5-2007

Marine Heritage Monitoring with High Resolution Survey Tools: ScapaMAP 2001-2006

Brian R. Calder
University of New Hampshire, Durham, brian.calder@unh.edu

Bobby Forbes
University of New Hampshire, Durham

Duncan Mallace
University of New Hampshire, Durham

Marine Heritage Monitoring with High-Resolution Survey Tools: Scapa Flow 2001-2006

Brian R. Calder[1], Bobby Forbes[2], and Duncan Mallace[3]

Abstract

Archaeologically, marine sites can be just as significant as those on land. Until recently, however, they were not protected in the UK to the same degree, leading to degradation of sites; the difficulty of investigating such sites still makes it problematic and expensive to properly describe, schedule and monitor them. Use of conventional high-resolution survey tools in an archaeological context is changing the economic structure of such investigations however, and it is now possible to remotely but routinely monitor the state of submerged cultural artifacts. Use of such data to optimize expenditure of expensive and rare assets (e.g., divers and on-bottom dive time) is an added bonus.

We present here the results of an investigation into methods for monitoring of marine heritage sites, using the remains of the Imperial German Navy (scuttled 1919) in Scapa Flow, Orkney as a case study. Using a baseline bathymetric survey in 2001 and a repeat bathymetric and volumetric survey in 2006, we illustrate the requirements for such surveys over and above normal hydrographic protocols and outline strategies for effective imaging of large wrecks. Suggested methods for manipulation of such data (including processing and visualization) are outlined, and we draw the distinction between products for scientific investigation and those for outreach and education, which have very different requirements. We then describe the use of backscatter and volumetric acoustic data in the investigation of wrecks, focusing on the extra information to be gained from them that is not evident in the traditional bathymetric DTM models or sounding point-cloud representations of data.

Finally, we consider the utility of high-resolution survey as part of an integrated site management policy, with particular reference to the economics of marine heritage monitoring and preservation.

1 Introduction

The last decade or so has seen a significant change to our attitudes towards submerged cultural resources. Once only the province of a few divers, the tremendous expansion of recreational diving in the late 80s and early 90s has seen the need for knowledge and access by this group and the general public rise exponentially. Management of this heritage is, however, a delicate balance between exploitation and conservation. The remains of the Imperial German Fleet (from World War I) in Scapa Flow, Orkney (Figure 1) are a particular example of this. The wrecks, due to their unique nature and relatively shallow scuttle depths, have been active recreational dive sites for more than twenty years and form a significant portion of the island economy. Even without intentional damage to the wrecks by divers (which is actively policed by the local dive-boat operators), dive pressure, like 'boot erosion' on land, has a significant effect on the wrecks, making them a non-renewable resource. Any reasonable management strategy must therefore understand the implications of restrictions on use of the resources, and temper this with a desire to preserve the resource for the future.

A basic aspect of management if, of course, monitoring of the current state of the resource. Under Scots law, a geographical delineation of the site to be protected is also required before any legislation can be passed. Due to the very nature of a marine archaeological site, direct monitoring is extremely difficult. At Scapa Flow, for example, the depth of the wrecks, temperature of the water and normal visibility mean that even experienced divers on advanced gas mixtures are capable of only approximately 30 min bottom time on each dive. It is therefore extremely difficult to investigate significant portions of the wrecks with

---

[1] Center for Coastal and Ocean Mapping and NOAA/UNH Joint Hydrographic Center, University of New Hampshire, Chase Ocean Engineering Lab, 24 Colovos Road, Durham NH 03824, USA. (`brc@ccom.unh.edu`)

[2] Scientific Underwater Logistics and Diving Ltd., 13 Franklin Road, Stromness, Orkney KW16 3AN, Scotland (`bobby.forbes@uk-sdsc.com`)

[3] NetSurvey Ltd., Oathill Farm, Claydon Road, Cropredy, Banbury, Oxon, OX17 1QA, UK (`duncan@netsurvey.co.uk`)

any speed, and consequently very expensive to assemble a composite picture of the state of preservation of the whole wreck. In practice, much of the investigation prior to the current work had been done on an *ad hoc*, anecdotal basis, interviewing recreational divers and making occasional targeted dives to the wrecks, recorded primarily through "Diver's impression" sketches (Figure 2).

Clearly, there is scope for improvement in monitoring strategies, and it was with this in mind that we started investigating how modern hydrographic survey methods might be applied to the problem. Starting with a sidescan survey in 1999, we co-opted one of the first ultra-high resolution Multibeam Echosounder (MBES) systems, a Reson 8125, to survey the bathymetry of the wrecks in 2001. Based on the success of this mission, we returned to Scapa Flow in 2006 with a Reson 7125 to re-survey the wrecks, adding to the data mix high-resolution backscatter and full volumetric (watercolumn) acoustic data recording. Between these two extremes, videography of the wrecks has provided detailed descriptions of sections of the wrecks, targeted based on prior bathymetric surveys. Our goals for the surveys, and hence for the current work, were to answer fundamental questions about the use of these technologies in marine archaeology, for example:

- How do we effectively delineate the active area associated with a wreck? If there is a debris field associated with multiple wrecks, how is this best defined?
- How do we best fuse data components in a multimodal survey?
- Is it possible to monitor the wrecks by carrying out a repeat survey? What is the best methodology for repeating a survey to make changes most readily apparent?
- What sort of data is required for scientific study of the site, and how does it differ from that required for public outreach and communication?

while taking advantage of the data to monitor and protect the particular wrecks in Scapa Flow.

In this paper, we investigate the properties of the selected survey instruments, evaluating their place in the spectrum of tools useful for marine archaeology. With very large wrecks, careful coordination of survey effort is required to achieve data that will support scientific study of the wrecks, and particular methods of data processing and visualization are required to optimally convert the data into information. We offer observations on how best to utilize these data, emphasizing the difference between standard hydrographic survey protocols and those used for processing data in a marine archaeological context. We illustrate what can be achieved with standard survey tools in a repeat survey, and the extra "value added" information that is available from such tools, including high-resolution backscatter and volumetric acoustic data. Finally, we summarize the role of these techniques in an integrated site management policy, with particular emphasis on their economic benefit.

2 Background

2.1 Genesis of the Scapa Flow Site

Scapa Flow is an almost totally enclosed expanse of water, bordered to the north by Mainland Orkney, to the east by the Holms, Burray and South Ronaldsay, on the west by Hoy and the south by Flotta. Its relatively sheltered 23 square miles of waters and strategic position therefore lead to Scapa Flow having an extensive maritime heritage from King Hakon's great Viking fleet in 1263, through to Hanseatic merchants in medieval times and convoys to the Baltic during the Napoleonic wars. This also led to Admiral John Jellicoe, RN establishing Scapa Flow as the home base for the British Grand Fleet prior to the outbreak of World War I. The Flow's place in history was cemented, however, by the scuttling of the Imperial German High Seas Fleet.

In November 1918, Germany having signed the armistice agreement at the end of World War I, the warships of the German Imperial Navy's High Seas Fleet entered Scapa Flow to be interned during the peace negotiations. The Allied Grand Fleet had earlier met the German ships in the North Sea and had escorted them to their new anchorage in the Orkneys where they could be guarded by the whole British fleet. Here they were to stay, moored in neat rows around the small islands of Cava and Fara, becoming a local tourist attraction; in total, the fleet consisted of 74 disarmed ships: 5 battlecruisers, 11 battleships, 8 light cruisers and 50 destroyers (Figure 1(b)).

By December the initial 20,000 crew had been reduced down to skeleton crews on all vessels, a total of around 4,800 officers and men. Over the next few months there were considerable discipline problems because of (communist) revolutionary elements in the German crews, with approximately 150 trouble

makers being sent home and Konteradmiral Ludwig Von Reuter, commanding officer of the internment flotilla, moving his flag to the light cruiser SMS EMDEN. By June 1919 the crews were reduced to Royal Navy caretaker levels, a total of about 1,700 men.

On Midsummer's day, 21 June 1919, cut off from his chain of command and fearing the imminent resumption of hostilities between Germany and Britain, Von Reuter ran up a signal hoist from SMS EMDEN reading "Paragraph eleven. Confirm", a reference to a German students'' drinking song calling for more liquid. Pre-prepared, all 74 of the fleet from torpedo boats to battlecruisers repeated the signal and proceeded to scuttle themselves in the British fleet anchorage. The British Grand Fleet had earlier put to sea, leaving only some small fleet auxiliary ships in the anchorage; without significant assets, the British commanders were only able to tow and beach a few of the ships, leaving the vast majority – approximately 437,000 tonnes – to sink to the bottom in depths ranging from 25-40m. In the confusion, eight German sailors lost their lives.

Most of the scuttled fleet did not stay where they had sunk. Those that were beached were removed almost immediately. In the 1920s, the firm Cox & Douglas began salvage operations, lifting many of the ships, a feat not surpassed to this day. This salvage continued until the advent of the Second World War, and only seven scuttled ships now remain on the seabed. After the war, Nundy Ltd started work on the remaining ships to remove the valuable non-ferrous metals by blasting holes in the engine room areas. More salvage took place in the 70s by Dougal Campbell who removed some of the armor plating when "pre-atomic" steel demanded a high price. All of the remaining wrecks have thus sustained heavy damage, but are now one of the best examples of ships from this era still in existence.

2.2 Marine Archaeological Preservation, Protection and Site Management

Due to the concentration of wrecks in close proximity, Scapa Flow is now a recreational diving destination of global renown. Numbers have continually increased throughout the 80s and 90s with approximately 3,000 – 3,500 divers visiting Orkney annually. A recent survey showed that this represented 2% of the tourists who come to Orkney annually but represented 8% of the total tourist revenue to the islands. Divers typically dive for six days doing two dives per day. Research suggests that 65% of all the diving carried out in Orkney is on the seven German wrecks, which represents around 20,000 dives per year

However, by the 90s, Historic Scotland had started receiving approaches from a variety of people and organisations, worried about souvenir hunting. Indeed, many of the ships fittings which had been common place in the early 80s were no longer present. Several articles then appeared on the web and in the diving press highlighting the need for action. A number of possible routes are available to facilitate protection of the sites under Scottish Law. The approach eventually taken was to Schedule the wrecks under the Archaeological Areas and Ancient Monuments Act 1979. This allowed continued access but made it an offence to tamper or remove anything from the sites. The Scheduling came into force on the 23 May 2001.

2.3 The Scapa Flow Marine Archaeological Project (ScapaMAP)

Public access and the occasional souvenir hunter are not the only threats faced by submerged resources, however. The physical and chemical environments also pose a continual mechanism for site formation processes. Coupled with the extensive salvage that had occurred on the remaining wrecks there was a need for high quality information in order that future management strategies could be formulated in an enlightened manner. Historic Scotland therefore provided funding for baseline survey work incorporating traditional archaeological diving techniques and remote sensing protocols used in other areas of marine science.

