NAME

mbfilter – Apply some simple filter functions to sidescan or beam amplitude data from swath sonar data files.

VERSION

Version 5.0

SYNOPSIS

mbfilter [**–A***kind* **–B***yr/mo/da/hr/mn/sc* **–C***mode/xdim/ldim/iteration* **–D***mode/xdim/ldim/iteration[/offset]* **–E***yr/mo/da/hr/mn/sc* **–F***format* **–I***infile* **–N***buffersize* **–R***west/east/south/north* **–S***mode/xdim/ldim/iteration[/threshold_lo/threshold_hi]* **–V –H]**

DESCRIPTION

mbfilter applies one or more simple filters to the specified sidescan or beam amplitude data. The filters include:
–S1: boxcar mean for low-pass filtering
–S2: gaussian mean for low-pass filtering
–S3: boxcar median for low-pass filtering
–S4: inverse gradient for low-pass filtering
–D1: boxcar mean subtraction for high-pass filtering
–D2: gaussian mean subtraction for high-pass filtering
–D3: boxcar median subtraction for high-pass filtering
–C1: edge detection for contrast enhancement
–C2: gradient magnitude subtraction for contrast enhancement

These filters are mostly intended for use with sidescan data. In particular, the low-pass or smoothing filters can be used for first-order speckle reduction in sidescan data, and the high-pass filters can be used to emphasize fine scale structure in the data. A combination of low-pass and high-pass filtering can effectively perform band-pass filtering. The contrast enhancing filters can, under certain circumstances, sharpen sidescan images of the seafloor. The low-pass and contrast enhancement filters are described in the paper by Sauter and Parson (1994) listed below.

The filtering used here is designed and applied in the same manner as spatial filters are applied to images. The swath data is treated as an image, with the x and y coordinates corresponding to pixel (or beam) number and ping number, respectively. The filters consist of M x N matrices which are convolved with the swath data image (M is the filter dimension in the acrosstrack direction and N is the filter dimension in the alongtrack direction). Missing or flagged data are ignored in the filtering process.

Low-pass filters are specified with the **–S** option, high-pass filters with the **–D** option, and contrast enhancement with the **–C** option. Users can apply up to ten filters in any order; the filters are applied in the order in which they are specified on the command line. Users can apply as many iterations of each filter as desired. The user specifies the window size used by the filters (3x3 or 5x5 are typical for using smoothing filters iteratively to reduce speckle; larger windows like 3x20 are typical for high-pass filtering).

The boxcar mean, gaussian mean, and median filter all achieve similar results in uniformly smoothing the data; the boxcar mean is faster to apply, the gaussian mean preserves the frequency content of the data better, and the median filter is least sensitive to spikes in the data. The inverse gradient filter applies averaging weights which depend on the inverse gradient of the data. This approach causes the filter to smooth regions without distinct edges much more than regions with edges, thus tending to preserve the sharpness of features more than the simpler smoothing filters.

The median smoothing filter can be set by the **–T** option to operate with low and high ratio thresholds (the value is changed only if the original value divided by the median value is less than the low threshold or greater than the high threshold). This allows the filter to preferentially despike the data. In particular, this approach is useful for suppressing "stripes" or "bad" pings which have amplitude or sidescan values differing significantly from surrounding pings.

The high-pass filters are constructed by calculating a low-pass filtered version of the data and then subtracting that from the original data. An offset value is added to the high-passed data so that it is positive (negative values are considered flagged as bad in some formats and not allowed in others). The high-pass filters can be used to remove long-wavelength variations in seafloor reflectivity in order to emphasize fine-scale structure.

The contrast enhancement filters are generally only successful when applied after smoothing because of their tendency to amplify noise. The edge detection filter enhances contrast by tending to shift values on either side of a boundary away from the average value across the boundary. A 5 X 5 filter or larger is generally required for success with the edge detection algorithm. The gradient filter increases contrast by subtracting twice the local gradient magnitude from each value.

