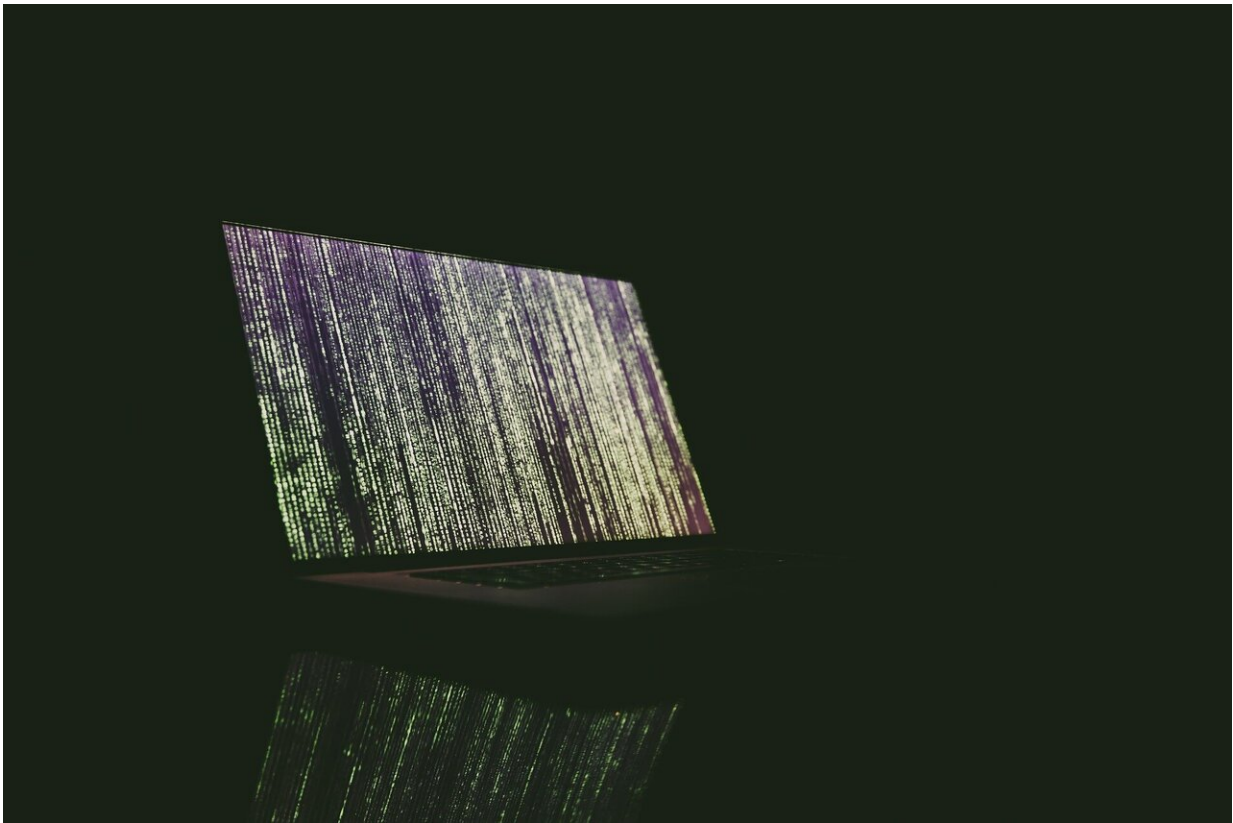# Positive reinforcements help algorithm forecast underground natural reserves

February 23 2021, by Vandana Suresh

Credit: Pixabay/CC0 Public Domain

Texas A&M University researchers have designed a reinforcement-based algorithm that automates the process of predicting the properties of the underground environment, facilitating the accurate forecasting of

oil and gas reserves.

Within the Earth's crust, layers of rock hold bountiful reservoirs of groundwater, oil and natural gas. Now, using machine learning, researchers at Texas A&M University have developed an [algorithm](#) that automates the process of determining key features of the Earth's subterranean environment. They said this research might help with accurate forecasting of our natural reserves.