A consortium of interested groups: Heriot-Watt University, University of New Hampshire, Archaeological Dive Unit, and Reson UK carried out high –resolution survey work in 2000 – 2001 with the final report submitted that year. One of the recommendations of the report was to carry out a repeat high resolution survey every five years. Historic Scotland again provided funds, and Scientific Underwater Logistics And Diving (SULA Diving) organized a second collaborative project in cooperation with the Centre for Coastal and Ocean Mapping, NetSurvey Ltd. and the UK Maritime and Coastguard Agency. The preliminary findings of that work are reported here.

3 Survey, Processing and Presentation Protocols

Multibeam Echosounders are the primary instrument for high resolution hydrographic survey in much of the world. Measuring depth at many (typically 100-256) points across a wide (typically 120-150° total angle) swath below the acoustic instrument at each measurement cycle, a MBES is typically operated to cover the area to be surveyed in a series of wide tracks that are overlapped by choice of ship position to ensure that all areas of the seafloor are ensonified.

Advances in signal processing technology towards the end of the 1990s resulted in a new generation of MBES systems being developed, typified by the Reson 8125. More powerful DSP capabilities allowed these systems to be dynamically focused (that is, to have the acoustic parameters adjusted continually to keep the seafloor in focus as the system receives backscattered energy), to form more individual beams, and to ping the seafloor more rapidly, resulting in a significant increase in data density for any given depth of water. With receive beamwidths on the order of 0.5° at nadir, this allows a nominal acoustic footprint of 0.9% of depth, making the system easily capable of resolving very small targets and therefore feasible for use in survey of wrecks and other objects where the goal is description of the configuration of the wreck rather than just determination of the shallowest point.

As the goals for the survey in an archaeological context are different from those in a hydrographic context, so the methods required for capture, processing and presentation of the results are different. In this section, we discuss the differences from typical hydrographic practice, and illustrate the results with data from the ScapaMAP surveys.

3.1 Organization of the survey

There are two primary goals in any archaeological survey: investigation of the site area as a whole, and documentation of any discrete objects within the site. The former is used to determine sites for the latter, as well as to provide background information. These two modes of operation require different strategies.

Wide scale systematic survey should be carried out essentially as a conventional hydrographic survey, with line spacing of each pass chosen to provide effective overlap with the next. Since the bathymetric data density with most MBES decreases as a function of off-nadir angle, object detection in the outer regions of the swath is less effective, and a higher overlap between swaths than might be used for hydrographic surveys is typical. The concomitant reduction in survey efficiency that this implies can be mitigated by use of appropriately processed backscatter data in some cases, although MBES backscatter is not always ideal for small object detection due to the relatively high grazing angle of observation compared to, e.g., a towed sidescan system.

Item investigation requires careful planning since the goal is to get the observing platform as close as possible to the target (in order to maximize resolution) without causing any potential damage to either survey platform or target. For small targets without significant vertical extent, this can be done simply by laying out a planned track over the centerline of the target (Figure 3(a)). For large targets without significant vertical extent, it is more effective to lay out two planned tracks to port and starboard of the target at approximately one third to a half of a swath width away from the centerline, assuming that two swaths are sufficient to cover the whole target (Figure 3(b)). For targets with large vertical extent, it is typically difficult to ensonify all areas of the target with a single pass since significant areas of acoustic shadows are formed. In this case, survey lines parallel to the centerline are required, laid out so that the track is over the outer extent of the primary wreck site on either side of the center (Figure 3(c)). This arrangement provides the maximum overlap between the two swaths, allowing for conservative processing of the data to maximize the visibility of features from one swath or the other once the combined object is constructed. Multiple-pass surveys of large objects are only effective if positioning is adequate both vertically and horizontally; this is considered in more detail in the following section.

Scientifically, the composite object formed by multiple swaths of data is most effective in understanding the overall shape of the target. For visualization, however, this is not necessarily the case. Irrespective of how well the merging of the swaths is done, the results are generally not particularly convincing (Figure 4). It is usually best to plan and execute a "beauty pass" survey line over the target in order to provide the most compelling image. The placement of this survey line depends heavily on the structure of the target. For example, if the target has an overhang, it might be more effective to survey outboard of the overhang, Figure 5(a), so that the superstructure of the target is imaged directly. If, on the other hand, the wreck is up-right, a survey line down the centerline would be more effective since shadowing should be minimized, Figure 5(b). Many possible alternatives exist. Whichever orientation is

mandated by the target, the survey line should be conducted as slowly as possible while maintaining the survey platform in control, since this maximizes the data density on the seafloor, which significantly improves the visibility of small targets. Regular, stable forward motion of the survey platform is essential in this, however, since excessive yawing or crabbing will result in a confused depiction of the wreck due to the primary geometric structure imposed by the asymmetric along-track and across-track sample spacing, leading to the data being seen primarily as swaths of points, rather than a uniform point cloud.

The sequencing of events during the survey depends on the state of knowledge of the targets before the survey commences. If sufficient information is available to identify all objects of primary interest and their centerlines *a priori*, the prudent approach is to survey the objects first in case of difficulty later in the process. If, however, little is known of the site, or if the object orientation and/or position are ill defined, a reconnaissance/development approach is more useful. That is, the large-scale survey is conducted first, targets are identified and development work is later scheduled on those targets. Whether this development must occur after the primary reconnaissance survey or if it may be interleaved depends strongly on the capabilities of the survey platform, and particularly her operational crew. (The latter is preferred since it develops the most information about known targets earliest in the process.) In either case, an immediate corollary of carrying out target developments is that the data being captured has to be capable of being processed in real-time onboard the survey platform so that the feedback loop can be closed. This has direct implications on hardware, software and personnel availability during the survey. Rapid feedback of results to guide the adaptive development of the survey plan, and sufficient flexibility in the survey plan to allow for this, are essential.

3.2 Data collection

The primary requirements for an archaeological survey are the same as a hydrographic survey: rigid mounting of the sonar transducers to avoid motion artifacts, use of an adequate motion sensor, and careful control and estimation of the offsets of the various components of the system (linear, angular and temporal). The primary differences in data collection are in degree of significance of some of the measurements that are made.

For all survey systems, understanding the sound speed structure of the water column is essential to application of appropriate refraction corrections. Often, it is possible to diagnose issues with multiple overlapping lines and determine empirical corrections if required. In an archaeological survey, however, it is frequently the case that single passes of objects are taken (as described above), so diagnosis of any unknown refraction due to a local micro-change in water mass can be extremely difficult. Since much of the analysis of the data is about shape, and particularly the likely changes in shape over a significant portion of the swath, inappropriate refraction corrections could result in different interpretations of the state of an object. Typically, therefore, an archaeological survey will require more frequent sound speed profile measurements than would be normal for a hydrographic survey, and especially before doing any high-resolution work around a particular feature of interest. The actual frequency required will depend of course on the particular survey area; if the area is particularly shallow and well mixed, fewer measurements may be required, but a protocol of one profile every hour and before each significant target is not unusual.

Hydrographic surveys typically spend a significant amount of time determining the appropriate corrector to reduce the depths measured to an appropriate datum, typically the chart datum for the local area (often Lowest Astronomical Tide or Mean Lower Low Water). Archaeological surveys have both looser and tighter requirements for water level correctors. In most cases, the absolute depth of any feature of the objects being investigated is unlikely to be particularly important. There is therefore no reason to tie the measurements to an absolute datum, since the majority of the information required is about the change in shape and relative position with time. Relative vertical positioning, however, is very important if the object being investigated cannot be covered in one pass of the sonar: small variations in the predicted or observed vertical correctors can significantly affect the coherence of the multiple passes being merged into a composite structure. In extreme cases, problems with vertical correctors could result in a misinterpretation of small structures within the object. Since these objects can be on the order of a few centimeters different in depth, very tight control of relative vertical correctors is required. For both of these reasons, it makes more sense to conduct archaeological surveys on the ellipsoid with Real Time Kinematic (RTK) or Post Processed Kinematic (PPK) GPS measurements, and only connect to a local datum if the data is intended to be dual-use, or needs to be connected to the nautical chart of the area for some reason. This

is not uncommon in archeological surveys, where a *quid pro quo* agreement with hydrographic surveying agencies can support or defray much of the cost of the survey.

Horizontal positioning is a significant concern irrespective of the surveying application. The use of RTK/PPK positioning for vertical correctors generally resolves this issue adequately, however, and no other constraints are required. If RTK/PPK positioning is not possible, however, then careful use of a Wide-Area Augmentation Service (WAAS) or Globally-Corrected GPS (such as StarFix, C-Nav, or GIPSY) would typically be required to achieve the sort of horizontal uncertainty required to adequately merge multiple passes into a coherent object structure.

Cleanliness of data is important in all surveying applications. In general hydrographic surveying, however, it may be considered acceptable to allow some beams of the swath to be intentionally noisy in order to optimize some other property of the system. A typical example is to artificially decrease the maximum range of the SONAR, sacrificing some of the outer beams, in order to increase the achievable ping rate and therefore increase the along-track data density over the center of the swath. Where the target of interest is a mostly smooth seafloor, this is generally acceptable since it is relatively straightforward to identify the anomalous soundings (Figure 6(a)). In the context of an object with large vertical extent and very complex structure, however, this can be significantly more difficult (Figure 6(b)), leading to a very difficult subjective decision making process (see Section 3.3). Given that the observation time for even large objects is very small (on the order of a few minutes), but the processing time for dealing with noise over such objects can be very large (on the order of an hour), if higher data density is required it makes economic sense to make multiple passes over the object rather than attempt to artificially increase the density of one pass. As long as sufficient vertical and horizontal positioning control is available (as outlined previously) then the data should still be able to be merged in this case into a composite object. Archaeological surveys therefore require even more attention to dynamic operational tuning of the MBES in order to ensure minimal outlier 'noise'.

Although the bathymetric measurements made by the MBES are the primary goal of the survey, other observational data is often available. Most system provide some measure of the backscattered energy from the seafloor, and modern systems often now provide the ability to capture acoustic backscatter data from ping transmission continuously (i.e., to image the watercolumn in addition to the seafloor). These measurements can be used to develop new data products that illustrate features of the sites that are not otherwise visible in the bathymetric data (see Section 5, for example). However, in many survey suites they are not routinely monitored during acquisition, and may not be fully preserved into the data set archived for post-processing. To a certain extent, choices made to optimize the data for bathymetric quality as outlined above run counter to the requirements for backscatter quality. Adjusting the transmit power in order to achieve reliable detection on the outer beams, for example, can cause the backscatter to be saturated, while rapidly changing the power and/or gain can result in backscatter artifacts that are difficult to recover later. Recognizing that bathymetric information is typically the most important, archaeological surveys should at least monitor the backscatter being developed during the survey, and should ensure that all relevant data and metadata are being archived. Frequently, this last is difficult to achieve without testing a component of the data being captured. It is therefore essential that all required processing tools are available in the field, and that the data is examined immediately after capture to ensure completeness.

