FINDING FLIERS: NEW TECHNIQUES AND METRICS

Matthew J. Wilson[1,*], Giuseppe Masetti[2,‡], and Brian R. Calder[3,‡]

[1]Tel: +1-(757)-441-6746, email: matthew.wilson@noaa.gov; [2]email: gmasetti@ccom.unh.edu; [3]email: brc@ccom.unh.edu

[*] NOAA Office of Coast Survey, Atlantic Hydrographic Branch, Norfolk, VA, USA

[‡] Center for Coastal and Ocean Mapping & NOAA/UNH Joint Hydrographic Center, University of New Hampshire, Durham, NH, USA

Abstract

Fliers in bathymetric grid products evidence an inaccurate portrayal of the seafloor, and must be remediated before final product creation. Traditional methods to ensure flier detection and removal are apparently not sufficient since nearly 25% of surveys received at NOAA in 2015 had final grid deliverables affected by fliers. These grids not only require time and effort to resolve, but also compromise the lineage of the original field submission. This work describes new automated techniques for flier identification and compares their effectiveness to traditional, manual methods of flier detection. The benefits of automatic identification being statistically significant, integration of the flier identification techniques could be in the gridding algorithm is suggested as a means to reduce the rate of fliers delivered on the created surfaces.

Introduction

Improved accuracy of hydrographic survey data and faster throughput from acquisition to chart application are priorities of the NOAA Office of Coast Survey (NOAA, 2015) but are common issues for any hydrographic office (HO). Challenges to these objectives exist during the review of hydrographic survey data to ensure accuracy of the final products, and of these, the removal of anomalous grid data, commonly known as fliers, is among the most significant. Fliers are particularly problematic to the objectives of accuracy and timeliness because by definition they represent inaccurate data in the final survey products, they require significant time and effort to fix, and in doing so new grid products are created, so that the lineage of the original field submission is compromised. This can create complications during the ensuing review (Wilson et al., 2016).

Unfortunately, manual methods of identifying fliers appear to be ineffective: in 2015, it was reported that fliers affected up to 25% of the surveys received at Coast Survey (Gonsalves, 2015). It therefore seems likely there may be some advantage in finding a means to automate the process.

QC Tools is a suite of applications that targets common HO challenges. It was developed within the HydrOffice project (https://www.hydroffice.org/), a collaborative effort led by the Center for Coastal and Ocean Mapping (CCOM/JHC) to ease the construction of ocean mapping tools by lowering the barrier to implementation. Although this paper focuses on the Coast Survey requirements, the modularized architecture of QC Tools facilitate customization to any HO's survey specifications and in-house best practices.

A component part of QC Tools is Flier Finder (Wilson et al., 2016), an algorithm to identify potential fliers. The fundamental idea of this tool (and many others in QC Tools), is to prioritize the parts of the data that require manual intervention, rather than spreading user effort equally across all the data. Flier Finder was enhanced in 2016 with several new flier-identification techniques, as well as auto-estimation of the sensitivity of the search based on the characteristics of the grid to be scanned.

Although it is almost axiomatic that an automated technique to assist in finding fliers would have benefit, this remains to be proven. This paper reports an experiment comparing the behavior of the recently updated Flier Finder algorithm to that of human operators with varying levels of expertise. In particular, it attempts to answer the question: is a computer-assisted Flier Finder more effective than a human operator?

Manual and Automated Methods of Flier Detection

Modern techniques of survey review are based on the fundamental idea that we no longer need to inspect every sounding acquired—instead we use metrics and visualization of the bathymetric grid to identify areas to “spot-check”. The most common methods used to find fliers involve a close examination of the various grid layers, extracting shoal-biased sounding selections, and usage of 3D/2D viewers. A variety of methods for flier detection are shown in Figure 1.

Examining available grid metrics offers a multitude of options for the user, and vertical exaggeration may be applied to the depth layer to accentuate potential anomalies. Statistics-based grids generally include standard deviation and uncertainty layers, and CUBE (Combined Uncertainty and Bathymetry Estimator)-based grids contain the number and strength of the related hypotheses (Calder and Mayer, 2003). Processing software generally includes the option to set view filters to the layers, apply customized color maps, and to create customized layers using grid math with the other layers available. All of this functionality can be extremely useful during grid review, but it does leave a lot of (subjective) variability to the data analyst workflow.

Visual inspection by 3D/2D viewers offers the great benefit of human intuition, which will always be better than any automated algorithm in the evaluation of a single spot-check. The user can detect anomalous data qualitatively should they happen to come across it. But therein lies the disadvantage of this method, and in the review of the various grid metrics: they are dependent upon the reviewer finding the anomaly. Rigorous manual grid review can be effective, but has a component of subjectivity. Furthermore, the user can never be entirely sure the entire grid was checked thoroughly enough. To illustrate the point, the grid shown in Figure 1 contains an estimated 68 million grid nodes, and the flier shown is just one of those. The flier is accentuated by visualization techniques, but with 68 million other nodes around it, it still might be considered the veritable “needle in a haystack.”

