Quantum Computing: A Brief History

With Applications of Quantum Computing in Automotive

David von Dollen, Daniel Weimer and Florian Neukart, Volkswagen Group

Quantencomputing: Eine kurze Geschichte – Anwendungen des Quantencomputings im Automobilbereich

In den letzten Jahren hat das Quantencomputing neue Erfolge erzielt, wie z. B. das Quantenüberlegenheitsexperiment von Google [1] und es zeigt sich, dass es von großen Industrieunternehmen zur Lösung komplexer Probleme eingesetzt wird. Aber was hat zu dieser Entwicklung geführt? Welche Arten von Problemen können wir in naher Zukunft erwarten, die mit Quantencomputing zu lösen sind? Auf welche Herausforderungen stoßen wir bei dieser Technologie und ihrem Einsatz im industriellen Umfeld?

Schlüsselwörter:
Quantencomputing, maschinelles Lernen, künstliche Intelligenz, Automotive, Preisgestaltung

In the last few years, quantum computing has achieved new successes, such as Google's quantum supremacy experiment [1], and has been showing adoption by large industrial firms to tackle complex problems. But what has led up to these developments? What kinds of problems can we expect to be able to solve in the near term with quantum computing? What are the challenges that we encounter with this technology and deploying within industrial settings?

Much of what we can achieve with quantum computing today can be attributed to 40 years of theoretical and technical development, kicked off by the renowned physicist Richard Feynman in 1982 in his work "Simulating physics with quantum computers" [2]. The main motivation for this work was the realization that in order to simulate physical states and interactions, you need to be able to compute an exponential number of probabilities for physical states. He reasoned that this would be highly inefficient using the types of computers available, which could only represent information in discrete bits of either 0 or 1, and that what was needed was a quantum computational system, which could leverage quantum mechanical properties such as superposition, entanglement, and interference, to efficiently simulate a quantum system and represent bits as 0, 1, or 0 and 1. In follow up works in the 1990s by other researchers, new quantum algorithms were discovered, which showed theoretical speedups by leveraging these properties, with applications in integer factoring, database search, and public key cryptography [3, 4]. In the late 1990s through the early 2000s, researchers implemented the first quantum computers with limited amounts of quantum bits (qubits) in the range of 2-7 qubits. In 2011, D-wave systems released the D-Wave One, a 128 qubit quantum annealing system.

Quantum Hardware and Software: Where We Are Today

We are currently in the midst of rapid development of quantum hardware and software technologies. Among these are chip types leveraging superconducting qubits, trapped ions, and quantum dots. Companies such as Google, IBM, Righetti are currently manufacturing chips and developing algorithms based on the universal gate model of quantum computing. This type of quantum computation is based on algorithms composed of quantum logic gates acting on quantum bits (qubits). There have been software packages designed to implement and experiment with this mode of computation, such as Google's TensorFlow-quantum and IBM's Qiskit [5, 6]. While many quantum algorithms have been developed using this paradigm, we are just beginning to see implementations realized, as these types of chips have limits in regards to qubit size and effects from noise. Improvements continue to be made, and this combination of hardware and software shows great promise in the near to medium term, as well as in the longer term as more systems with larger numbers of fault tolerant qubits become available.

Another type of device which leverages quantum effects is the D-Wave quantum annealer, which was designed to solve combinatorial optimization problems formulated as Ising spin glass models. In this mode of quantum computation, the system starts in a state of superposition, and follows an evolution described by quantum mechanics. If the evolution is slow enough, the system will remain close to a ground state, allowing for an effect known as "quantum tunneling" where barriers in the optimization landscape are tunneled through such that the evolution process finds a global minimum. This is in contrast with non-quantum, or classical computation, which walks, or samples from, the optimization landscape without leveraging quantum effects. D-Wave has steadily been releasing chips with growing numbers of qubits with the latest generation, D-Wave Advantage, consisting of 5000

David von Dollen is Lead Scientist at Volkswagen Group of America in San Francisco, California, United States.

Dr. Daniel Weimer is Head of Artificial Intelligence at Volkswagen Group of America in Detroit, Michigan, United States.

Prof. Dr. Florian Neukart is Director of Volkswagen Data:Lab in Munich, Germany and Assistant Professor at Leiden University, Netherlands.

david.vondollen@audi.com
www.volkswagengroupofamerica.com

qubits on a single chip [7]. It is important to note that a difference between quantum annealers and gate model chips is that error correction is not needed for the quantum annealing regime as it is for gate model. Quantum annealers also operate on fewer energy levels than gate model chips.

Quantum Computing: Opportunities for Applications in Industry

In recent years various industries have started to experiment with quantum applications for industrial problems [8, 9]. Areas which show promise for this include optimization, machine learning, and simulation. What types of problems can we expect to be able to apply quantum computing to in the near term within various industries?