3.3 Processing strategies for wreck data

As with data collection, the processing strategies for wreck data typically follow the protocols used in modern hydrographic processing schemes. This typically involves a data flow path where raw data is transformed quickly into a surface representation, frequently with auxiliary data layers such as standard deviation, data density or uncertainty among others, which is then used to guide the effort of removing the data observations that are not consistent with the hydrographer's interpretation of the configuration of the seafloor. In some instances, automatic or semi-automatic methods of processing are used to construct the surfaces; in other cases simple distance weighting is used. After remediation of the inconsistent data by manual or semi-automatic means, the data is summarized either as a surface or as a collection of 'raw sounding' observations as dictated by the hydrographic agency contracting the data.

For archaeological survey, these methods are sufficient for areas where general reconnaissance survey is being undertaken, prior to any detailed investigation of an object. We have found that remediation utilizing a three-dimensional representation of data points in an area-based editor is by far the most effective method for dealing with wreck data, primarily because it allows for the detailed visualization of

the vertical structure of the wrecks in context with the dubious observations. For investigation of particular objects, it is often the case that the most complex part of the editing task is deciding whether a particular sounding is erroneous, or simply a very small part of the wreck: wrecks often have small pieces of superstructure or hull only marginally attached to the main wreckage, depending on the state of decay. In this instance, we have found it effective to use small 'subset' slices through the wreck's hull aligned with the primary longitudinal axis of the wreck. This helps to establish 'inside' from 'outside' of the wreck, and delineates the primary hull or superstructure more clearly. As with all subset editing applications, it is essential to ensure that sufficient context remains that objects are not removed as erroneous because only the part within the slice is considered. Having an 'overview' of the data in 2D or 3D (which is preferred) helps to ensure that the detailed view is not too selective.

In hydrographic data processing, the most common goal is to retain the shoalest point in any data set, which preserves the observation most significant to surface navigation. Some leniency in removing deeper coincident points is therefore natural. In archaeological data processing, however, deeper points could represent sub-decks that are exposed by gaps in the upper plating and imaged by some other beam. It is essential, therefore, to be especially careful when reviewing the data within the wreck structure to ensure that significant detail is not removed. This is not always trivial. Good practice for this is to prioritize time for investigation of the detailed structure of the wreck within the processing stream, to visualize the structure appropriately as outlined above, and to build a mental picture of the structure 'through the noise' before starting on the editing task. We have found that this last is usually possible due to the human ability to see structure clearly past visual distracter points, and greatly aids in decisions as to what to keep.

Visualization of the structure of the wrecks is essential in making suitable decisions for editing. In particular, there are significant limitations to DTMs of a wreck to represent many of the overhanging objects that wrecks frequently manifest, such as lifeboat davits or masts, Figure 7 (see Section 3.4). We have found that the level of detail of the bathymetric data generated from moderns surveys is more than sufficient to allow details of objects to be identified from pictures of the ships before they were sunk, if available, and editing with respect to the known structure of the wreck prior to sinking can be extremely beneficial in deciding what is likely noise, and what significant archaeological detail.

3.4 Presentation of wreck or object data

The essence, and primary difficulty, of visualization of the data from archaeological surveys is in balancing efficiency and easy of interaction against scientific veracity and interpretability. The two primary modes of display are as either a point cloud (i.e., a representation of direct observations with one glyph for each) or as some form of surface model, whether DTM or TIN. The correct answer is properly a combination of the two types of visualization depending on the goals and users.

The simplest form of visualization for spatial data is as a surface model. Simple to compute and fast to render, they readily provide geo-spatial context for the site, and are easy to interact with using common tools. There are some limitations, however, the most important of which is the implicit assumption that there is a continuous surface to model. In the case of general bathymetry, this is acceptable: in most cases there is a continuous seafloor, and if there is an overhang the primary interest is in the shallowest part of it. Even with wrecks, in the hydrographic case, the shallowest part of the wreck at each location of interest can be summarized by the shoalest point, or shoalest probable depth according to preference. For an archaeological investigation, however, we would like to preserve interior inclusions and overhangs as inherent properties of the object, which cannot be done with a simple surface model of the type commonly used, Figure 8. More complex models are certainly possible, such as full CAD models with photo-realistic rendering, but the time required to generate such models, the costs of the associated software, and the limits on interaction with the data due to rendering delays make them unlikely as adjuncts to rapid scientific investigation of sites. These limitations suggest that while a surface model may not be the final source for detailed investigation of the interior of an object (or at least such of the interior as can be seen), it does have a role to play to provide for rapid interaction with data and site context in the larger sense.

In order to generate effective surface models, however, some other potential limitations have to be considered. The primary concern is one of resolution: at what level of detail should the object be constructed? To a certain degree there is a free hand with this choice, since we acknowledge that the surface object is primarily for large scale context, so preservation of small detail does not overly concern us. (We note in passing that algorithms based on a simple mean of all points in a neighborhood will almost always result in poor renderings since any interior inclusions will cause 'pits' in the data where the upper

surface is actually intact, changing entirely any estimate of the state of the wreck. A simple shoal biased surface resolves this issue, although it can cause others due to unedited outliers.)   However, if the resolution is set inappropriately low, then significant details of the object will be lost, and along with them the visual cues that would incite further investigation of the object. In Figure 9, for example, the detail of the frontal hull collapse of the SMS BRUMMER would result in significantly different interpretation in the version constructed at 1.0m resolution; possibly sufficiently that the area might not receive detailed investigation as would be the case from visualization of the 0.5m version. Both, of course, fail to show the detail in this area of the point cloud version of the same area from approximately the same vantage point, but it is interesting to note that while both surface representations fail to properly highlight the mid-line hull collapse points clearly visible in the point cloud, which are significant to the understanding of the hull stability (and hence likely longevity) of the wreck, they do make it significantly easier to see the linear structures on the seabed off the bow of the wreck, which are likely the remains of the foremast. Complementary roles of surfaces and point clouds are clearly indicated.

Points resolve many of the issues of surfaces outlined above, but come with a significant performance penalty. Counter-intuitively, although each sounding point is much simpler than a surface, to allow each point to be illuminated it has to be drawn in most cases as a geometric cube rather than a simple point, requiring that all six faces be lit and rendered; the combination of a few hundred thousand of these can give even modern graphics cards sufficient work to slow the achievable frame rate below that which is acceptable for an interactive data manipulation experience. This may be reduced to some degree by suitable visualization techniques such as only drawing a subset of all points while interacting with the data, and then redrawing the rest when the viewpoint stops moving, or by careful selection of which cubes to draw and optimization of the drawing primitives. At some point, however, interactivity of point clouds becomes limited, making them suitable for small-scale detailed investigation of parts of an object, or pre-rendered non-interactive views of large objects (e.g., as a video showing particular parts of the data).

For points, the principal visualization controls are point size and color. Point size is mostly a matter of taste, although too large a point size results in merging of structures, and can obscure details in the data (Figure 10(a)-(b)); some interactive choice is appropriate since this varies with object. Coloring of points depends on the application. In publicity work for outreach applications, monochromatic rendering (Figure 10(c)) can be very atmospheric, but does not convey the scientific information that can be color-coded onto the points such as depth, Figure 10(d). Technically, the depths of the observations should be evident from the size of the points in the projected visualization space. It is common to have difficulty in determining this in close points that differ in depth significantly but which are rendered close together due to the perspective of the interactive viewpoint. One promising technique for providing a halfway-house between a fully rendered model and a point model is the use of oriented facets, Figure 11. Here, normals are computed for each sounding based on the vector mean of the normals for the triangular facets between the sounding and its immediate neighbors (i.e., immediately adjacent beams within the same ping, and the same beams in the previous and next pings). The soundings are then rendered as small quadrilateral patches, colored by depth and oriented with respect to the mean normal. The effect is to generate a pseudo-surface which is readily lit and renders quickly, and which can also be used to occlude soundings which are 'behind' the nearest surface to some extent (simply by culling soundings with eye-point angles greater than 90°). Since these are still soundings, however, they can be time-tagged and therefore combined in a 4D sense with other data, such as the watercolumn and ship trajectory information shown here (c.f. Figure 16 and Section 5).

Visualization difficulties with points can be reduced considerably by the use of animation. Relative motion of near and far points against each other (motion parallax) triggers strong depth cues in the brain, allowing the viewer to build a mental impression of three dimensional shape of the wreck that allows for more complete understanding of the structure even when the animation stops. So strong is the effect that it is even effective to 'shake' the viewpoint slightly around a nominal view vector in order to understand the local structure being examined. Similar effects can be had with pre-programmed animation sequences summarized in a video clip, but our experiments with these show that they are not as effective as an interactive experience. Heuristically, we believe this is probably because the pre-planned flight path does not allow the user to focus on the data that interests them, or that which they find particularly confusing. Since these effects are different for all users, a pre-planned sequence is less than ideal. Even if the data are lower resolution during the interaction, or are not rendered as well as they might possibly be in a non-real time method, information transfer about the shape of the object is higher when interactive.

These observations highlight the differences between data products generated for scientific use, and those generated for communication to the public (a very important mission in most archaeological surveys). While the latter are required to be visually compelling, but static; the former are required to be scientifically informative, quantitative and interactive. These requirements dictate the constraints of visualization. The monochrome representation of Figure 10(c), for example, does not carry the information inherent in Figure 10(d), but it does have more visual impact. Our experiments suggest that color differences are less important for static visualization, and may even be deleterious to the 'solid surface' illusion in point clouds that make the objects much more 'solid' in a static view. For scientific work, color-coding and the ability to rapidly change coloring is essential. In both cases, animation is essential, although in the case of scientific work easier to arrange.

4 Site Monitoring through Repeat Surveys

The primary aim of the ScapaMAP II survey in 2006 was to answer the fundamental question of whether there is sufficient repeatable detail in the remotely sensed description of a wreck or other object that the data could be used for monitoring of the object over time. This might be considered to be an obvious assertion: the remote sensed data are extremely compelling because they are instantly recognizable as wrecks, but a great deal of this information is filled in by the human observer, rather than being inherent in the data. It was not initially clear that the repeatability of survey would be sufficient for scientific investigation of the sites, and therefore monitoring.

On the macroscopic scale, sites of any size above approximately 1m^2 can trivially be recovered using conventional navigational equipment, and in that limited sense, the survey is repeatable. Direct comparison of data at the macroscopic scale, Figure 12, clearly shows that similar features are seen in the two surveys which are five years apart in this case. In many instances, even small details can be compared side by side; in Figure 12, the breaks in the hull plates along the central longitudinal axis of the SMS BRUMMER are evident in both cases, an example where a detail that could be dismissed as a data anomaly in one survey is confirmed by the second. (Hull plate breaches are very significant in estimating the structural integrity of a wreck, and hence the level of preservation and likely decay rate, since they form a weak-point which can concentrate decay, leading to a cascade of collapse.)