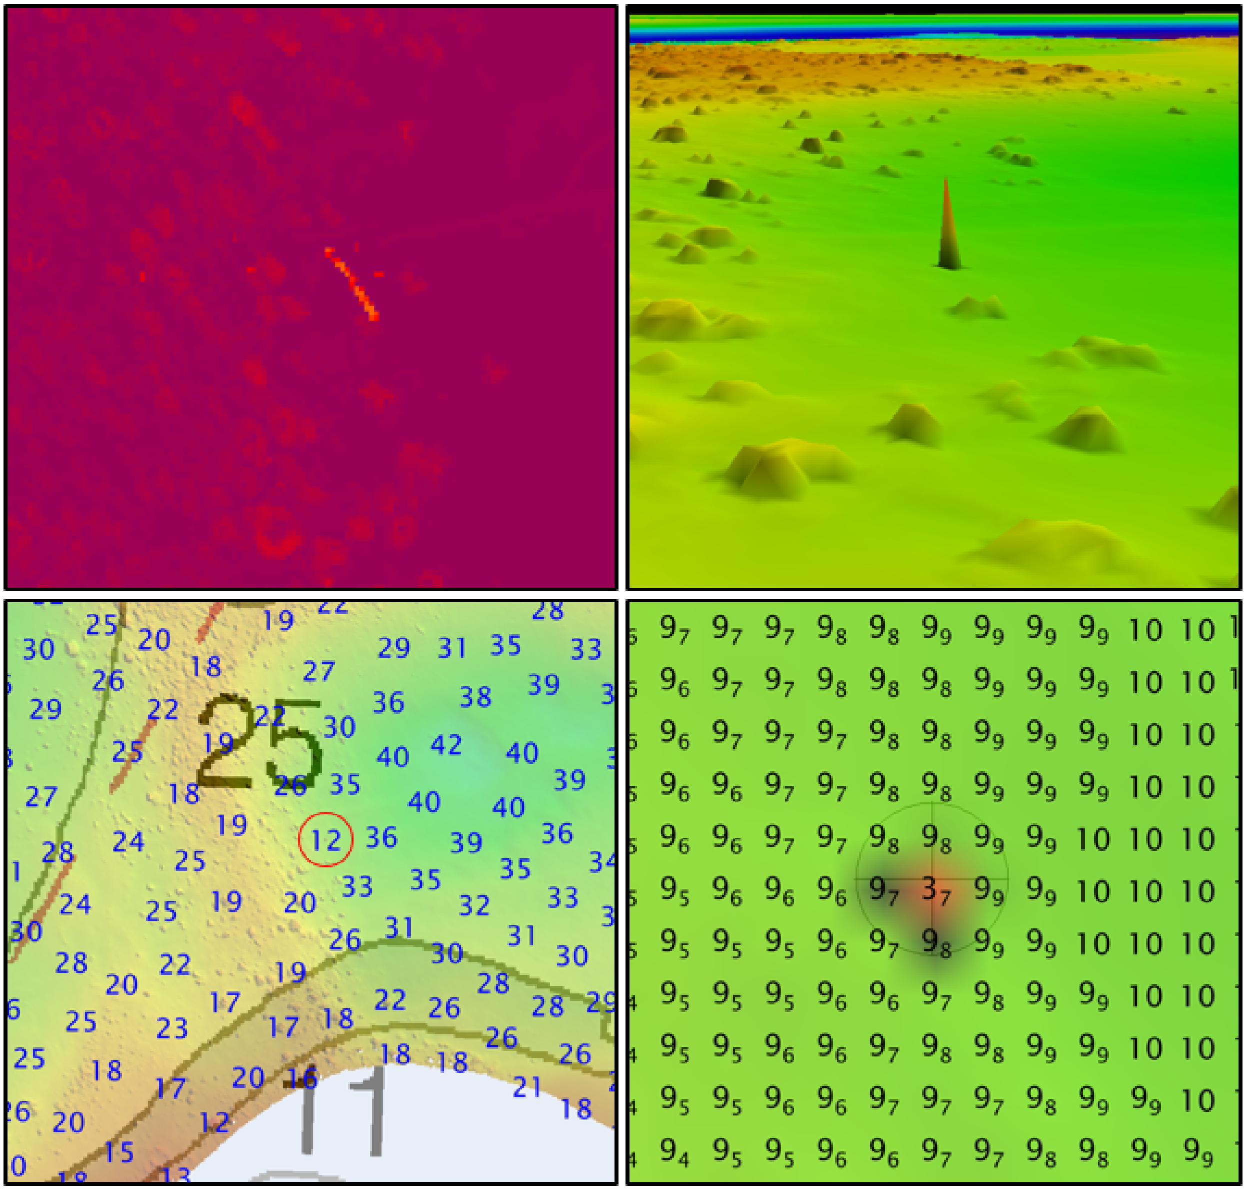

Figure 1. Various means of detecting the same grid flier: an examination of the various grid metrics, which in this case is the node standard deviation that clearly shows the burst of noise that birthed this flier (top left); a 3D view of the grid provides a very intuitive suggestion that the depth in question is not real (top right); a shoal-biased sounding selection (in feet) shows a 12 foot sounding well offshore of both the 12- and 18-foot contours, which should suggest a closer look as a potential flier or barring that a danger to navigation (bottom left); and an automated scan of the grid depth layer (node depths in meters) set at a six meter height or less would easily flag the node for review (bottom right).

A shoal-biased sounding selection contains considerably fewer data points than the grid from which it was derived, therefore the user can feel more confident that each data point was checked. This is a considerable advantage over the previous methods, although still can require a significant amount of time to review. Shoal-biased sounding selections also offer the advantage of chart context: soundings atop anomalous data may differ considerably from the chart and so draw the attention of the reviewer. There is a side benefit in the chart comparison exercise and potential discovery of dangers to navigation. Disadvantages of this method are that “deep” fliers have no chance of detection, nor will fliers that do not significantly disagree with the chart.

Unlike each of the manual techniques, the automated scan conducted via Flier Finder allows the user to feel absolutely assured that all of the nodes in the grid were evaluated objectively, and this alone is a tremendous advantage. The disadvantage is that the evaluation is only as good as the algorithms applied during the scan, which might either miss the anomaly, or it may yield excessive "false positives" (i.e., declarations of problems where none exist, which require the reviewer to commit some time to disprove).

