 25x Python and >100x MATLAB or R.

Third party access:

The species identification could be improved by the use of machine learning methods.

The FPOD app is structured to provide the data exports you could use and the means to put your classification back into the data file so that the viewing, analysing and exporting tools can be used on it.

So if you are interested in doing that please get in touch.

All feedback, especially negative, is welcome!

Contact: nick.tregenza@chelonia.co.uk