The results of the survey also show that detailed descriptions of the wrecks may be formed from pairs of surveys compared side by side. In the detailed view of the bow of the SMS BRUMMER in Figure 12, for example, it is possible to estimate the progress of the hull plate and subdeck collapse into the body of the hull, and towards the superstructure further aft. Repeated sufficiently often, such surveys would allow estimates of the rate of collapse as well as continuous monitoring of the current configuration of the wreck. Care is required in interpretation of the results of the repeated surveys, however. Increases in MBES performance generally result in higher data density, which is evident in Figure 13 in the short period of time between 2001 and 2006. Elements of the wreck which seem to "appear" in subsequent surveys should be treated with suspicion, therefore, since they may just represent details which were not evident in the prior survey. In addition, slight changes in survey platform trajectory can result in very different shadowing patterns on objects with significant vertical extent (such as wrecks), which might be (over-) interpreted as important differences. (Both of these effects are evident in Figure 13(b) where the data from 2006 shows better definition of, e.g., the thin longitudinal remains of the sub-deck support members, which therefore anomalously "appear" in the later survey, but is also missing the main deck bulkheads between the forward gun director and the main superstructure, most likely due to shadowing rather than actual collapse.)

Methods for more quantitative estimates of the differences between two surveys are quite limited. The simplest is to compute the difference between two coincident surfaces constructed from the separate surveys, although this inherits all of the problems of dealing with surfaces as outlined previously and in general is only useful for gross differences. Directly computing a difference of two point sets is not well defined since there is no direct definition of interior and exterior spaces in the data. Experiments in using stereo rendering of pairs of surveys show that it is difficult to 'fuse' two images that are not from the same survey, since small changes in shadowing (e.g., due to trajectory), data density and swath orientation significant affect the rendered images. In our experiments, the viewer fixated one or the other survey, rather than seeing any differences between them. More complex visualization techniques might render better quantitative methods for comparison, but it is likely that this problem is formally as difficult (and

closely related) to the problem of forming the semi-random points into a coherent 3D surface described previously: something that is not readily soluble by automatic methods.

The primary limitations in doing any comparisons are the achievable horizontal and vertical positioning accuracy, both of which limit the degree of fusion that is possible with pairs of surveys; small changes will be completely masked if there is a systematic bias in either dimension such as might be formed through variation in GPS satellite constellations or differing vertical datums. The positioning requirements for this type of repeat survey are in general even stricter than they are for standard hydrographic surveys because here the measure of success is the relative uncertainty rather than the (typically more generous) absolute uncertainty that is commonly required. That is, in a hydrographic survey, we might be content to position a shoal to within 5m (2drms) repeatably, allowing an adequate margin by which satellite constellation changes may be accommodated. In an archaeological survey, however, a constellation change that resulted in a horizontal offset of 2.5m in a random direction would result in a complete inability to compare the results reliably. For similar reasons, care in correction for motion of the platform is paramount, since small residual motion artifacts (e.g., induced heave on the centimeter to decimeter scale) can result in obscuration of significant features. To a certain extent the wider application of RTK or PPK GPS measurements and tightly coupled motion sensors and features such as delayed heave estimates will result in these difficulties being of lesser significance in the future. The requirements for the survey that they imply will remain, however.

5 Survey with New Data Products

Backscatter data, typically from a towed sidescan system, has been a staple of wreck investigation for many years. There is a difference in the task of detecting an object, for which this is ideal, and some careful investigation of its morphology, however: in the case of very large wrecks, the prevalence of shadowing is so high that much detail of the hull shape is often obscured, Figure 14. In addition, the lack of bathymetric data results in difficulties in correcting this data for positioning which further distorts the morphological indicators that are useful for monitoring the object's condition.

Addition of bathymetric information does not significantly improve behavior of MBES backscatter over large wrecks, however; the fundamental limitation of shadowing is still present, and attempts to geocode backscatter coherently over a large wreck are unconvincing.

In our experiments, the most compelling use of backscatter is in its conventional survey mode; that is, to map changes in sedimentation around the areas of interest, or significant objects. In Figure 15, for example, we show a section of Gutter Sound, Orkney where the backscatter shows a significantly anomalous return for which there is no bathymetric explanation (even when the backscatter is heavily vertically exaggerated and strongly shaded as here so that centimetric artifacts become obvious). It is unknown whether this anomaly represents a difference in seabed surficial sediment, or whether it is a result of material left on the seabed due to a pollution event at some time in the past. However, the ability to map these changes, especially since this data is essentially 'free' from the bathymetric survey, can help in the rapid-response and baseline mapping segments of archaeological site investigation.

A newer capability in MBES systems is to capture data for the entire water column, rather than just the bathymetry and seabed backscatter. Potentially, this allows for greater detail of investigation of the data in a post-processing mode, since the data is not reduced to just the seabed backscatter, or one data point per beam, Figure 16. Comparison of the water column data to the detected bathymetric data is also useful, particularly where small features (e.g., masts, lifeboat davits or guns) are present, since the water column data may show more detail, or finer detail, than can be resolved even in high resolution bathymetry. There is a significant cost in this type of analysis, however, since the data volumes for this type of data can be very high (on the order of several gigabytes per minute). The cost of storing, processing and presenting this data will mean that for the immediate future only very small sections of data are likely to be collected and processed, implying that very careful targeting of this type of data is required. For archaeological surveys, therefore, this is most likely to mean targeted passes on already well established objects, rather than use in rapid reconnaissance.

Methods for processing and display of water column data are still very much in their infancy. Future techniques might include volume rendering of semi-transparent displays, multiple object detection and tracking within the water column, and volumetric reconstruction of data. It remains to be seen, however, whether this can be done in anything like real-time, and what the computational cost will be.

6 The Role of Survey in Site Management Plans

The reality of most marine archaeology is that there will rarely be sufficient funding for an investigation of the site to the level of detail that is really required to fully document the state of preservation. This is not to say that funding is not available; for some sites, such as the wreck of the MONITOR, large, well-funded projects have been launched. For every site of this kind, however, there are maybe hundreds more with just as significant a cultural impact, which are unlikely to receive serious study. The cost of any marine investigation, and the difficulties of mobilizing human observers to the location means that there is a necessity for methods to rapidly, but perhaps approximately, gather information on a site.

Survey technologies fill this niche. Remote sensing of bathymetry has been shown, here and elsewhere, to produce compelling descriptions of the 3D structure of archaeological sites, and can be done in significantly less time than would be required for a diver-led investigation of the site. In addition, the precisely geo-referenced data can act as a means to prioritize use of more expensive or limited resources, such as diver time. For example, in the case of the SMS BRUMMER, Figure 13, knowledge that the front deck plates have changed significantly would be cause to vector divers to the site for more detailed investigation *in situ*. Repeatability of survey methods is now sufficiently good that repeat surveys make sense, and provide an enticingly economical method to monitor particular objects as a function of time.

There are therefore two primary roles for high resolution survey in a marine archaeological context: rapid pre-investigation surveys to establish a baseline efficiently, and continuous monitoring of the site to establish decay rates or the effects of intermittent external drivers, e.g., hurricane damage. The arguments for both of these are essentially economic, although efficiency of time in observation and safety of the observers are also factors.

Clearly, however, remote sensing is not a complete solution for marine archaeology; there remains no substitute for a human observer for the finest scale observations. This may, of course, be mediated through a remote technology such as a Remotely Operated Vehicle with cameras to do the observing, but the purposeful investigation of a trained observer is still necessary. In the future, more exotic technologies such as AUVs might extend the limits of what is now possible, but no matter how close the remote platform is brought to the site, and the frequency of the observing system, there are fundamental limits to the resolution and precision of the data that can be achieved. Like all tools, therefore, survey systems are only a part of a coordinated site management plan.

7 Summary

Our experience with ScapaMAP shows that remote survey provides a rapid alternative to more conventional marine archaeological investigations, and that recovery of targets is sufficient to allow effective repeat surveys to be carried out without extreme measures. The requirements for marine archaeological surveys are basically those for hydrographic survey, although we are frequently more interested in relative, rather than absolute, error, and therefore we need to pay more attention to factors such as offsets between components of the survey system, or horizontal positioning uncertainty, rather than things like water levels or other vertical correctors. Careful planning of survey lines is also required to ensure efficient and effective imaging of objects with large vertical extent.

Processing methods for marine archaeological data are driven primarily by the desire to maximize resolution, rather than preserving shoal points. This is especially difficult where small over-hanging features are observed with little reliable redundancy. Standard hydrographic tools are effective, however, when paired with a sufficiently observant, trained operator.

Visualization and display of objects with large vertical extent and complex morphology is problematic, and depends on the intended goal of the display. Scientific investigation and public outreach demand different approaches in data type, color-coding and lighting among other factors. Our experiments indicate that point-cloud type displays are generally more effective than surface type displays because of the observer's ability to fill in the gaps between the points with an inferred surface, and that interactive displays, or at least animated versions of data, are more effective because they provide much better depth cues for the observer due to the effects of motion parallax, making it easier to interpret the 3D structure of the objects.
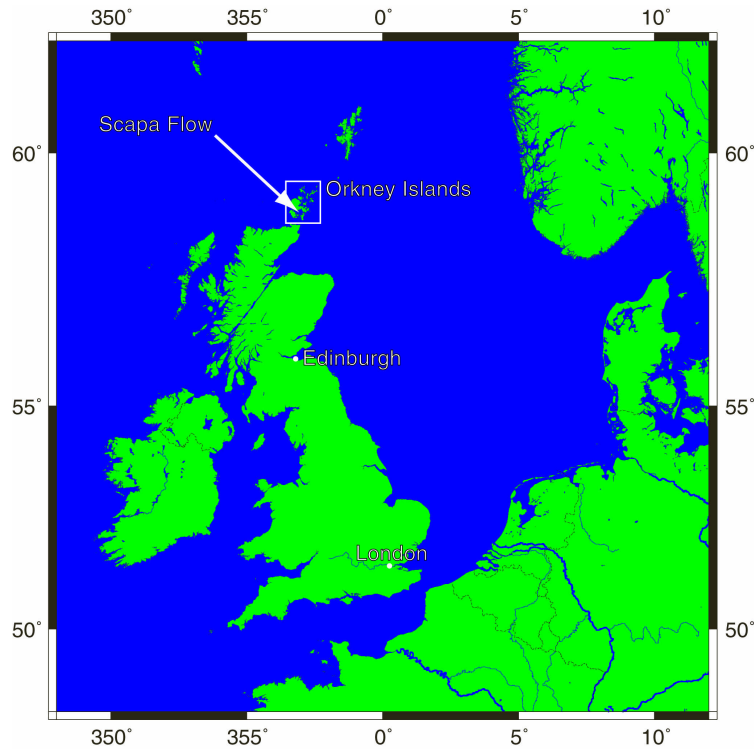
(a) Location of Orkney Islands to the north east of Scotland

(b) The Imperial German fleet at anchor in Scapa Flow in early 1919. View is looking south towards Cava from Swanbister Bay. (Source: Orkney Library Archive)