Specifically, the researchers' algorithm is designed on the principle of reinforcement or reward learning. Here, the [computer algorithm](#) converges on the correct description of the underground environment based on rewards it accrues for making correct predictions of the pressure and flow expected from boreholes.

"Subsurface systems that are typically a mile below our feet are completely opaque. At that depth we cannot see anything and have to use instruments to measure quantities, like pressure and rates of flow," said Siddharth Misra, associate professor in the Harold Vance Department of Petroleum Engineering and the Department of Geology and Geophysics. "Although my current study is a first step, my goal is to have a completely automated way of using that information to accurately characterize the properties of the subsurface."

The algorithm is described in the December issue of the journal *Applied Energy*.

Simulating the geology of the underground environment can greatly facilitate forecasting of oil and gas reserves, predicting groundwater systems and anticipating seismic hazards. Depending on the intended application, boreholes serve as exit sites for oil, gas and water or entry sites for excess atmospheric carbon dioxide that need to be trapped underground.

Along the length of the boreholes, drilling operators can ascertain the pressures and flow rates of liquids or gas by placing sensors. Conventionally, these sensor measurements are plugged into elaborate mathematical formulations, or reservoir models, that predict the properties of the subsurface such as the porosity and permeability of rocks.

But reservoir models are mathematically cumbersome, require extensive human intervention, and at times, even give a flawed picture of the underground geology. Misra said there has been an ongoing effort to construct algorithms that are free from human involvement yet accurate.

For their study, Misra and his team chose a type of [machine-learning](#) algorithm based on the concept of reinforcement learning. Simply put, the software learns to make a series of decisions based on feedback from its computational environment.

"Imagine a bird in a cage. The bird will interact with the boundaries of the cage where it can sit or swing or where there is food and water. It keeps getting feedback from its environment, which helps it decide which places in the cage it would rather be at a given time," Misra said. "Algorithms based on reinforcement learning are based on a similar idea. They too interact with an environment, but it's a computational environment, to reach a decision or a solution to a given problem."

So, these algorithms are rewarded for favorable predictions and are penalized for unfavorable ones. Over time, reinforcement-based algorithms arrive at the correct solution by maximizing their accrued reward.

Another technical advantage of reinforcement-based algorithms is that they do not make any presuppositions about the pattern of data. For example, Misra's algorithm does not assume that the pressure measured

at a certain time and depth is related to what the pressure was at the same depth in the past. This property makes his algorithm less biased, thereby reducing the chances of error at predicting the subterranean environment.

When initiated, Misra's algorithm begins by randomly guessing a value for porosity and permeability of the rocks constituting the subsurface. Based on these values, the algorithm calculates a flow rate and pressure that it expects from a borehole. If these values do not match the actual values obtained from field measurements, also known as historical data, the algorithm is penalized. Consequently, it is forced to correct its next guess for the porosity and permeability. However, if its guesses were somewhat correct, the algorithm is rewarded and makes further guesses along that direction.

The researchers found that within 10 iterations of reinforcement learning the algorithm was able to correctly and very quickly predict the properties of simple subsurface scenarios.

Misra noted that although the subsurface simulated in their study was simplistic, their work is still a proof of concept that reinforcement algorithms can be used successfully in automated reservoir-property predictions, also referred as automated history matching.

"A subsurface system can have 10 or 20 boreholes spread over a two- to five-mile radius. If we understand the subsurface clearly, we can plan and predict a lot of things in advance, for example, we would be able to anticipate subsurface environments if we go a bit deeper or the flow rate of gas at that depth," Misra said. "In this study, we have turned history matching into a sequential decision-making problem, which has the potential to reduce engineers' efforts, mitigate human bias and remove the need of large sets of labeled training data."

He said future work will focus on simulating more complex reservoirs and improving the computational efficiency of the algorithm.

  More information: Hao Li et al, Reinforcement learning based automated history matching for improved hydrocarbon production forecast, *Applied Energy* (2020). DOI: 10.1016/j.apenergy.2020.116311

Provided by Texas A&M University