Can Quantum Computing bring R2D2 to Life?

Nearly everyone knows R2D2, the beloved droid from the Star Wars universe. Such a sophisticated droid may not exist in real life yet, but scientists have started developing techniques to train robots for simple but essential tasks, such as walking through a room filled with obstacles. This is an easy task for most humans, as our eyes help us to avoid the obstacles and our sense of orientation will lead us to the goal. A robot, however, does not possess these natural abilities, so how can we teach them to do the same?

A sophisticated, yet intuitive method is called reinforcement learning. This machine learning technique lets the robot learn information about its environment through trial-and-error. Simply put, the robot (or a simulation of it) is put in the obstacle-filled room and has some way to observe it (e.g., using distance sensors). Next, it is allowed to move through the room and collect rewards by avoiding obstacles and moving closer to the goal. As time passes, the robot learns how collect as many rewards as possible, which ultimately leads it to the goal.

To process the observations and rewards, the robot uses a deep neural network. This works a bit like a simplified, artificial brain, that gets better through training. Zooming in, the neural network contains many 'trainable parameters' that are connected to each other – somewhat like brain cells. During the training process, the parameters are updated in a coordinated manner, considering the observations and rewards of the robot. When the training ends, the parameters are optimized such that the robot can exactly predict the outcome of each step it takes. As a result, the robot can reach the goal by simply by choosing the steps that lead to the best outcome.

Optimizing the parameters of the neural network can be done with a regular computer, but recent research showed that quantum computers, a new type of computers based on different physical properties, might also be good at this task. Quantum computers use qubits, instead of bits, to process information. Simply put, a bit is the smallest piece of information on a regular computer which has a value of 1 or 0. In contrast, a qubit is the quantum version of a bit and can both be 1 and 0 at the same time with a certain probability. In addition, qubits can be 'entangled' with each other, which means that manipulating one qubit may instantly influence one or more of the other qubits. These two properties make quantum computers potentially very powerful.

This brings us to the purpose of this master's thesis. A simple simulated robot equipped with some distance sensors was trained to walk through a room filled with obstacles. The neural network, functioning as the brain of the robot, was trained using qubits instead of regular bits. It turned out that the best quantum-based neural network needed twenty times fewer trainable parameters to learn how to navigate, compared to neural networks trained with regular computer. This is a promising result, as more complex tasks easily require millions of trainable parameters, which can be difficult to optimize.

However, many questions still remained unanswered. Firstly, the arrangement of the qubits influenced the performance of the quantum-based neural network substantially, although it was unclear which differences contributed the most. In addition, the obstacle-filled room was relatively small, so it is unclear how well the quantum-based neural network performs in larger spaces. Finally, the quantum-noise that is present in quantum computers may also impact the results, but that has not been investigated accurately enough yet. Taking all of this into consideration, we probably won't have a quantum-R2D2 anytime soon, but future perspectives look promising nevertheless!