Figure 1: The British fleet anchorage of Scapa Flow was first used immediately prior to World War I. The intent was to blockade the North Sea from Britain to Norway and thereby stop the Imperial German Navy from commerce raiding in the Atlantic. The Imperial German fleet was interned in Scapa Flow following the Armistice of 11 November 1918 until a negotiated agreement on their fate could be signed.
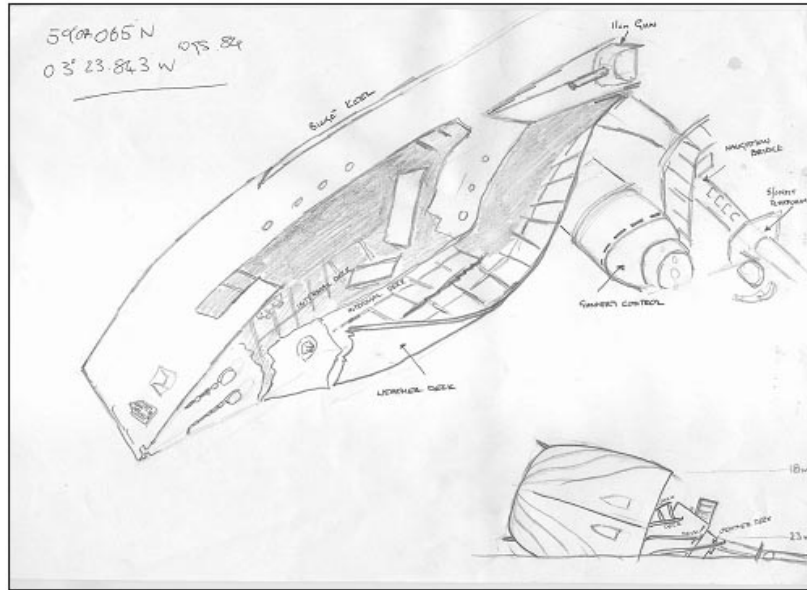
Figure 2: Diver's drawing of the bow of the SMS DRESDEN, based on numerous dives on the wreck in 30 min sections. The level of detail is as good as the diver's memory, but may vary depending on level of experience, visibility level, etc. (Source: Steve Liscoe, ADU)

(a) Small object with low relief, one line.

(b) Large object with low relief, two lines.

(c) Large object with large relief, two lines.
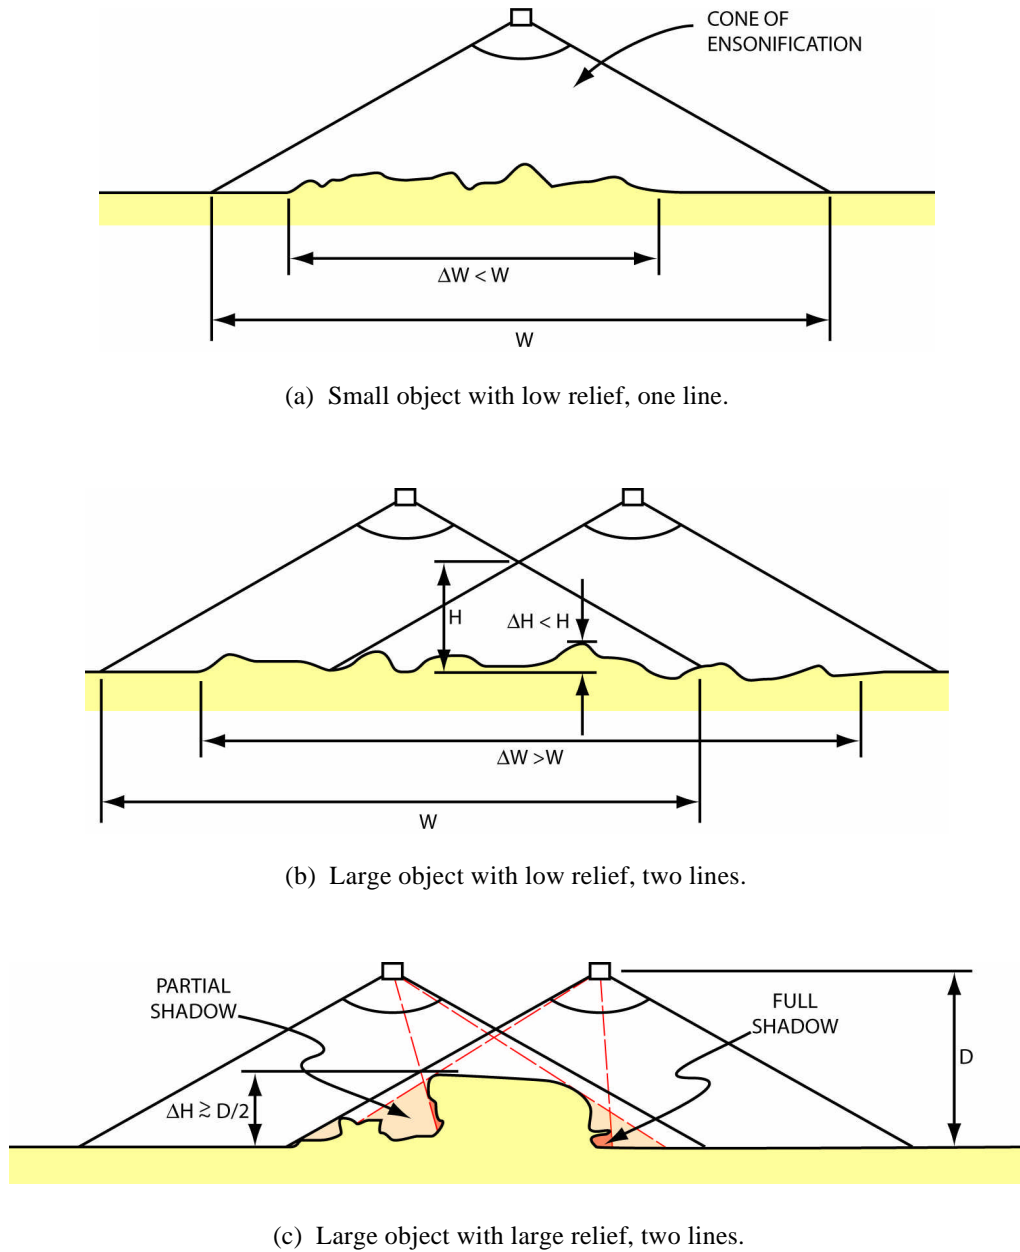
Figure 3: Potential line organization for wreck surveys, which depend on the characteristics of the object being surveyed. For objects with low relief, (a), (b), layout depends on horizontal side; for objects with large relief, (c), layout is done to minimize potential for acoustic shadowing.
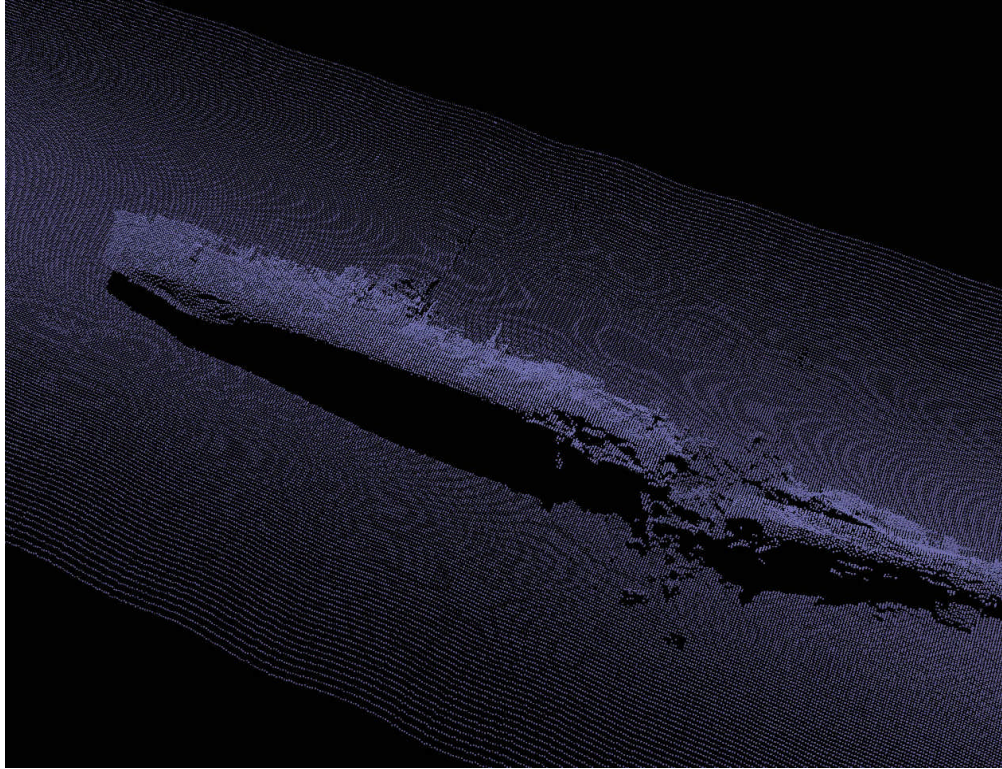
(a) A single pass of the MBES on the SMS KÖLN.
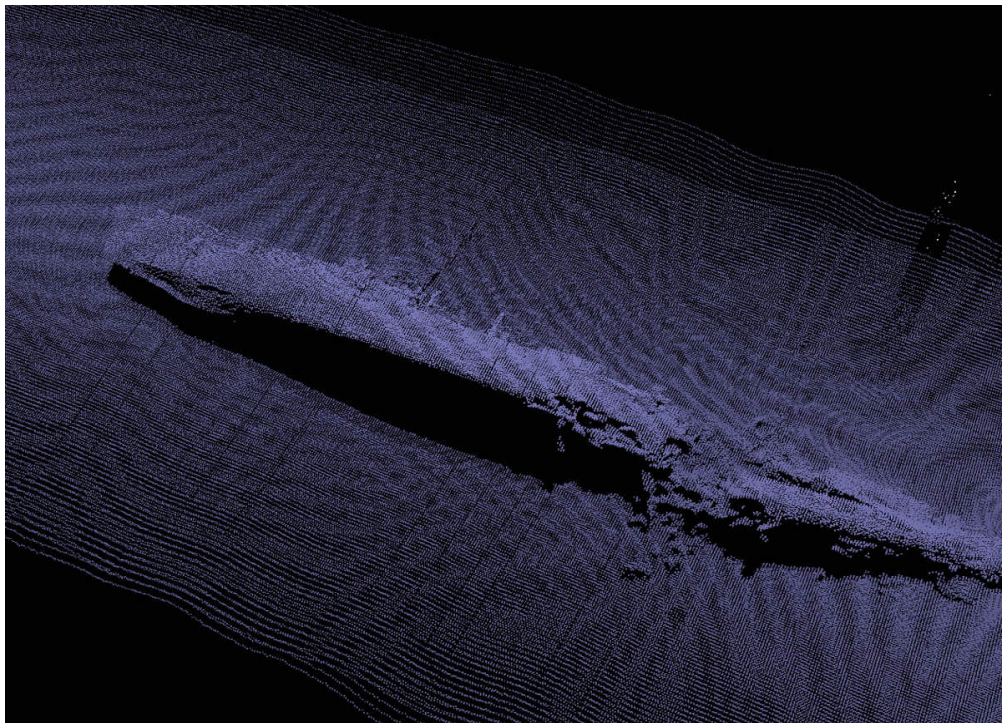
(b) Two passes of the MBES on the SMS KÖLN.