The filtered amplitude or sidescan data are written to ancilliary files located parallel to the input swath data files. Filtered amplitude data will be in files named by adding ".ffa" to the input file name, and sidescan data will be in files with an ".ffs" suffix. Filtered amplitude or sidescan data can be plotted with **mbm_plot** by appending "F" to that macro's mode argument (e.g. **−G***4F* for filtered amplitude data, and **−G***5F* for filtered sidescan data). If accessing **mbswath** directly rather than through **mbm_plot**, plot the filtered data by appending "F" to the program's mode argument (e.g. **−Z***4F* for filtered amplitude data, and **−Z***5F* for filtered sidescan data). When using **mbmosaic** to generate mosaics of amplitude or sidescan data, the filtered data can be accessed by appending "F" to the data kind argument (e.g. **−A***3F* for amplitude and **−A***4F* for sidescan data.

AUTHORSHIP

David W. Caress (caress@mbari.org)
Monterey Bay Aquarium Research Institute
Dale N. Chayes (dale@ldeo.columbia.edu)
Lamont-Doherty Earth Observatory

OPTIONS

−A *kind*
Determines whether beam amplitude (*kind* = 1), or sidescan (*kind* = 2) data will be processed. Default: *kind* = 2.

−B *yr/mo/da/hr/mn/sc*
This option sets the starting time for data allowed in the input data. The **−E** option sets the ending time for data. If the starting time is before the ending time, then any data with a time stamp before the starting time or after the ending time is ignored. If instead the starting time is after the ending time, then any data between the ending and starting time will be ignored. This scheme allows time windowing both inside and outside a specified interval. Default: *yr/mo/da/hr/mn/sc* = 1962/2/21/10/30/0.

−C *mode/xdim/ldim/iteration*
Turns on contrast enhancement filtering and sets the filter parameters to be used. Here *mode* specifies the filter type:
mode = 1 : Edge Detection Filter
mode = 2 : Gradient Magnitude Subtraction Filter
The dimensions of the data window used are set using *xdim* (acrosstrack dimension) and *ldim* (alongtrack dimension). Values of *xdim* = 5 and *ldim* = 5 are typical, but larger or smaller dimensions can be used. The *iteration* value specifies the number of times the filter is applied; there is no limit to this value. Default: contrast enhancement off, *xdim* = 5, *ldim* = 5, *iteration* = 1.

–D *mode/xdim/ldim/iteration/offset*
Turns on high-pass filtering and sets the filter parameters to be used. Here *mode* specifies the filter type:

mode = 1 : Boxcar Mean Subtraction Filter
mode = 2 : Gaussian Mean Subtraction Filter
mode = 3 : Boxcar Median Subtraction Filter

The dimensions of the data window used are set using *xdim* (acrosstrack dimension) and *ldim* (alongtrack dimension). Values of *xdim* = 3 and *ldim* = 3 are typical, but larger dimensions can be used. The *iteration* value specifies the number of times the filter is applied; there is no limit to this value, but high-pass filters are generally only applied once. The *offset* value is added to each high-pass filtered value to force the results to be positive; this value should be chosen according to the range of values allowed in the data type and data formats being used (e.g. for SeaBeam 2100 sidescan as represented in format 41, the sidescan values can range from 1 to 65535, so an *offset* = 1000 is appropriate). Default: high-pass filtering off, *xdim* = 3, *ldim* = 10, *iteration* = 1, *offset* = 1000.

–E *yr/mo/da/hr/mn/sc*
This option sets the ending time for data allowed in the input data. The **–B** option sets the starting time for data. If the starting time is before the ending time, then any data with a time stamp before the starting time or after the ending time is ignored. If instead the starting time is after the ending time, then any data between the ending and starting time will be ignored. This scheme allows time windowing both inside and outside a specified interval. Default: *yr/mo/da/hr/mn/sc* = 2062/2/21/10/30/0.

–F *format*
Sets the MBIO integer format identifier for the input file specified with the –I option. By default, mbfilter derives the format id from the mbpro- cess parameter file associated with the input file (-I option) or, if nec- essary, infers the format from the "*.mbXX" MB-System suffix convention.

–H This "help" flag causes the program to print out a description of its operation and then exit immediately.