Flier Finder

The Flier Finder algorithm works by scanning the depth layer of the grid and searching for potential anomalies in the form of shoal or deep "spikes". If a spike is found, a flag is registered, indicated the depth node that triggered the flag has been marked for further review. A record of the flag is written to an S-57 file that can be overlaid on the grid in order to guide the subsequent review.

The current QC Tools implementation of Flier Finder offers three different algorithms designed to flag shoal and deep spikes in the depth layer of the grid ("Laplacian Operator", "Gaussian Curvature", and "Adjacent Nodes"), and two more designed to flag nodes separate from the grid altogether, either adjacent to the edges ("Edge Slivers") or far detached ("Isolated Nodes"). These may be enabled or disabled by the user as desired, and a flier search height may be estimated automatically, or can be manually set by the user. These user options are shown in Figure 2.

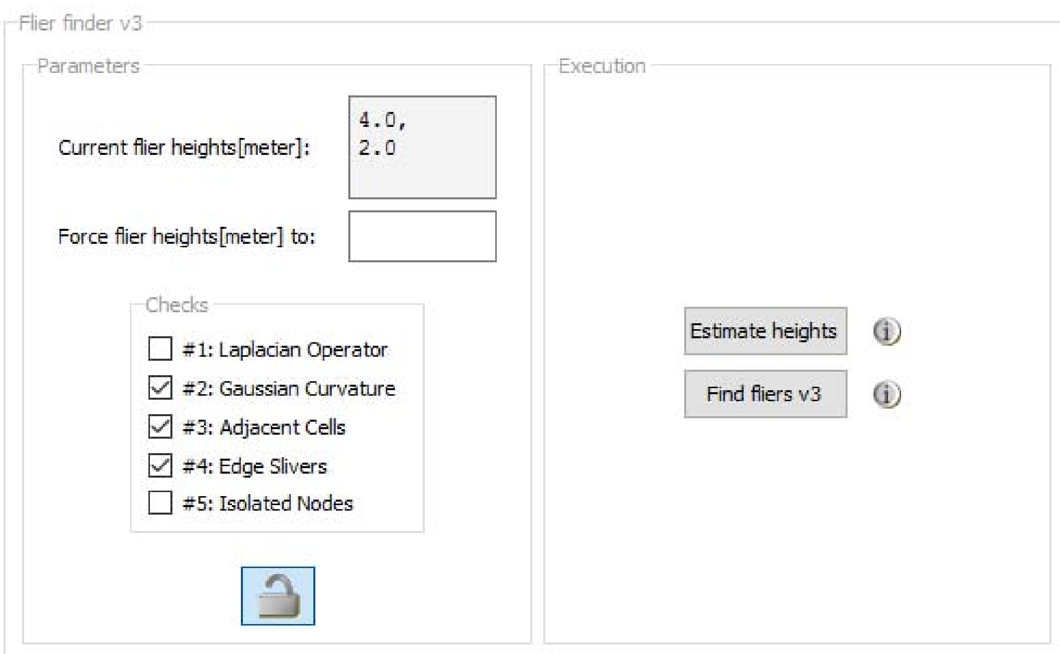

Figure 2. After one or more grids have been loaded, the graphical user interface of Flier Finder v3 is available in QC Tools. "Estimate heights" will automatically suggest a flier search height for each grid input, and in the example above heights of 4.0 and 2.0 meters are estimated for the two grids loaded. Default use of Flier Finder v3 is the execution of algorithms #2, #3, and #4 at these estimated heights, but a user has the option to "unlock" the parameters to enable or disable the algorithms for a customized search. The user may also force a flier search height, which will override the estimated one.

Algorithm Descriptions

The Laplacian Operator is a measure of curvature of the depth surface at each node, equivalent to summing the depth gradients of the four nodes neighboring the node of interest. Absolute values greater than four times the flier search height are flagged. This algorithm is not enabled by default because testing showed it was sometimes prone to excessive false positive detections. It can be used on an as-needed basis as a means of increasing the sensitivity of Flier Finder.

The Gaussian Curvature is a measure of concavity at each node. Gradients are taken in the x and y direction to establish g_x and g_y, and repeated on each gradient (again in the x and y directions) to establish the mixed second derivatives g_{xx}, g_{xy}, g_{yx}, and g_{yy} (note that $g_{xy} = g_{yx}$). These terms are the inputs into the calculation for Gaussian Curvature,

$$K = \frac{(g_{xx} \times g_{yy} - g_{xy} \times g_{yx})}{(1 + g_x^2 + g_y^2)^2} \quad (1)$$

Values of K greater than six are flagged, which was deemed an effective threshold during testing. Note this algorithm is independent of flier search height, and it is enabled by default.

The Adjacent Nodes algorithm will difference the depth at the nodes surrounding a node of interest (up to a maximum of 8 neighbors). The number of instances in which the depth difference exceeds the flier search height is tallied, and if the ratio of the tally to the total number of surrounding nodes is 0.8 or greater, the node is flagged. This threshold was chosen because of its observed effectiveness during testing, and the algorithm is enabled by default.

Examples of algorithms #1, #2, and #3 are shown in Figure 3.