Figure 4: Effects of visualizing more than one MBES pass over a wreck with high detail level and insufficiently accurate positioning. Small inconsistencies in horizontal and vertical positioning result in a 'smearing' of the wreck, (b), relative to a single pass, (a).

(a) Low relief objects, one line.

(b) High relief object, multiple lines with one to fill in the superstructure data if possible
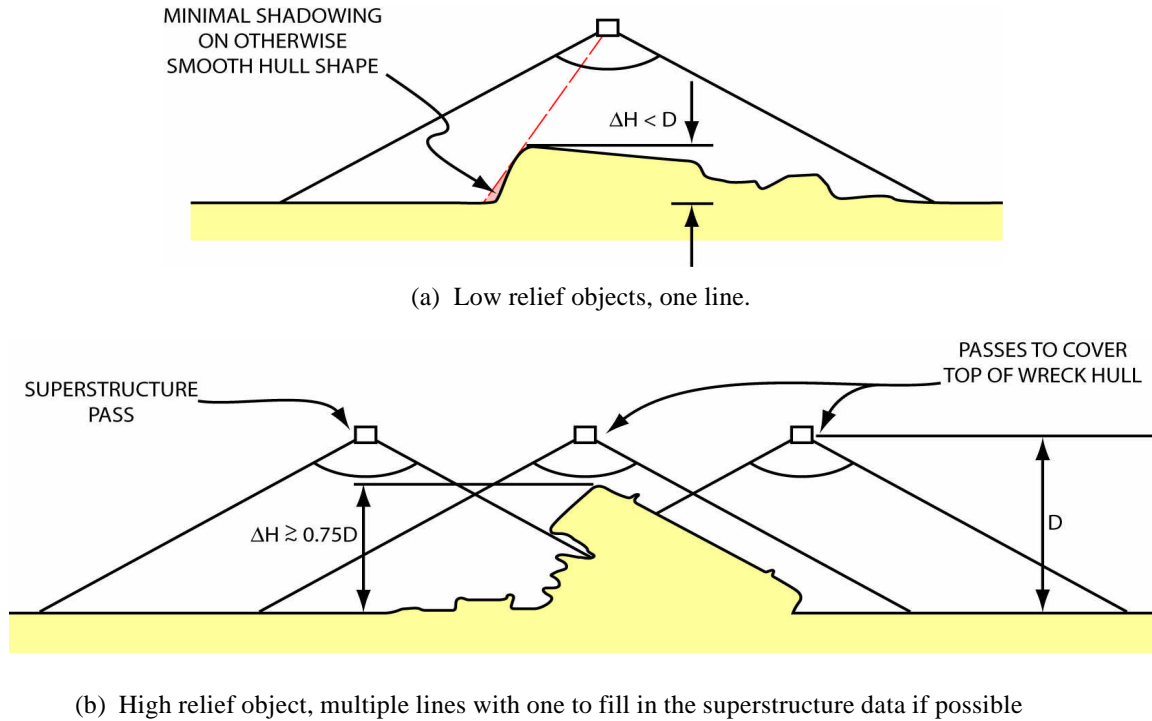
Figure 5: Layouts for 'beauty passes' over objects, intended primarily for public outreach visualization. In the case of high relief, (b), the best view will depend strongly on the configuration of the wreck, and may be of the superstructure but only part of the hull.

(a) Noise readily distinguished from data

(b) Noise too close to object of interest
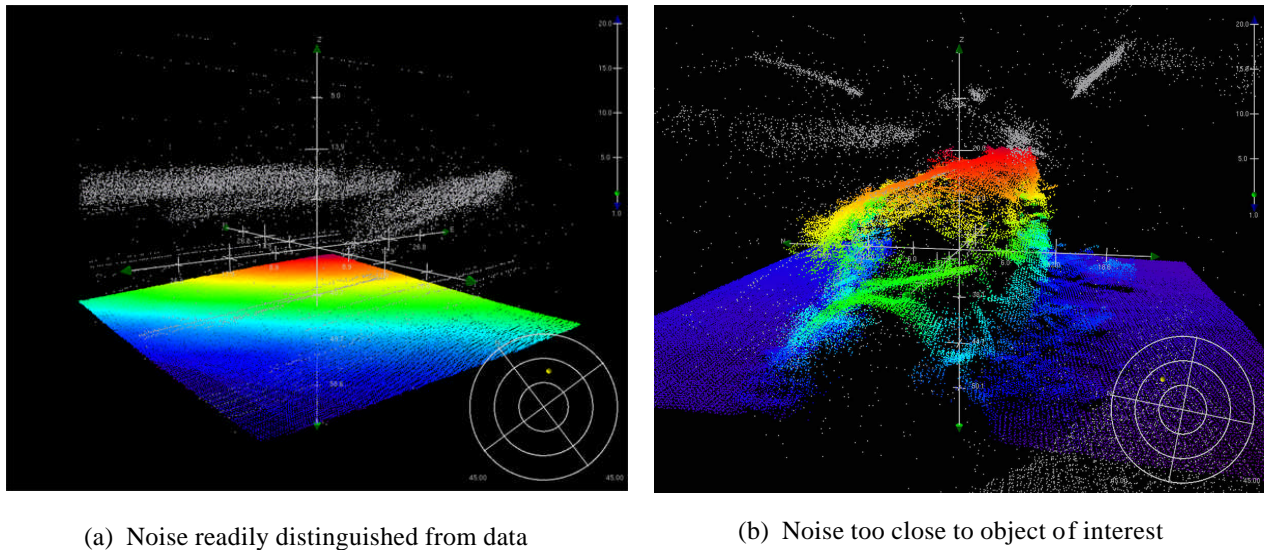
Figure 6: Examples of hydrographically 'acceptable' data noise (a) which is readily distinguished from the true surface data, and archaeologically challenging data noise (b) which is very difficult to subjectively separate from the 'true' structure of the wreck due to the very small nature of the structure components being considered.
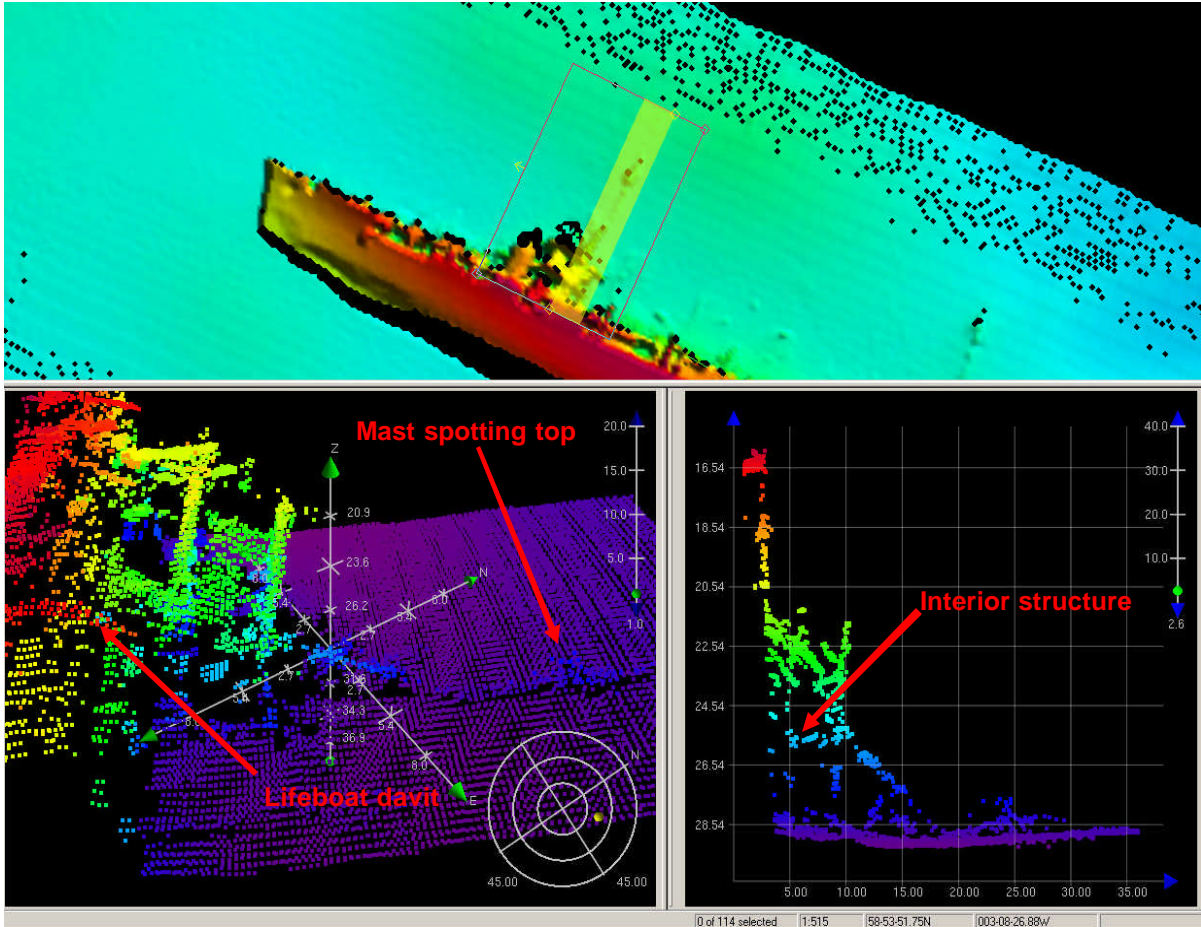
(a) The SMS KÖLN in 2006
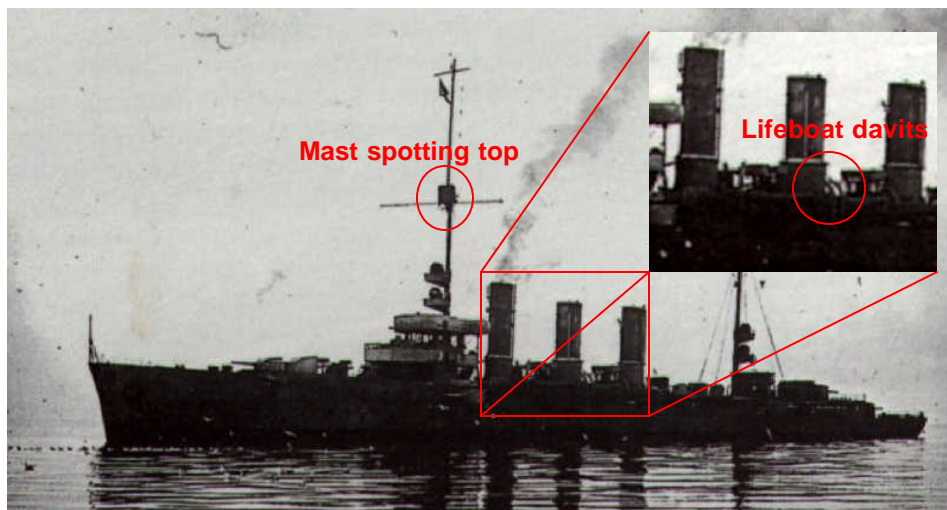
(b) The SMS KÖLN in 1919