–I *infile*
Swath data file from which the input data will be read, or a datalist file containing a list of input swath data files and/or other datalist files. If infile is a datalist file, then mbprocess will attempt to process all data files identified by recursively reading infile. Default: *infile* = "datalist.mb-1"

–N *buffersize*
Sets the maximum number of data records which can be read into the buffer. In general, data records may be of several different types (e.g. parameter, position, comment) in addition to survey data records. Many data formats include many more position data records than survey data records. Thus, a large buffer may be required to access a reasonable number of survey data records. However, on memory limited machines large buffer sizes can lead to poor performance due to memory swapping. The default value of *buffersize* = 500 is appropriate for most cases, but users can set the buffer size as required. The absolute maximum buffer size is 5000. Default: *buffersize* = 500.

–R *west/east/south/north*
Sets the longitude and latitude bounds within which swath sonar data will be read. Only the data which lies within these bounds will be copied. Default: *west*=–360, east=*360*, *south*=–90, *north*=90.

–S *mode/xdim/ldim/iteration*
Turns on low-pass smoothing filtering and sets the filter parameters to be used to smooth the data. Here *mode* specifies the filter type:

mode = 1 : Boxcar Mean Filter
mode = 2 : Gaussian Mean Filter
mode = 3 : Boxcar Median Filter
mode = 4 : Boxcar Inverse Gradient Filter

The dimensions of the data window used are set using *xdim* (acrosstrack dimension) and *ldim* (alongtrack dimension). Values of *xdim* = 3 and *ldim* = 3 are typical, but larger dimensions can be used. The *iteration* value specifies the number of times the filter is applied; there is no limit to this value. Default: *mode* = 1, *xdim* = 3, *ldim* = 3, *iteration* = 1.

−T *threshold_lo/threshold_hi*
This option causes the boxcar median smoothing filter to operate with low and high ratio thresholds (the value is changed only if the original value divided by the median value is less than *threshold_lo* or greater than *threshold_hi*). This allows the filter to preferentially despike the data. In particular, this approach is useful for suppressing "stripes" or "bad" pings which have amplitude or sidescan values differing significantly from surrounding pings. This option only works with the median smoothing filter.

−V Normally, **mbfilter** works "silently" without outputting anything to the stderr stream. If the **−V** flag is given, then **mbfilter** works in a "verbose" mode and outputs the program version being used, the values of some important control parameters, and all error status messages.

EXAMPLES

Suppose one has a SeaBeam 2100 data file called test.mb41 which contains bathymetry (121 beams in a 120 degree swath), beam amplitude (121 beams coincident with bathymetry), and sidescan data (2000 pixels, roughly a 150 degree swath) which has been corrected for the amplitude vs grazing angle variation using the program **mbanglecorrect**. Plots of the corrected sidescan often show a large amount of speckle that was suppressed in plots of the raw data by the large contrast between the specular and non-specular regions of the swath. To reduce the speckle and make coherent features of the data clearer, the user can apply any of the smoothing filters available in **mbfilter**. One iteration of the boxcar mean filter can be applied as follows:

    mbfilter −F41 −Itest.mb41 −Otest_mean.mb41 -S1/3/3/1

Five iterations of the inverse gradient filter can be applied as follows:

    mbfilter −F41 −Itest.mb41 −Otest_igrad.mb41 -S4/3/3/5

To remove large scale variations in seafloor reflectivity, one can apply a high-pass filter to the data:

    mbfilter −F41 −Itest.mb41 −Otest_hipass.mb41 -D1/3/10/1/10000

To first apply a high-pass filter to emphasize fine-scale structure and then apply a low-pass filter to reduce speckle:

    mbfilter −F41 −Itest.mb41 −Otest_hipass.mb41 -D1/3/10/1/10000 −S4/3/3/5

To first reduce speckle by smoothing the data with Gaussian mean filter and apply an edge detection contrast enhancement filter:

    mbfilter −F41 −Itest.mb41 −Otest_hipass.mb41 -S3/7/7/1 −C1/5/5/1

SEE ALSO

mbsystem(1), **mbmosaic**(1), **mbm_plot**(1), **mbbackangle**(1), **mbanglecorrect**(1)

REFERENCES

Sauter, D., and L. Parson, Spatial filtering for speckle reduction, contrast enchancement, and texture analysis of GLORIA images, *IEEE J. Ocean. Eng.*, **19**, 563-576, 1994.

BUGS

All the filtering in the world won’t make bad data good.