9	9	9	9	0	1	0	3
9	8	9	6	1	4	**10**	**9**
9	9	3	9	0	7	**24**	**9**
9	9	9	9	0	0	6	0
-1	0	-1	-0.09	0	0	0.2	0.33
0	-2.3	0	-0.14	0	0.13	0.25	**1**
-2.5	0	**20**	0	0	0	**1**	0.4
0	-9	0	-36	0	0.2	0.2	0.33

Figure 3. Generic depth values are given (top left), and a flier search height of two meters is enabled. The Laplacian Operator (top right) sums the depth gradients in the x and y direction and flags any absolute values greater than four times the search height. Values of Gaussian Curvature (bottom left) greater than six are flagged, regardless of flier height. Up to eight surrounding nodes are differenced with each node in Adjacent Nodes (bottom right), and if the ratio of number of times the difference exceeds the flier search height to the number of surrounding nodes is greater than 0.8, the node is flagged. In this example, algorithms #1, #2, and #3, at a two meter search height, result in 4, 1, and 2 flags, respectively (shown in bold red font). Note that the Laplacian Operator is most sensitive to adjustments in search height; conversely, the number of flags returned by the Gaussian Curvature never changes.

Edge Slivers identifies small groups of interconnected nodes that are detached from, but within a certain horizontal distance to, the grid. If the depth difference between the nearest detached nodes and grid is greater than the flier search height, a flag is registered. This algorithm is focused on detecting the fliers that frequently occur in the areas of sparse sounding density on grid edges, and it is enabled by default.

Isolated Nodes flags any remaining nodes outside of the five-node horizontal distance, regardless of the flier search height. The intent is to flag those nodes far detached from the grid and perhaps unnoticed by the user, because these often result in problems later, likely when combining individual grids of various resolution. It is not enabled by default and is used on an as-needed basis as an independent check for detached grid nodes.

Finally, the outcomes of these methods are merged to provide the survey analyst with a set of candidate fliers.

Flier Search Height Estimation

A shortcoming of previous versions of Flier Finder was the required user input of flier search height. This input was quite subjective, and the output was very much dependent on this parameter: finer search heights in areas of deep and dynamic bathymetry might result in excessive false positives, but a coarse search height along relatively shallow and flat seafloor might result in no flags at all. This was problematic, especially during the Rapid Survey Assessment (RSA), a 2015 Coast Survey initiative that utilized QC Tools to quickly identify significant deficiencies in survey deliverables, returning surveys to the field for rework if necessary (Berkowitz, 2015). The subjective flier search height input parameter led to varying output obtained from Flier Finder between field and office, which was problematic. Ideally, for such acceptance purposes, standard user parameters should exist or be mandated so the output is easily repeatable.

To mitigate this issue, a flier search height is now estimated automatically from the grid. The search height estimation attempts to find the “sweet spot” for the most efficient use of the tool, so that a reasonable number of potential fliers are returned in order to guide the review, while not overwhelming the user with excessive false positives.

The flier search height is estimated from the depth layer of the grid, using the median depth, depth variability, and roughness. A base height is assigned by the median depth, and increases are made incrementally according to the level of depth variability and roughness. Depth variability is estimated by the normalized median of absolute deviation, which is the ratio of the absolute difference of mean and median depth to standard deviation. Roughness is estimated by the standard deviation of the Gaussian Curvature of all the nodes. Increases are made along a rigid scale of heights for a standard range of estimated flier heights. The exact input, output, and decision criteria for common Coast Survey depth ranges is shown in Figure 4, and are the result of extensive testing on the bathymetric grids accompanying real data submissions to Coast Survey. The end result has shallower, flatter seafloors tending toward a very fine search height, sensitive to anomalies of small magnitude, whereas deeper, dynamic, and rocky seafloors tend toward a coarse search height, making the algorithm less sensitive to rocks and steep slopes.

Depth Median	< 20	20 - 40	40 - 80	80 - 160	> 160
Base Height	1	2	4	6	8

Normalized Median of Absolute Deviation	# of increases
> 0.2	0
0.2 - 0.1	1
< 0.1	2

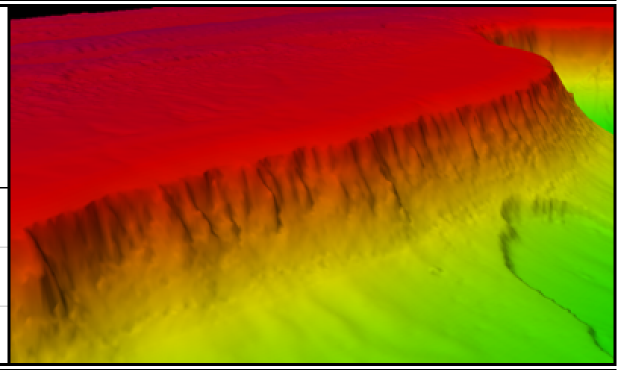

Standard Deviation of Gaussian Curvature	# of increases
< 0.01	0
0.01 - 0.10	1
> 0.10	2

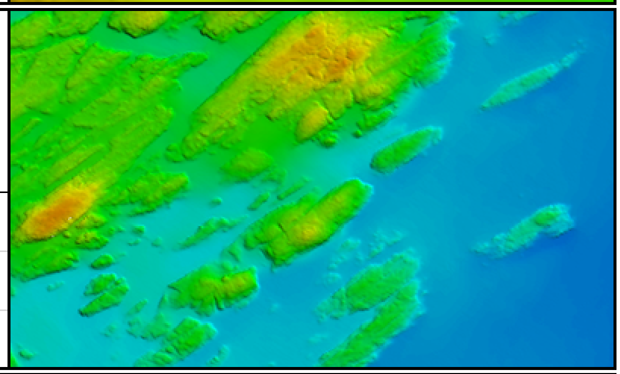

Advance base height along the scale below by the total # of increases

Final Height	1	2	4	6	8	10	12	14	16

Figure 4. QC Tools Flier Finder will estimate a flier search height from the depth layer of a bathymetric grid with the goal of flagging a reasonable number of locations for further inspection, while not overwhelming the user with excessive output. The depth median of the grid sets the base flier height (meters), and increases are made based on estimates of depth variability and roughness. The normalized median of absolute deviation, the ratio of the absolute difference between depth mean and depth median to depth standard deviation, is a proxy for depth variability, and the standard deviation of the Gaussian Curvature of all the nodes is a proxy for roughness. The number of increases based on these values is shown above, and corresponds to the number of advances along the scale to produce a final estimated flier search height (meters).

