Schlumberger: Deep Reinforcement Learning to Automate Acoustic Data Processing

Team: Sofiane Hadji and Joey Khoury El Aramouni
Schlumberger Advisors: Sandip Bose and Lingchen Zhu

1. Scope

A major part of Schlumberger's revenue has to do with the imaging and processing of data for its client to evaluate oil wells. For that reason, a wide variety of measurements are collected encompassing a number of modalities such as electromagnetic and acoustic often with tools lowered into wells in the ground and scanning thousands of meters of rock formations. This diverse data provides answers that are important to understanding the subsurface geology. Currently, there is still a great amount of manual intervention in data processing, which limits the value that can be extracted from the measurements as well as the quality of the subsequent workflows. Therefore, in this project Deep Reinforcement Learning (DRL) was leveraged in to order automate acoustic data processing and decrease its turnaround time by 30%.

2. Data

The raw waveform of the acoustic data and the dimensionally reduced features extracted from this data using a variational autoencoder algorithm served as inputs to the DRL scheme.

DRL Algorithm Picks Correct Echoes

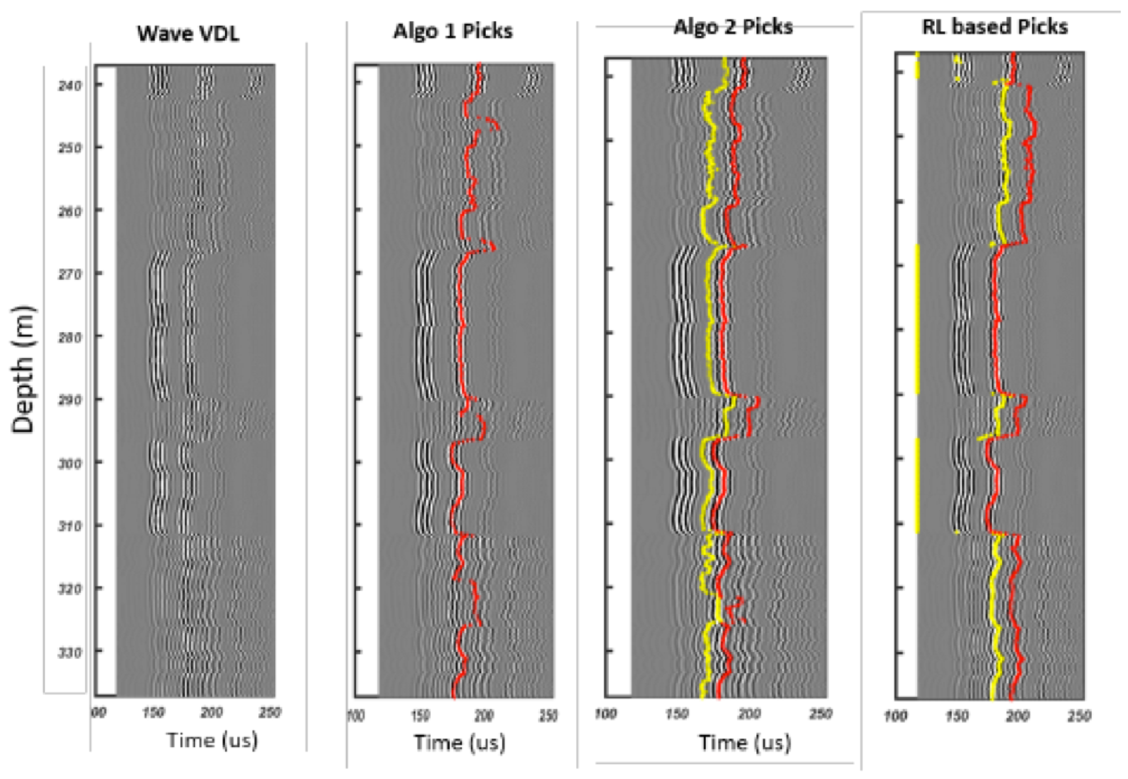

3. DRL Cycle

The agent picks up an input from the environment (VAE output or waveform), takes an action (choice of echo tracking method and additional parameter), receives a reward (using the actual echoes or the output of a data derived quality control (QC) function), and starts over again.