Figure 7: The ability of a DTM to represent overhanging features is limited, but their importance can be archaeologically very significant; visualization of these from, e.g., images of the ship before it was sunk can be very important in understanding the visual structure and therefore what to retain in the data during the editing process.
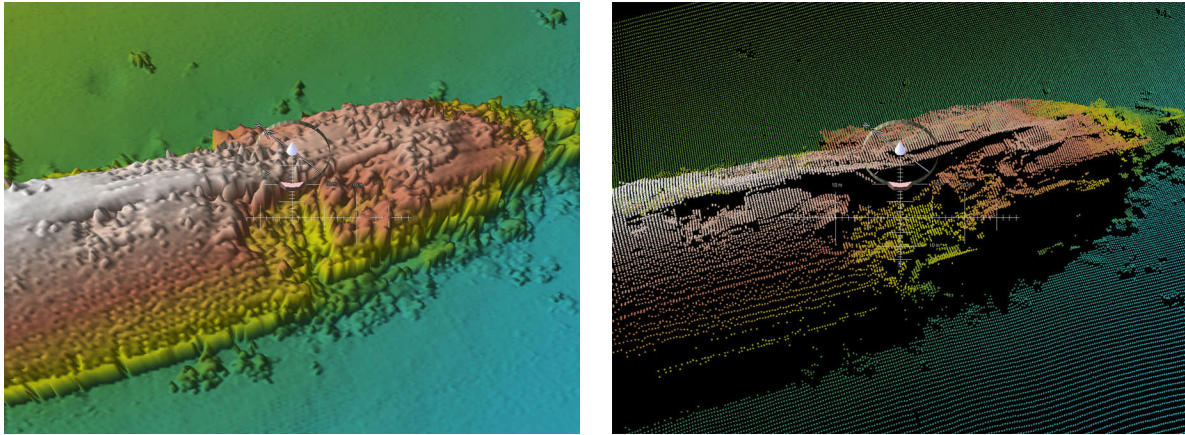
Figure 8: Surface representations of wrecks (here, the SMS KRONPRINZ WILHELM) and other data with significant interior inclusions or overhangs cannot maintain fidelity due to the assumption of a single contiguous surface, which does not map well to this type of data. Point clouds allow this, but are difficult to manipulate.
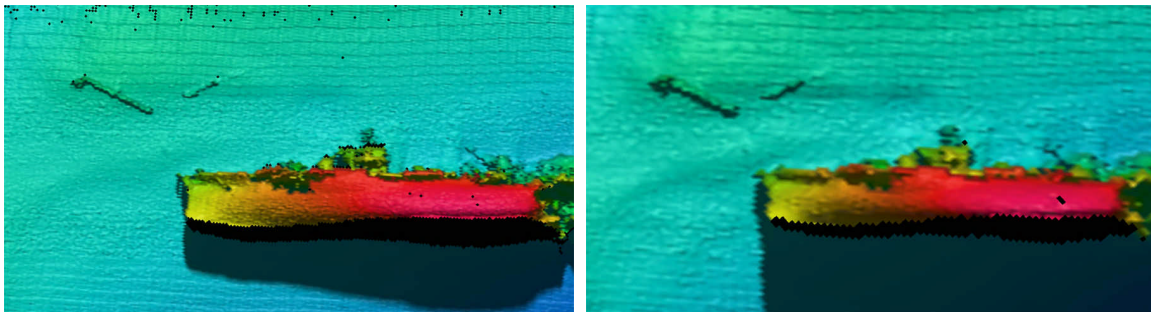
(a) Grid resolution 0.5m

(b) Grid resolution 1.0m
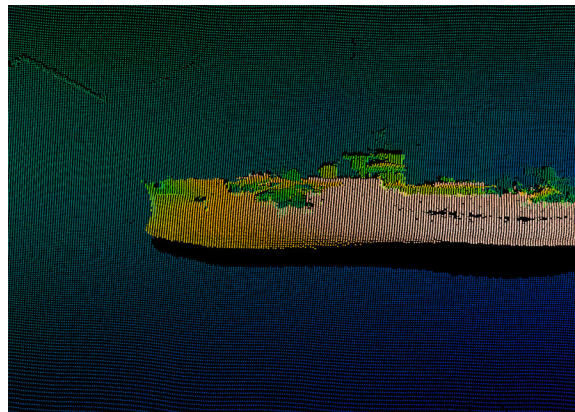
(c) Point cloud

Figure 9: Surface resolution is essential in visual identification of areas of the object that require further investigation in the point cloud model of the data. Grids at 0.5m and 1.0m tell a very different story about the state of preservation of the bow of the SMS BRUMMER, with the 0.5m being closer to the detail found in the point cloud.
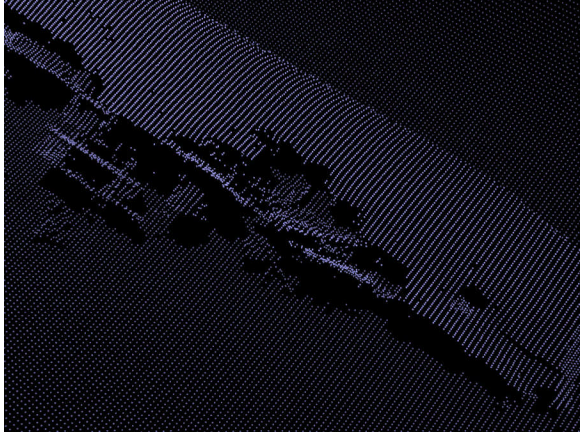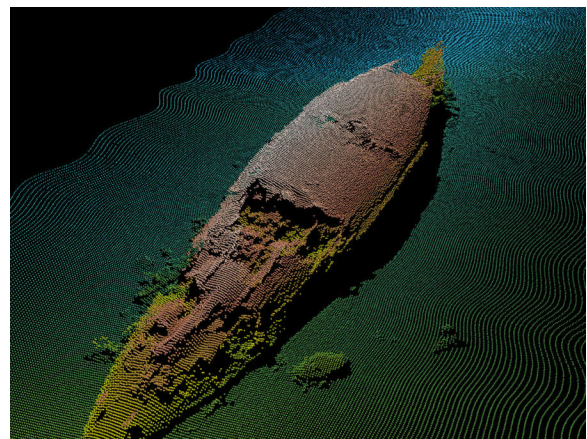
(a) SMS BRUMMER, 0.1m
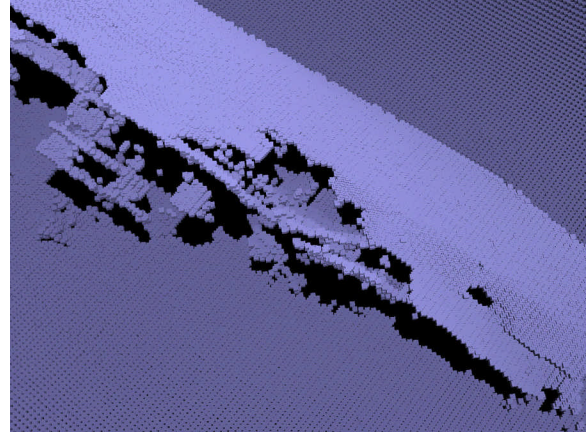
(b) SMS BRUMMER, 0.4m
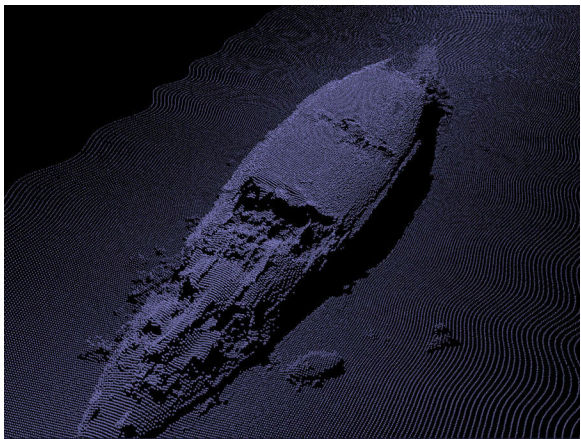
(c) SMS MARKGRAF, duotone

(d) SMS MARKGRAF, colored by depth

Figure 10: Variation of point size and color can results in significant differences in visibility and utility of the data. Increase of point size (a), (b) can improve some rendering, and/or obscure some details of the object. Monochromatic colorings can be very dramatic, (c), but do not contain the scientific data inherent in a color-coded data set, (d), here showing depth.
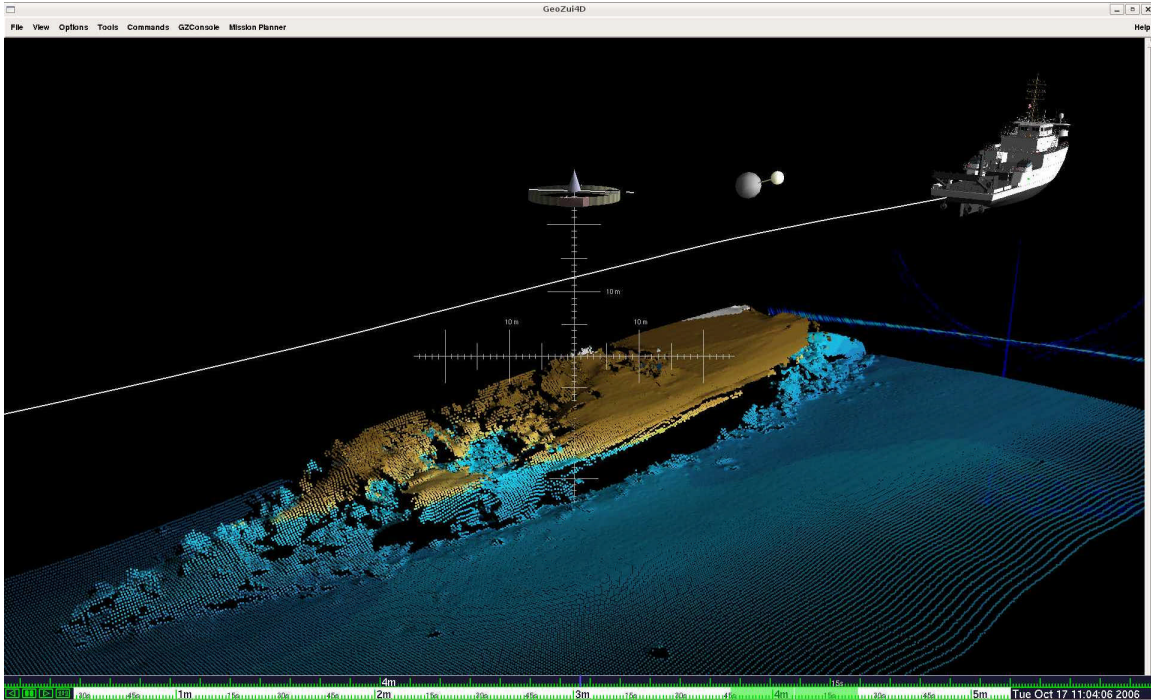
Figure 11: Alternative rendering of point data (here, of the SMS KÖNIG) as oriented facets using the mean normal for triangulated patches between soundings and their immediate nearest neighbors. The data points, color-coded by depth, are represented as small quadrilateral patches which are fast to draw and readily lit and rendered. The data is still inherently point-based, however, and therefore can be integrated with other 4D (i.e., spatio-temporal) data. (Image source: Roland Arsenault, Data Visualization Research Lab, CCOM/JHC).
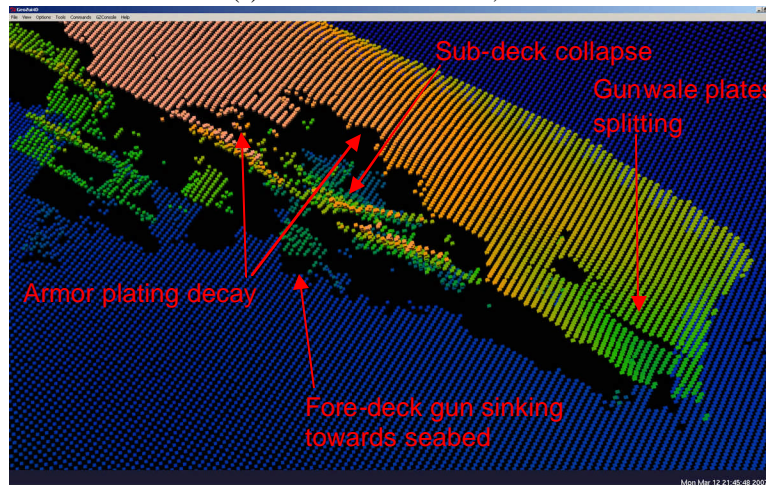