Advantages and Disadvantages

Flier Finder evaluates all the nodes in a grid and actively identifies potential anomalies. This is a great advantage, and particularly useful for finding fliers of smaller magnitude that are more difficult to detect manually. The disadvantage of Flier Finder is that the evaluation of all the nodes is only as good as the algorithms applied. The likelihood that a flier is successfully detected by Flier Finder depends only on whether it meets the user-specified search criteria. Setting these parameters at a level too sensitive might overwhelm the user with false positive flags; conversely, setting it too coarse might miss real fliers.

Flier Finder adds the most value when used on high resolution bathymetric grids that are tedious to evaluate manually. Over relatively flat seafloors, the parameters can be set to a very sensitive level to flag even small anomalies. Increasing depth and roughness of the seafloor requires less sensitive parameters, and so only the most egregious of the anomalies might be flagged, while those of a lesser magnitude will not be detected. There are some regions, particularly offshore of the Alaskan coast, where the steep slopes associated with the local bathymetry may prohibit the use of the current implementation of Flier Finder altogether. The steep slopes in these regions can cause false positive flags on the order of hundreds, or even thousands, an amount of output far too cumbersome to add any value for the user.

The issue of high false positive rate in such regions of steep slopes can be mitigated with future algorithm development in Flier Finder that implements an adaptive flier search height. Many of the false positives are the result of undesirable "edge effects" resulting from depth thresholds applied to the grids. Performing Flier Finder prior to the application of depth thresholds greatly alleviates the false positives; however, this would require a routine to either ignore those nodes with depths outside of the required interval for that particular resolution. While these ideas show encouraging results in testing, they have not been further developed, as reconfiguring the algorithms toward compatibility with variable resolution grids, anticipated for use in 2017 by Coast Survey (Eisenberg, 2017), may potentially resolve many of these issues.

Flier Metrics

It is critical to objectively evaluate the effectiveness of automated scanning to detect grid fliers. A series of experiments were therefore conducted comparing the performance of automated detection techniques (in Flier Finder) with the behavior of human operators working on the same data.

Test Data

In order to establish a baseline for comparison, two control grids were established with several known fliers. Several personnel attempted to detect the fliers using manual methods, and their results were then compared to the results of automated scanning in terms of time required and the success rate of identification.

The two control grids are subsections of a bathymetric grid submitted to Coast Survey in 2015 as part of a standard survey submission. Several fliers were found during the review of this data. The control grids used for these metrics are the same data as was submitted, so that they can be considered representative of actual data quality issues often experienced during the review process. That said, the grids chosen are very conducive to automated scanning, which has been found to be more effective in certain types of bathymetry. In this case, it is an area that required NOAA "object detection" multibeam coverage, i.e., an area of critical under-keel clearance requiring high resolution to detect all significant features (NOAA, 2016). As such, the grids chosen for this exercise were those constrained to depths between 0 and 22 meters, and with a resolution of 50 cm, with the goal of detecting all features measuring at least 1m x 1m x 1m.

The sounding data accompanying the control grids was used to definitively identify 42 instances in which spurious soundings affected the grids. Though more grid fliers may potentially exist, it is believed that the 42 grid fliers identified represent uncontroversial anomalies of two meters or more (in some cases substantially more). In this group of 42 identified fliers, there is a clear distinction between edge fliers (26), and the remaining fliers that reside within the grid (16). The edge fliers are quite common due to the sparse sounding density on the grid edges. Of those fliers located within the grid, six are shoal fliers considered to be of most significance because of their impact to shoal-biased sounding selections and contours later extracted during chart application. Though these are most significant, all of the 42 fliers identified in the control grids are important, because their presence in submitted surveys could potentially result in RSA rejection (Berkowitz, 2015).

Examples of edge, shoal, and deep fliers identified in the control grid are each given in Figure 5.

