Embedded Deep Learning for Underwater Acoustics

by Fabio Frazao, Oliver Kirsebom, Bruno Padovese, and Stan Matwin

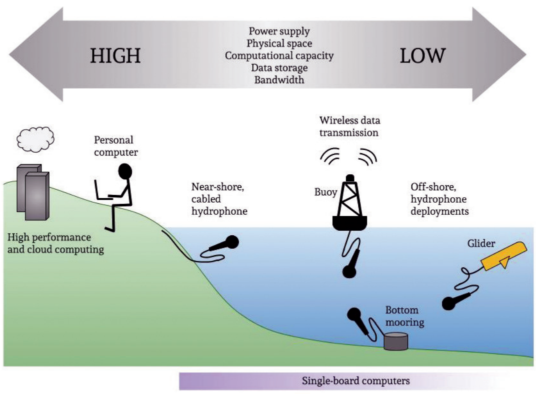

Figure 1: Usage landscape for embedded deep learning in underwater acoustics.

In this contribution, we argue that the advent of low-cost, single-board computers (SBCs) capable of running deep learning algorithms offers great prospects for advancing ocean research. Combining low power requirements with powerful processing units, the new generation of SBCs makes it practical to use deep learning algorithms on platforms like buoys and ocean gliders, e.g., for in-situ analysis of acoustic data. Moreover, the user friendliness of these devices makes them accessible to a broader cohort of ocean researchers.

SBCs are now sufficiently powerful to run standard operating systems and mainstream software, making their use very close to that of a modern personal computer (PC). Moreover, in the last few years, computing modules specifically designed to run deep learning algorithms on dedicated processing units have been introduced to the market, e.g., Jetson Nano and Google Coral. These modules are available on development boards that adopt the familiar interfaces of SBCs such as the very popular Raspberry Pi.

SBCs are small, affordable, and fill a gap between PCs and microcontrollers. A PC requires more power than a buoy or an ocean glider can afford. A microcontroller, on the other hand, uses very little power, but also has limited computational capacity. Devices like the Raspberry Pi and the Google Coral board sit in between and can be deployed in situations where the processing requirements cannot be met by microcontrollers, but there is insufficient power and/or space for a PC.

Deep neural networks are finding increasing use in ocean acoustics, especially for detecting and classifying sounds. A key feature of deep learning models is that they require only minimal feature engineering compared to conventional machine-learning approaches. A deep learning model is trained on the raw acoustic data (or a high-dimensional representation, like a spectrogram), learning the features most conducive to solving a given task. While this enables improved performance, it comes at a cost: the models require large amounts of data to train and the training is computationally expensive and is best performed on dedicated computing infrastructure. However, once trained, a deep learning model requires far fewer resources to run allowing it to be deployed on modern SBCs.

As illustrated in Figure 1, ocean acoustics present a range of use cases for embedded deep learning which vary in terms of severity of constraints. At one extreme, one finds cabled, near-shore hydrophones, which typically are not constrained in terms of power supply. This is in contrast to other types of deployments (buoys, bottom moorings, gliders), which all have limited power budgets, and also impose constraints on data storage and/or transmission during operation. Indeed, precisely because hydrophones generate more data than can be stored or transmitted, it is highly desirable if not necessary to have embedded detection and classification algorithms as part of such systems.

For extreme applications with minimal power budgets (e.g., ocean gliders), microcontrollers with custom-designed circuit boards may still be required. Although some deep learning algorithms can run on such devices, they are limited in terms of the complexity of the tasks that they can perform. In these situations, the software and hardware development also require a higher skill level.

Since SBCs are closer to regular computers than microcontrollers, their use requires a lower skill level, making them accessible to a broader cohort of ocean researchers without advanced knowledge of electronics or firmware development. With modern deep learning frameworks, an increasing variety of models can be developed on PCs and then deployed on devices like the Raspberry Pi and the Jetson Nano with minimal modifications. For applications where the power constraints are not extreme, these devices can bring deep learning to field instruments in their off-the-shelf form, which makes development simpler, faster, and less costly. SBCs typically require in the range of 1 to 10 W, which poses an obstacle to systems powered by non-rechargeable batteries but can be accommodated in platforms with rechargeable power banks.

Finally, the computing modules can also serve as the deep learning unit in custom-made hardware. The Raspberry Pi Foundation offers a compute module, designed to be embedded in custom-designed carrier boards. Both the Jetson and Coral computing modules are available without their development boards, so they can be embedded into hardware specifically designed for a particular application. In underwater acoustics, for example, the custom hardware can deal with duty-cycles, audio recording, power management, and data transmission, only using the deep learning computing module when necessary. The advantage is that this custom board can be more efficient, since it will only include what is necessary, while also incorporating useful components not available in their off-the-shelf counterparts.

SBCs and dedicated hardware are making the use of cutting-edge deep learning algorithms in field deployments more accessible. Coupled with modern deep learning software frameworks, researchers can deploy the same trained model on a growing number of platforms. Although there are scenarios where custom hardware development might be required and power constraints mandate the use of microcontrollers, there are many cases in which off-the-shelf computing modules and development boards can be adequate. As software and hardware continue to develop, advanced deep learning techniques will become even more accessible to ocean scientists.

Fabio Frazao is the Deep Learning Developer at MERIDIAN. Dr. Oliver Kirsebom is the Senior Staff Scientist at MERIDIAN and an Adjunct Professor in the Faculty of Computer Science at Dalhousie University. Bruno Padovese is the Data Augmentation Specialist at MERIDIAN. Dr. Stan Matwin is the Principal Investigator of the MERIDIAN project, Director of the Institute for Big Data Analytics, and Professor in the Faculty of Computer Science at Dalhousie University. https://meridian.cs.dal.ca