(a) SMS BRUMMER, 2001   (b) SMS BRUMMER, 2006

Figure 12: Overview of the SMS BRUMMER surveys from 2001, (a), and 2006, (b). Although the data densities are different, and the motion compensation from 2001 is not ideal (the MBES was deployed on a pole mount with limited rigidity), the two surveys are directly comparable, and details such as the hull plate failures on the longitudinal centerline and salvage recovery points are maintained between surveys.
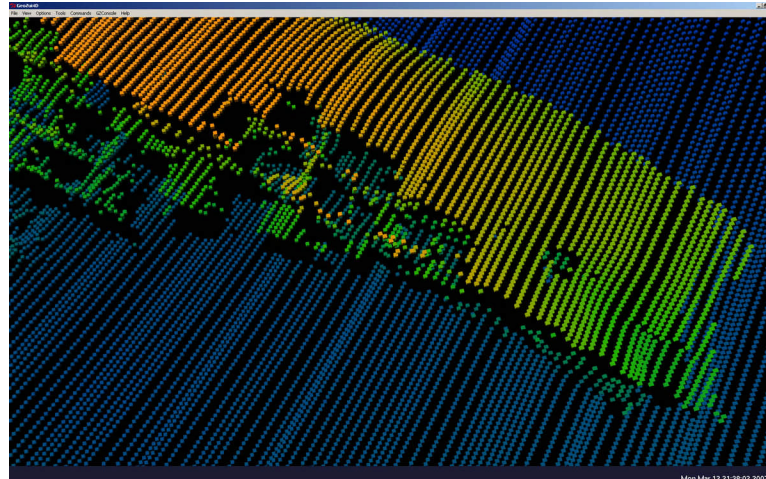
(a) SMS BRUMMER bow, 2001

(b) SMS BRUMMER bow, 2006

Figure 13: Detailed view of the bow of the SMS BRUMMER in 2001, (a), and 2006, (b). The progress of the hull plate and sub-deck collapse towards the keel and superstructure are evident, showing the sequential monitoring of the state of preservation of the wrecks can be supported by remote sensed data of this type.
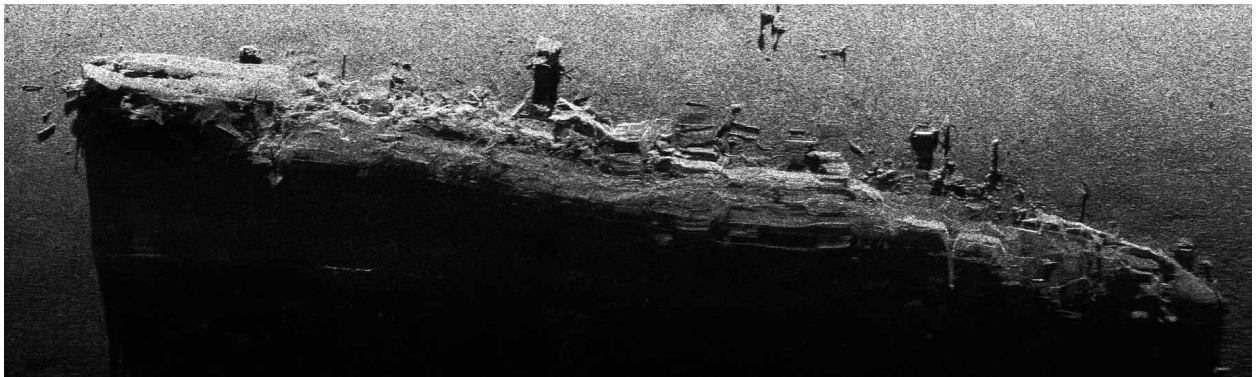

Figure 14: Sidescan imagery of the SMS KRONPRINZ WILHELM, collected with a Klein 2000 in 1999. The vertical extent of the wreck results in significant shadowing, and consequent difficulty in recognizing morphological indicators useful for assessing the state of the wreck. Although ideal for detecting wrecks, towed sidescan is not always ideal for monitoring.

(a) Backscatter

(b) Color-coded, shaded, bathymetry

Figure 15: Multibeam backscatter and bathymetry for the same section of Gutter Sound. The backscatter shows a significant anomaly that has no bathymetric expression, although the data agree on the objects in the top right of the imagery. Use of backscatter to identify sediment variabilities, for example due to differential erosion or pollutant absorption is a useful 'collateral' tool derived from baseline bathymetric surveys.
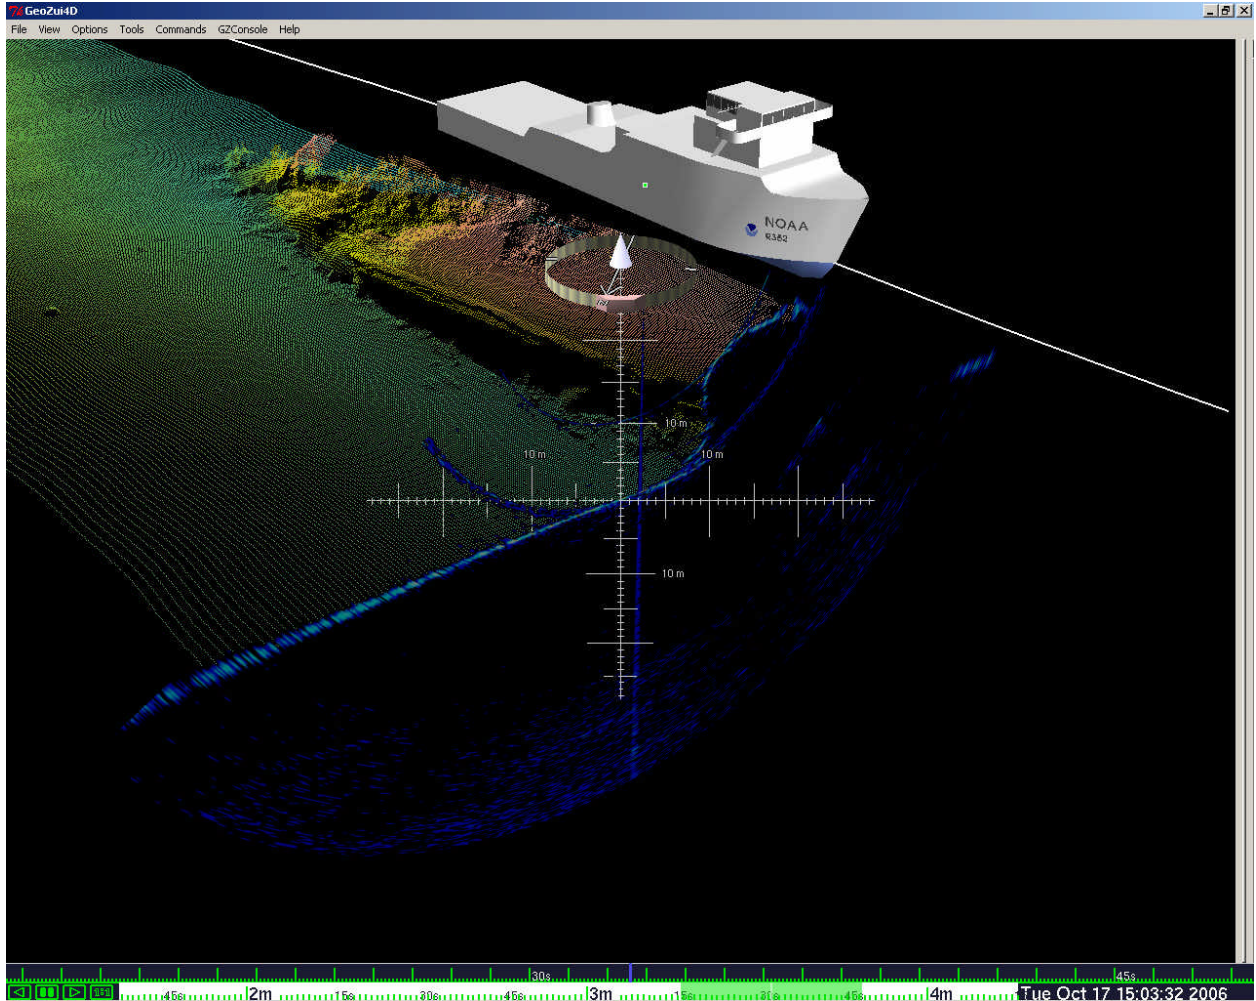
Figure 16: Snapshot of a video sequence showing the bathymetric points and water column data with respect to the survey platform while imaging the SMS KÖNIG. Water column data can reveal more about the structures of a wreck and allows for comparison of detected bathymetric points against observed acoustic data, but volume of data and difficulties in processing mean that this data type needs to be targeted to known objects rather than being applied uniformly across the survey area.