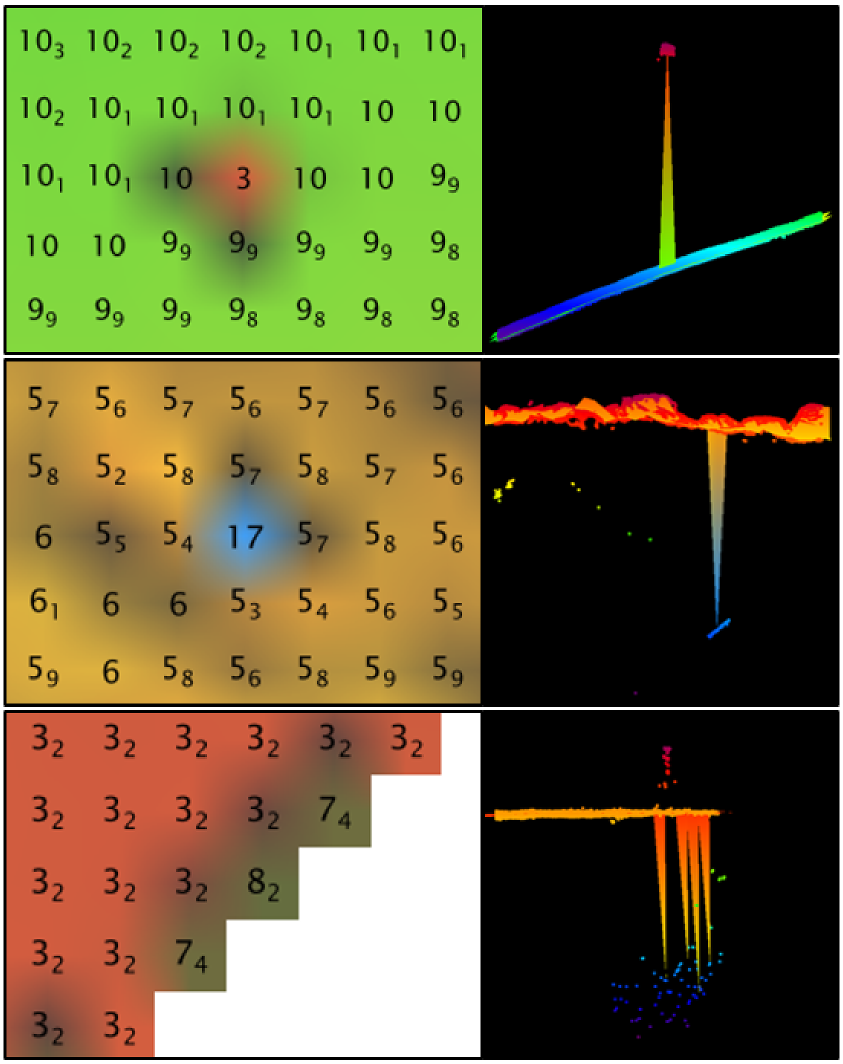

Figure 5. Examples of anomalous data "fliers" from the control grids include shoal fliers (top), deep fliers (middle), and edge fliers (bottom), which can be shoal or deep. Each example shows an overview of the grid (left) with the depth of each node (meters), and a 2D view of the soundings and grid (right).

Method

The two control grids were distributed to seven volunteer reviewers who had no previous experience with the dataset, with instructions to manually examine the grids and place flags atop any potential anomalies or areas they believed worthy of further inspection. The reviewers were also given the corresponding nautical chart for context. The sounding data was not supplied, however, so the reviewer could not verify potential anomalies by looking at the original sonar data. Because of this, the time and effort spent mimics the first part of the flier identification, which is to use the various forms of grid review to determine a list of areas to inspect more closely. This same part of the process is automated in grid scanning techniques such as Flier Finder, which was used separately, so the output could be compared with that from the manual reviews.

The second part of the grid review—examining the sounding data to verify potential anomalies—was not part of the exercise, as this part can only be done manually. However, the time and effort required to complete this verification might be considered proportional to the number of flags made by the reviewer (or the number of flags returned by the automated scanning). Note, this process would be lengthened with excessive false positives.

Reviewers returned their list of flags in a GIS layer, with locations marked where they believed further examination necessary from the manual methods of grid review. They also recorded the time spent to generate the list, as well as their years of experience in hydrography.

When performed on the same two control grids used in the manual reviews, Flier Finder estimated heights of four meters (control grid "A") and two meters (control grid "B"). Though the depths within each grid are quite similar, control grid A contains expansive rocky areas and is therefore more challenging for manual review, while also requiring a slightly less sensitive flier search height due to the roughness. The default algorithms (#2, #3, and #4 from Figure 2) were used on both control grids. Thus, Flier Finder was operated with the standard, default parameters.

Results and Discussion

For each of the 42 verified fliers, Table 1 lists the control grid of origin, the flier type, the approximate magnitude of the anomaly, the number of times and percentage the flier was detected manually, and whether the anomaly was flagged by the various components of the Flier Finder algorithm when operated with default parameters. Also given are the overall percentages of detection for the three different types of fliers, for both the manual and automated reviews.

Table 1. Results of both manual and automated detections for 42 verified fliers

Flier #	Control Grid	Flier Type	Magnitude of flier (m)	Manual Detections							Automated Scan[1]			
				# of times detected (7 reviewers total)					%		Algorithm #			%
											#2	#3	#4	
1	B	edge*	1	✓					14%				✓	
2	B	edge	4	✓					14%			✓		
3	B	edge*	4	✓	✓	✓			43%				✓	
4	B	edge	3	✓	✓	✓			43%			✓		
5	B	edge	2	✓					14%					
6	B	edge	2	✓					14%			✓		
7	B	edge*	2	✓	✓				29%				✓	
8	B	edge*	3	✓					14%				✓	
9	B	edge	1	✓	✓	✓			43%					
10	A	edge*	2						0%				✓	
11	A	edge*	2	✓	✓				29%				✓	
12	A	edge*	3	✓					14%				✓	
13	A	edge	3	✓	✓	✓	✓		57%	**26%**				**85%**
14	A	edge	9	✓	✓	✓	✓		57%			✓		
15	A	edge*	2	✓	✓	✓	✓		57%				✓	
16	A	edge*	2	✓					14%				✓	
17	A	edge	4	✓					14%			✓		
18	A	edge*	2	✓					14%				✓	
19	A	edge	5	✓					14%			✓		
20	A	edge	4	✓	✓	✓			43%			✓		
21	A	edge	4	✓	✓	✓			43%			✓		
22	A	edge	4	✓	✓	✓			43%			✓		
23	A	edge	16	✓					14%		✓			
24	A	edge	2	✓	✓				29%					
25	A	edge	11						0%			✓		
26	A	edge*	2	✓					14%				✓	
27	A	deep	4						0%			✓		
28	A	deep	5						0%			✓		
29	A	deep	11						0%			✓		
30	A	deep	3						0%					
31	A	deep	4						0%			✓		
32	A	deep	10						0%	**0%**	✓	✓		**90%**
33	A	deep	11						0%		✓	✓		
34	A	deep	4						0%			✓		
35	A	deep	4						0%			✓		
36	A	deep	6						0%		✓	✓		
37	B	shoal	4	✓	✓	✓	✓		57%			✓		
38	B	shoal	4	✓	✓	✓	✓		57%			✓		
39	B	shoal	4	✓	✓	✓	✓		57%	**45%**		✓		**83%**
40	B	shoal	6	✓	✓	✓	✓	✓	71%		✓	✓		
41	A	shoal	3	✓					14%					
42	A	shoal	6	✓					14%		✓	✓		
									22.8%				**85.7%**	