Optimization

Many industries deal with hard optimization problems along their value chains. This can take the form of manufacturing processes, route planning, resource planning and allocation in response to supply and demand, as well product development, such as optimizing components and materials. This is where quantum computing may be applied to certain problems and have an impact in unlocking value and driving down costs.

Simulation

Materials simulation is another area in which quantum computing shows great promise. In fact, it is one of the original applications for which quantum computing was devised. Various teams are working towards implementations of chemical and materials simulations on near term quantum computers. Leveraging quantum computing to simulate systems could lead to breakthroughs in drug discovery, materials engineering and fabrication, generative design and optimization [1].

Machine Learning

Machine learning is another area in which quantum computing may play a role in generating performance gains. Machine learning algorithms have been shown to benefit from data preprocessing routines leveraging quantum computers, which can filter data to select input which may have more relevance and increased signal for a predictive model. There is ongoing research into leveraging quantum computers to encode data and assist in linear algebra subroutines, which may help in model training, and representation capability. Finally, there exists software libraries for implementing variational quantum circuits in the form of quantum neural networks, which may be able to learn decision boundaries for classical and quantum data.

Case Study: Quantum-Assisted Feature Selection for Vehicle Price Prediction

Predicting vehicle prices is a common task in the automotive industry, whether it is for new, used, or leased vehicles. Many times, this prediction problem can be modeled on data drawn from distributions of attributes such as vehicle weight, year, miles per gallon etc. But, how well do these features predict the target variable, price? Are there combinations of features which have stronger signal than other combinations or all features available in the data set?

This is where quantum computing is uniquely capable and can assist in the machine learning training regime. With a quantum computer, our team at VW showed that we were able to search the space of features such that we were able to find an optimal combination of features which maximized the relevancy to the price, while minimizing the redundancy between features. Less features means that the model costs in terms of training times and complexity were minimized as well. We also were able to

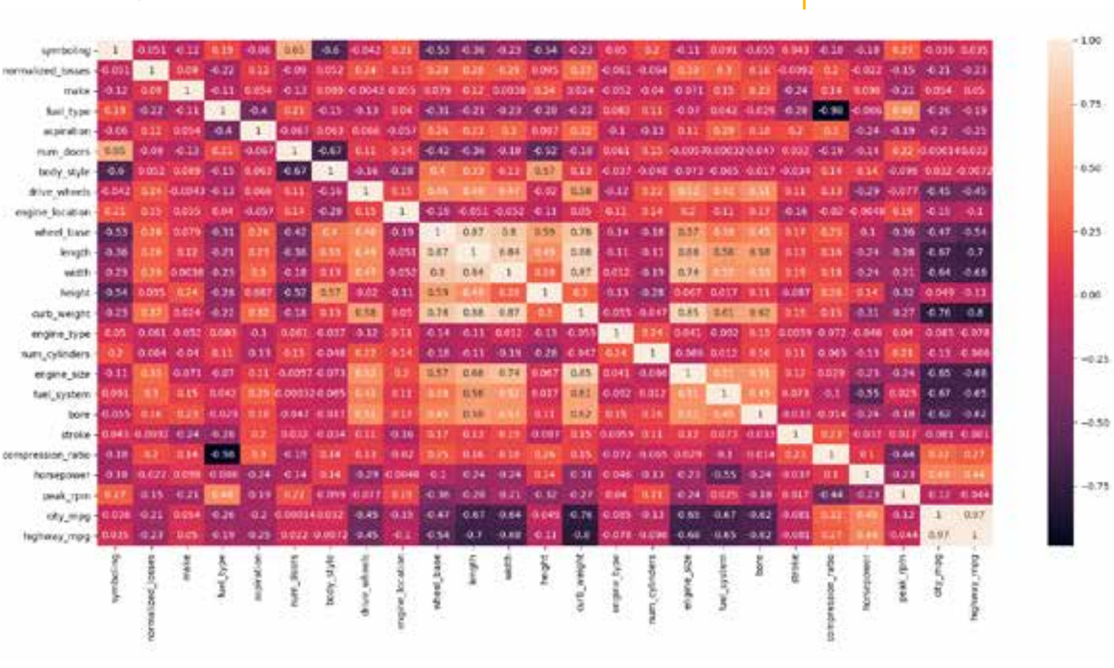

Image 1: Visualization of correlation between automobile data. Quantum computers may be leveraged to select specific combinations of features which finally improve predictive model performance.

capture the explainability in our modeling process, by being able to show which features had a greater weighting with respect to our prediction of price.

We used the open-sourced UCI auto dataset in order to demonstrate our process [2]. The dataset contained 26 attributes such as vehicle weight and mpg, as well as our target variable of price which was in the range of [51118, …, 45400]. We used a similar model on internal data with more attributes, which we could not publish, but achieved similar results.