DRL Summary

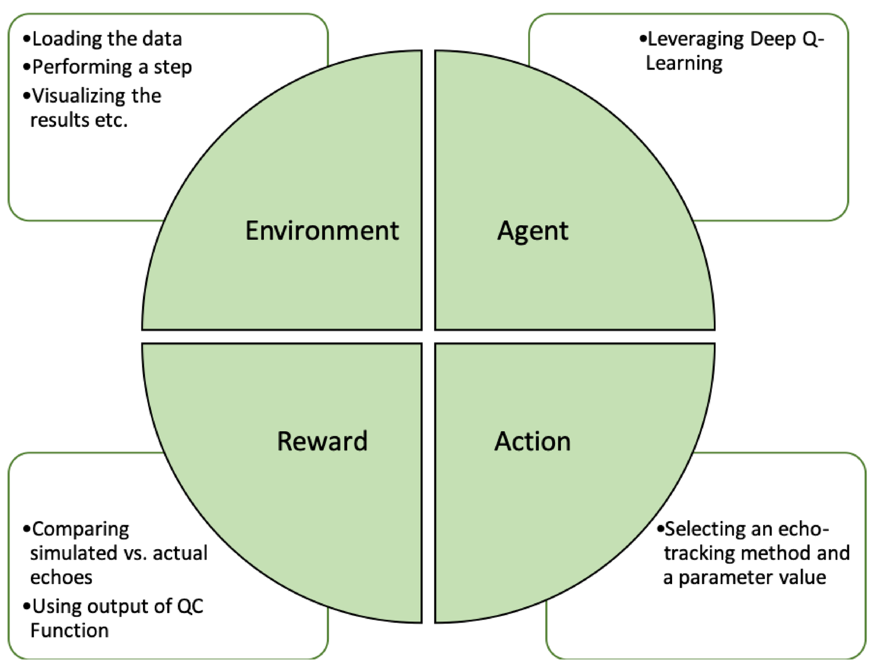

DRL Cycle

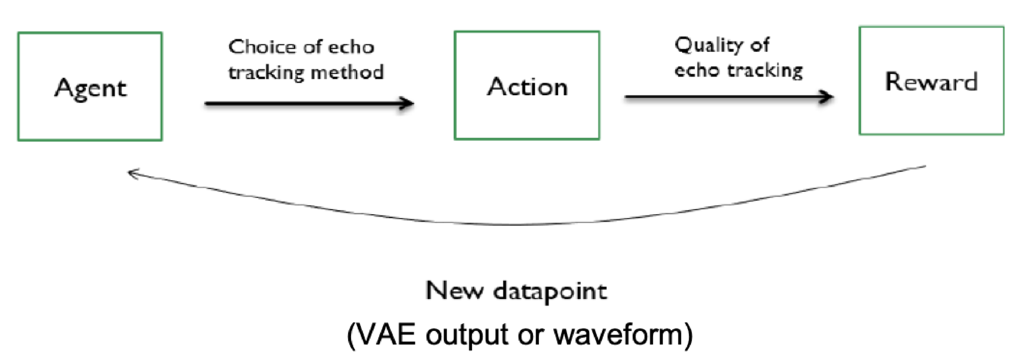

4. Exploration vs. Exploitation

Initially, it took about 10 hours to run the algorithm. Since we were not working in a very complex environment, we did not need to spend a considerable amount of time in exploration mode in the beginning of training. We thus decided to decrease the initial value of epsilon and were able to decrease the processing time to about 25 min.

Processing time vs. initial epsilon

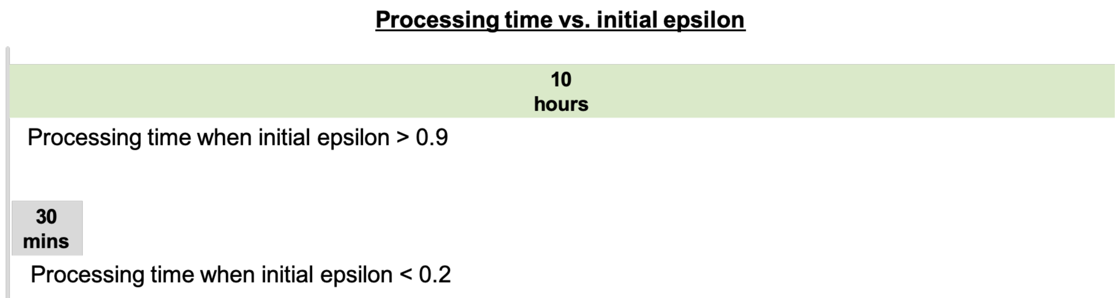

5. DRL Performance

Reward converged to 6,000 in only 3 episodes when epsilon was set to 0.2 while training using synthetic dataset 1. Moreover, the performance of the DRL algorithm was consistent when training using one environment and testing on a different environment (using real echoes). Finally, the algorithm was able to track the correct echoes when trained using the QC function.

Model performance when trained and tested on different environment

0%

% Error when training on synthetic 1 and testing on synthetic 2

2.8%

% Error when training on synthetic 3 and testing on synthetic 2

Reward vs. Episode (Actual Echoes)