[1]performed at auto-estimated search heights (4m for grid A, 2m for grid B) and using default algorithms
*flier is actually detached from grid

As shown in Table 1, the overall rate of manual detection was 22.8%. As expected, shoal fliers of the highest significance were detected more often. There was also an observed bias toward those fliers occurring along the grid edges (both shoal and deep). However, the internal deep fliers were not found by any of the reviewers, though some had significant magnitudes. The manual reviewers showed considerable variation in the number of flags registered, with a minimum of 32, a maximum of 779, and an average of 224 flags submitted. These yielded estimated false positive percentages in the range of 75% to greater than 99% with an average of 84%.

The overall average rate of detection for Flier Finder operated at default search heights and parameters was 85.7%, with each of the enabled algorithms contributing to the detections. Decreasing the flier search height in Flier Finder will increase the detection rate to 100% with the disadvantage of increasing the number of false positives. Operating with the default parameters returned a total of 43 flags, with an estimated false positive rate of 16%. Adjusting the parameters affects the output considerably, as shown in Table 2, for control grid A, and Table 3, for control grid B. The numbers in red font represent the number of flags returned at the default settings. With forced lower search heights, and 42 verified fliers in the control grids, one can infer both the potential for high false positives, as well as the value of the auto-estimated search parameters. Note, the Laplacian Operator was used in this assessment of output sensitivity to flier search height, but Isolated Nodes algorithm was not, because it flags any detached nodes and is independent of search height.

Table 2. Flags returned from various Flier Finder settings applied to control grid A.

Flier Search Height	#1 - Laplacian Operator	#2 - Gaussian Curvature	#3 - Adjacent Cells	#4 - Edge Slivers
1	4187	6	477	50
2	99	6	34	19
3	11	6	20	9
4	4	**6**	**12**	**7**
5	4	6	10	4
6	2	6	7	2

Table 3. Flags returned from various Flier Finder settings applied to control grid B.

Flier Search Height	#1 - Laplacian Operator	#2 - Gaussian Curvature	#3 - Adjacent Cells	#4 - Edge Slivers
1	819	1	153	32
2	9	**1**	**8**	**9**
3	4	1	5	3
4	2	1	4	3

Reviewers spent as little as 43 minutes conducting the review, and as much as 4 hours, for an average time of 2 hours and 36 minutes. Again, this was the time needed to inspect the control grids and create a list of marked locations that would then require further manual inspection. The time required for Flier Finder to estimate heights and run the various algorithms (to arrive at a list of marked locations) is just a few minutes (and it might be automatically executed as a background process since human intervention is not required). Thus, there is considerable time efficiency gained in using Flier Finder to help generate this list. The remaining part of the review (manually inspecting the sounding data associated with potential grid anomalies) is correlated to the overall number of flagged locations, and so in this case may actually be faster with the Flier Finder output. This is not always the case, however, as Flier Finder often results in much higher rates of false positives when used on certain types of bathymetry.

Reviewers ranged from three to 14 years of experience in hydrography, with an average of nine years, but no correlation between experience level and effectiveness of their review could be found. Likely, the sample size is too small, but there also appeared to be considerable variation in the methodology used, and the number of flags recorded manually. This suggests a need for better training and standardization in the methods of flier detection, and to a greater extent, grid review. To this end, highly repeatable methods such as Flier Finder will be useful.

It should be reiterated that the control grids established in this exercise are those very conducive to the automated scanning performed by Flier Finder. Additionally, the selected control grids are quite difficult to review manually due to their high resolution. Thus, this is the very best case scenario for automated scanning, chosen purposefully to demonstrate its effectiveness. Note, however, that such a scenario is not at all unusual, particularly in the bathymetry prevalent along the U.S. east coast and Gulf of Mexico.

It is important to note that none of the flier detection methods presented in this paper are guaranteed. Flier #42 (from Table 1) is a very significant, shoal flier—it is the same flier shown in Figure 1, where it is presented via several different detection methodologies. However, it was detected only one time during the manual reviews. This is not to suggest any shortcomings of the reviewers, but rather to suggest that their objective—and what has been asked of them for several years—may be much more difficult to achieve than it is commonly believed. Flier #42 is a single node, residing in a grid of 50 cm resolution among roughly 68 million other nodes. For this single node to be detected manually—each and every time this situation arises, and by personnel of various levels of experience—may not be a reasonable expectation. High occurrences of data quality issues in Coast Survey resulting from anomalous grid data fliers (Gonsalves, 2015) clearly reflect this.

