



Deep Learning

Applications for Multibeam Survey **By Burns Foster**

As in many industries, the last few decades of hydrographic surveying has seen exponential leaps in technology. Multibeam and phase discrimination swath sonar systems have evolved to the point where they collect several thousand samples every second, providing unprecedented coverage and quality in surveying the world's oceans and inland waters.

With this increase in collected data volume comes a corresponding increase in the time required to post-process this data. A recent trend has focused on automation for the entire processing chain, involving immediate geo-referencing and basic quality filters in real-time survey, and the creation and live update of downstream products, e.g., elevation models. However, a significant

data processing challenge still exists; despite noteworthy advances in technology for both the acquisition hardware and processing software, the marine environment is noisy. Any acoustic remote sensing system will, to some degree, be susceptible to interference from the environment, particularly for an inexperienced operator. With the monotonous data processing task handled mostly through automation, the last remaining significant bottleneck is manual review of the acquired data: to locate, identify, and appropriately track any features of interest, and also to remove any noise present in the data. Depending on the environment, sensor quality, and other factors, identifying and removing noise from the processed depth values can be by far the most lengthy and arduous task for the user.

Given the physical characteristics of acoustic remote sensing, most systems follow similar designs and operating characteristics. As an extension to this, any noise in the data tends to follow similar patterns, not only within a given hardware vendor's product line, but across vendors as well. Fortunately, research dedicated to locating patterns in massive datasets is something that has just recently become a very large, dynamic industry: Deep Learning. With this knowledge of patterns of noise in any multibeam sonar dataset, it should be possible to apply Deep Learning techniques to automatically locate, classify, and remove noise without operator intervention.

Deep Learning has very recently disrupted many aspects of our lives, including how we shop, how we travel, what news we read, and more. Much of this ubiquity can be attributed to the resources available to the average user, as many of the libraries which have been developed by companies like Google and Facebook are freely available online for anyone with an Internet connection and the willingness to learn them. These same companies often provide access to their computing centres for end users for a reasonable fee, allowing nearly anyone access to incredible computing power for processing immense datasets in a short period of time. These offerings provide an incredible variety of tools and techniques to solve a variety of problems, to the point that perhaps the most significant challenge is not the data processing itself, but choosing which approach to apply to your particular data problem.

Things become even more challenging when dealing with 3D point clouds. There are plenty of techniques, and even entire libraries built, for processing 2D raster images, whether they be colour images, elevation models, or any

other type of raster information. Point clouds represent a much tougher challenge, due to the extra dimension involved (height), as well as the irregular nature of the data. There has recently been some focus on this data type; for example, several groups of researchers working on data provided by small, portable LiDAR sensors for indoor scenes, small outdoor scenes, or even real-time navigation from a moving platform. Generally speaking, this research has been directed toward object detection and classification from these points, e.g., find all pedestrians, or a wall versus a chair, etc. However, these relatively new algorithms have a few drawbacks. They generally all involve fairly large, complex algorithms and networks, which require significant processing times on fairly high-end graphics cards. In addition, they are generally focused on small, local-area scenes from a stationary laser scanner, or from a single pass of a mobile scanner. Looking at a typical multibeam survey, operators are collecting tens of millions of points per day with a single system, and the data volumes involved would quickly overwhelm most of these approaches.

We find some useful techniques in a separate but related industry with the analysis of magnetic resonance imaging (MRI), where recent research in identifying and classifying various types of cancers from these 3D scans have brought about very robust mechanisms for Deep Learning on three dimensional data using a convolutional neural network. For our (current) application, we are only concerned with one data type: noise. This helps us limit the size and scope of our neural network. In addition, we can apply some simple data processing techniques which can down-sample our datasets, while still preserving enough information from the source data that the network can make an intelligent classification.

With a promising Deep Learning approach in place, a few challenges present themselves. First, as with any Deep Learning algorithm, we need a fairly significant amount of “labelled” data available to train the network. This means we need access to a significant amount of processed multibeam data where the noise points have already been identified as such, either through some existing filter technology or by a human. As a software provider for some of the largest hydrographic data producers in the world, Teledyne CARIS is fortunate to have knowledge and access to vast public datasets which have been “labelled” (i.e., cleaned) by hundreds of data processors over many, many years (although the quality and precision of this labelling, it turns out, can vary quite a bit). The second challenge is the variant nature of acoustic remote sensing; generally speaking, the survey platform is a surface vessel which is incapable of adjusting its own distance relative to the seafloor to maintain a consistent scanning density (not more than once anyway). As the survey platform moves into deeper water, the range to the seafloor increases, and the resolution at which we can sample the seafloor decreases with it. In terms of a geo-referenced point cloud (i.e., X, Y, and Z information), we need to be able to account for this change in scale, either in a pre-processing stage or within the network itself. CARIS has invested several years’ worth of research into a very similar problem, in developing algorithms which automatically determine the ideal sampling resolution over a given region of points. We can leverage this technology directly, such that the network only needs to learn noise characteristics across a fixed, scale-less region of points provided by these algorithms.

It turns out that, among a variety of data providers, in different countries, with varying end users, the definition of “noise” can vary quite a bit. In several cases we encountered, manmade structures like quay walls, bridge

piers, and piles have been removed with the same classification as typical acoustic noise, presumably because these structures also appear above the water and their position is already well-known and recorded. However, in practice, a general-purpose noise classifier should likely preserve these manmade structures and allow an operator to decide what classification they should be assigned (perhaps in the future, they may even receive a separate, non-noise classification). The consequence of this is, even with a previously “labelled” dataset retrieved from these data warehouses, great care must be taken with every piece of data that is provided to the algorithm for learning.

After several months of design, testing, and evaluation, we now have a robust, scalable (both in data volume and depth range) algorithm capable of identifying and labelling the majority of acoustic noise, even in very challenging and complex datasets. One of our employees testing it for the first time remarked that the first result that came back was “spooky.” While it may never be “perfect,” the main intention of this algorithm is to handle a significant majority of the work around removing noise from acoustic data, so that a human operator can focus on the wild, weird, and unusual results that sometimes crop up when surveying the world’s oceans.

Burns Foster is the Product Manager for New Product Initiatives with Teledyne CARIS, focusing on developing new products and services outside of CARIS’ core competencies. He previously spent five years as the Product Manager for HIPS and SIPS, CARIS’ flagship processing software suite.