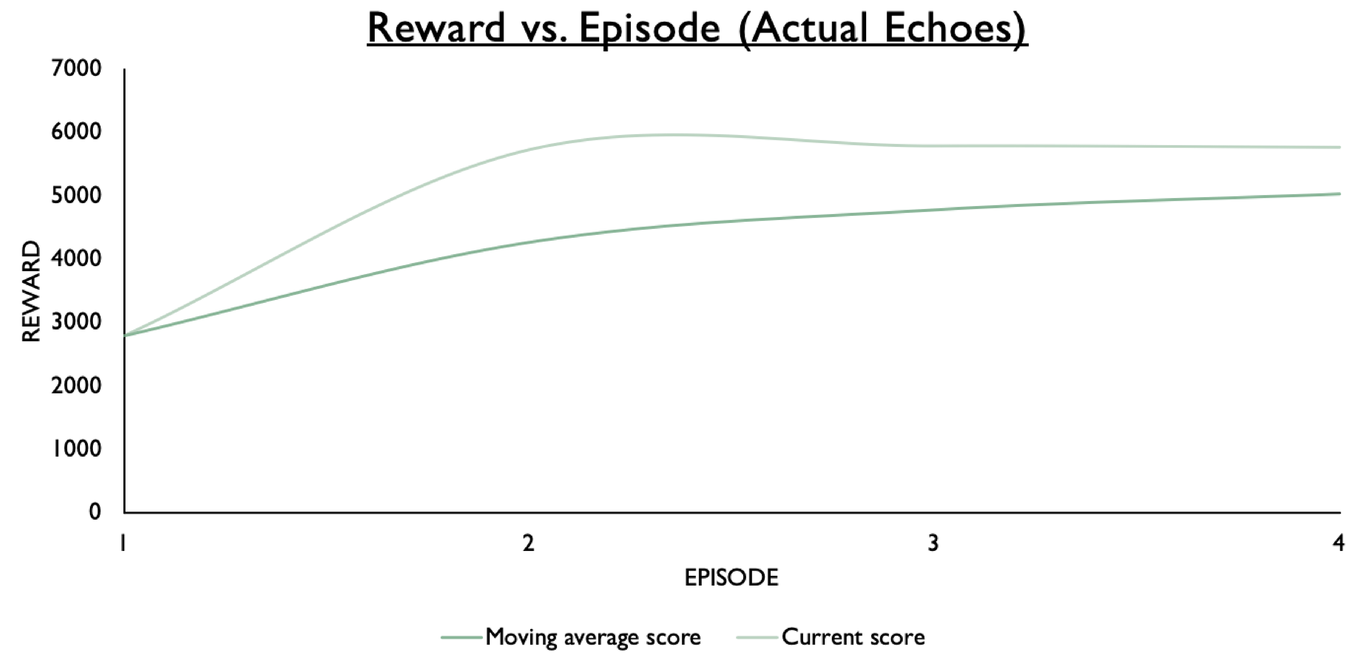

QC reward behavior on individual frame (1-echo frame)

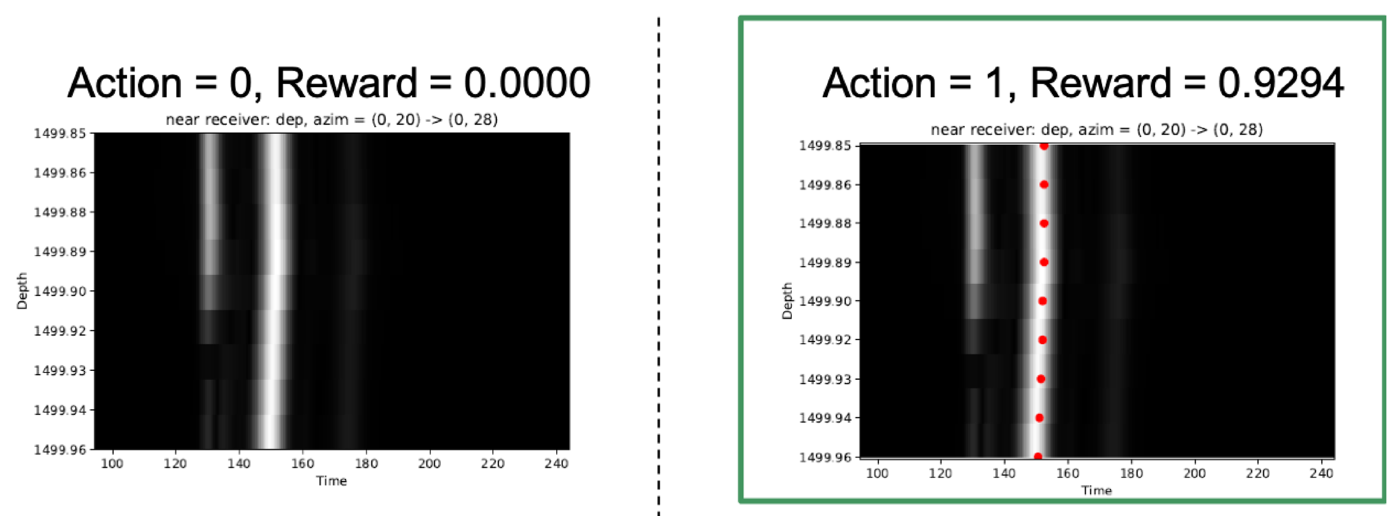

6.DRL Flexibility

Two approaches were explored in order to increase the flexibility of the DRL algorithm.

Methods explored to increase the flexibility of the DRL algorithm

1. Training the DRL algorithm on multiple datasets
2. Continuous training: saving a trained DRL model and continue training on the dataset that needs to be tested

7. Business Impact

The cost of acoustic data processing is close to $30M per year. The DRL algorithm we created will result in 30% cost saving, which is equivalent to $9M saved every year.

Business Impact of the project

$30M	Up to 30%	Up to $9M
Data Processing Cost ($)	Cost saving (%)	Cost saving ($)