The manual methods of flier detection presented in Figure 1 represent various means of grid visualization and presentation of grid metrics, but each of these are indirect methods of flier detection, i.e., these methods are not actively seeking out anomalies. Conversely, the automated scanning via Flier Finder is a method for direct flier detection: it is the only method that will actively examine all of the nodes and flag anomalies that fail a particular evaluation. The results of this case study clearly demonstrate their potential effectiveness in certain conditions, as do the recent observations that suggest improved timeliness and data quality in Coast Survey, in-part due to wide adoption of QC Tools (Evans, 2017), both in the field and in the office.

However, Flier #41 (from Table 1) is also a significant shoal flier, found by only one reviewer, and not found by Flier Finder when operated with the default parameters. Had the search height been forced to three meters (rather than the estimated height of four meters), it would have been flagged, but there is no guarantee a user would have done so. The automated scanning is reassuring, in that the user knows all the nodes are scanned; however, the scan is only as effective as the algorithms applied. In this case, Flier #41 was not captured in the scan using the default parameters, but at least one human happened to find, recognize, and flag it. This also provides a clear hint on where to focus for the future development of the Flier Finder algorithm.

Clearly, automated tools such as Flier Finder should never totally replace qualitative assessment (Evans, 2017), as this remains a critical part of survey review. Ideally, automated tools are a supplement to the manual methods, so that in this case it would be recommended that Flier Finder be used in-conjunction with manual methods. In this manner, automated scanning is simply one of several tools in the toolbox. Whether a manual method or an automated one, no method of flier detection presented in this paper is guaranteed to find all the fliers. The key is to use each of them, and in doing so, achieve the highest likelihood of detecting all such anomalies, with the important side effect of increasing confidence in the reviewed survey data. Finally, once identified as statistically significant, some of the automated flier identification techniques could be directly integrated in the processing algorithm, reducing, at the source, the rate of fliers delivered in the created surfaces.

Conclusions

Anomalous grid data fliers cost time and effort to resolve, and if unnoticed, they reduce accuracy and data quality. They are a problem for many hydrographic offices, Coast Survey included, with an estimate from 2015 of 25% of all survey submissions affected by fliers.

Successfully detecting anomalous grid data fliers requires multiple forms of input and different inspection methods, some of which are tedious and monotonous for survey reviewers, which makes them ideal candidates for automated scanning. Find Fliers does not aim to completely replace qualitative, human review, but instead to make it more effective by providing bulk scanning of large data sets. The tests performed and the metrics compiled in this case study comparing human to machine review processes suggest that in certain bathymetry such automated scanning is not just effective, but might actually be critical in achieving higher data quality through the elimination of anomalous grid data fliers. Recent increases in data quality and timeliness in Coast Survey has been attributed in part to the adoption of HydrOffice QC Tools, which implements the Flier Finder algorithm. Both the encouraging results and the highlighted challenges represent natural feedback with which to drive the future developments of QC Tools.

Acknowledgements

Thanks first and foremost to the brave NOAA personnel who participated in this case study, which by design was very challenging. Thanks also to NOAA Coast Survey and UNH CCOM/JHC for actively supporting new products and innovation, and to the Coast Survey Hydrographic Systems and Technology Branch for their help in the integration and distribution of software. Lastly, thanks to all the NOAA users for their enthusiasm and feedback for QC Tools, and to a greater extent, HydrOffice. CCOM/JHC participation in this work was supported by NOAA grant NA15NOS4000200, which is gratefully acknowledged.

References

Berkowitz, E., 2015, Rapid Survey Assessment (RSA) Workflow, OCS QMS Controlled Document, 10 July 2015

Calder, B. R., and Mayer, L. A., 2003, Automatic processing of high-rate, high-density multibeam echosounder data: Geochemistry Geophysics Geosystems, v. 4, no. 6, p. 1048.

Eisenberg, J., Variable Resolution Implementation, *in* Proceedings NOAA Coast Survey Field Procedures Workshop, Virginia Beach, VA, USA, 24-26 January 2017.

Evans, B., What are our Shared Challenges, *in* Proceedings NOAA Field Procedures Workshop, Virginia Beach, Virginia, USA, 2017.

Gonsalves, M., Survey Wellness, *in* Proceedings NOAA Coast Survey Field Procedures Workshop, Virginia Beach, Virginia, USA, 27-20 January 2015.

NOAA, 2015, Strategic Plan 2015-2019. Navigate with Confidence: Office of Coast Survey.

NOAA, 2016, National Ocean Survey Hydrographic Surveys Specifications and Deliverables. March 2017.

Wilson, M., Masetti, G., and Calder, B., NOAA QC Tools: Origin, Development, and Future, *in* Proceedings Canadian Hydrographic Conference, Halifax, Nova Scotia, Canada, 2016.

Author Biographies

Matthew Wilson is a physical scientist with the NOAA Office of Coast Survey. He received an MS degree in Ocean Mapping (UNH, USA) in 2012, and an MBA degree (Penn State, USA) in 2016. He works at the Atlantic Hydrographic Branch in Norfolk, Virginia, involved in data verification, software development, and ship support.

Giuseppe Masetti received an MS degree in Ocean Engineering (UNH, USA) in 2012, and a Masters in Marine Geomatics (2008) and Ph.D. (2013) in System Monitoring and Environmental Risk Management (University of Genoa, Italy). As a Research Assistance Professor at CCOM/JHC, he is focusing on hydrographic data processing and acoustic backscatter.

Brian Calder is a Research Associate Professor and Associate Director at CCOM/JHC (UNH, USA). He has a M.Eng (Merit) and Ph.D. in Electrical and Electronic Engineering, with a thesis that focused on Bayesian methods in sidescan processing, from Heriot-Watt University, Scotland (1994, and 1997, respectively). He is currently focusing on statistically robust automated data processing approaches and tracing uncertainty in hydrographic data.