In the first step of the quantum assisted routine, we calculated correlation statistics between each feature variable and the target price, as well pairwise correlation between features. These statistics became input for what is known as a quadratic binary optimization problem. The input represented a landscape, for which the quantum computer returned solutions in the form of 0 or 1, or binary indicators for whether or not each feature was to be included in the optimal set. Then, we sent this input to the quantum computer to search the optimization landscape. The quantum computer projected all configurations of features into a superposition, and then following a time dependent evolution process governed by quantum mechanics, slowly evolved the system. The output of the quantum computer took the form of samples, which were ranked by their energies. We selected the solution with the lowest energy, as that represented the configuration with the maximum relevancy and minimum redundancy, or the ground state of the quantum-mechanical evolution of the system.

Using this solution, we were able to filter our data set down and feed into a predictive model, in this case a gradient boosted regression tree algorithm, and achieve higher performance than using all features or other feature reduction techniques. The gradient boosted trees algorithm works by sequentially creating trees from the features of the training data sets, and the structure of the combined trees form an ensemble of trees with which to output a prediction given new data. The loss metric for training the algorithm was the mean absolute error. Other feature reduction techniques included for comparison included greedy selection and recursive feature elimination. We showed an average reduction in mean absolute error of price predictions (MAE = \$1471.02 +/- 135.6) of around 45 % versus all features with a simple linear regression model (MAE =2690.5 +/- 460). This cost savings, when applied to tens or hundreds of thousands of automobiles, is the type of value that can be unlocked with quantum computing and machine learning.

Challenges

In the near term, challenges can take the form of hardware implementation, in that the chips being designed and produced today have small numbers of qubits available, limited coherence times, low resiliency to noise and interference from outside environments, and high manufacturing costs. Many chips require extremely low temperature environments, with operating temperatures at around 15 millikelvin, and costly dilution refrigerators to operate. Companies are just beginning to figure out how to deploy these systems into production, and tools are being created in order to monitor and serve these deployments in the form of hybridized architectures which may call an external API for a quantum sub routine.

Quantum computers have the potential to unlock value for companies by being applied to hard real-world problems and lowering costs and opening new revenue streams. But like other areas of technology such as artificial intelligence, it will take effort to translate business problems down to a level where the problem can be solved on a hardware level. Further effort will need to be undertaken to deploy, integrate, and continuously monitor these solutions.

Additional challenges may also take the form of resources in regards to trained experts. At present the skills required to understand how to translate a business requirement in the form of an optimization or machine learning problem include: knowledge of programming frameworks for the quantum processing unit, knowledge of quantum mechanics, linear algebra, probability, and the mechanisms and advantages of quantum algorithms and how to apply them in the correct settings. A future challenge that companies looking to leverage this technology may be access to trained experts in this field.

Summary

Quantum computing shows promise in application to hard challenges in regards to optimization, machine learning, and simulation that are encountered across various industries. While there are certain constraints in regards to the size of the problems that quantum computers can tackle, and quantum chips continue to be limited in coherence times and by noise, there are areas where quantum computing may be applied to unlock value and cut costs for companies across industries today.

Keywords:
quantum computing, machine learning, artificial intelligence, automotive, pricing

References

[1] Nature: Hello quantum world! Google publishes landmark quantum supremacy claim. URL: www.nature.com/articles/d41586-019-03213-z, Date of access 01.06.2021.
[2] Feynman, R.: Simulating physics with computers. In: Int J Theor Phys 21 (1982), pp. 467-488.
[3] Shor, P.: Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer. URL: arXiv:quant-ph/9508027, Date of access 01.06.2021.
[4] Grover, L.: A fast quantum mechanical algorithm for database search. URL: arXiv:quant-ph/9605043, Date of access 01.06.2021.
[5] M. B. et al.: TensorFlow Quantum: A Software Framework for Quantum Machine Learning. URL: arXiv:2003.02989, Date of access 01.06.2021.
[6] IBM: Qiskit. URL: https://qiskit.org, Date of access 01.06.2021.
[7] TechRepublic: D-Wave announces 5,000-qubit fifth generation quantum annealer. URL: www.techrepublic.com/article/d-wave-announces-5000-qubit-fifth-generation-quantum-annealer, Date of access 01.06.2021.
[8] F. N. et al.: Traffic Flow Optimization Using a Quantum Annealer. URL: www.frontiersin.org/articles/10.3389/fict.2017.00029/full, Date of access 01.06.2021.
[9] S. Y. et al.: Quantum Shuttle: Traffic Navigation with Quantum Computing. URL: https://arxiv.org/abs/2006.14162, Date of access 01.06.2021.