

Received March 29, 2022; revised May 5, 2022; accepted May 16, 2022; date of publication May 23, 2022; date of current version June 17, 2022.

Digital Object Identifier 10.1109/TQE.2022.3176806

Quantum Kernels for Real-World Predictions Based on Electronic Health Records

ZORAN KRUNIC¹ , FREDERIK F. FLÖTHER² , GEORGE SEEGAN¹,
NATHAN D. EARNEST-NOBLE³, AND OMAR SHEHAB³ 

¹Amgen, Thousand Oaks, CA 91320 USA

²IBM Quantum, IBM Switzerland Ltd., CH-8010 Zurich, Switzerland

³IBM Quantum, IBM Thomas J. Watson Research Center, Yorktown Heights, NY 10598 USA

Corresponding author: Zoran Krunic (e-mail: zkrunic@amgen.com).

ABSTRACT Research on near-term quantum machine learning has explored how classical machine learning algorithms endowed with access to quantum kernels (similarity measures) can outperform their purely classical counterparts. Although theoretical work has shown a provable advantage on synthetic data sets, no work done to date has studied empirically whether the quantum advantage is attainable and with what data. In this article, we report the first systematic investigation of empirical quantum advantage (EQA) in healthcare and life sciences and propose an end-to-end framework to study EQA. We selected electronic health records data subsets and created a configuration space of 5–20 features and 200–300 training samples. For each configuration coordinate, we trained classical support vector machine models based on radial basis function kernels and quantum models with custom kernels using an IBM quantum computer, making this one of the largest quantum machine learning experiments to date. We empirically identified regimes where quantum kernels could provide an advantage and introduced a terrain ruggedness index, a metric to help quantitatively estimate how the accuracy of a given model will perform. The generalizable framework introduced here represents a key step toward *a priori* identification of data sets where quantum advantage could exist.

INDEX TERMS Artificial intelligence, digital health, electronic health records (EHR), empirical quantum advantage (EQA), machine learning, quantum kernels, real-world data, small data sets, support vector machines (SVM).

I. INTRODUCTION

Over the last years, real-world data have been increasingly used to generate medical evidence and progress precision medicine. This includes sources such as electronic health records (EHRs), claims and billing data, product and disease registries, and data from wearables and health applications [1]. Powerful data mining techniques have been applied to such data sets, particularly to EHRs, to predict a broad range of medical conditions and events [2]–[5]. However, classical machine learning and data science techniques have limitations with regard to learning some of the most complex patterns; for instance, the predictive power of genetic risk scores derived from genome-wide association studies has plateaued over the last years [6]. As a result, quantum machine learning has been explored as an alternative and for certain general problems, it has already been proved that quantum machine

learning algorithms can provide benefits beyond the scope of classical ones [7]. The complexity of the correlations and patterns in EHRs and (real-world) medical data sets make such data sources prime candidates for the application of quantum algorithms [8].

The study of supervised machine learning problems with quantum techniques is an active area of research [9]. In early work on classification with near-term quantum algorithms [10]–[12], the proposed quantum feature maps typically encode the datapoints into inner products or amplitudes in the Hilbert space. The quantum circuit used to implement the feature map is of a length that is typically a linear or polylogarithmic function of the size of the data set and the number of qubits is a function of the number of features. In subsequent work, the advantage of a quantum feature map was rigorously proved for a carefully chosen synthetic data

set [7]. Recently, a body of work [13], [14] implementing quantum feature maps for small-scale coarse-grained practical data sets has emerged. Whereas there have been studies of different feature maps [15], none have been discovered so far with rigorous advantage for general or practical data sets. The capability limits of near-term quantum computers have also been pushed in work where the data set was less coarse-grained [16], [17]; furthermore, efforts have begun to study how hyperparameters affect the potential advantage of a given quantum classifier [18]. Another stream of research has emerged on finding a suitable quantum feature map for a given data set with [19] providing a recent review of quantum classification algorithms. Studies on quantum feature maps involve both the study of the kernel function and the study of the quantum circuits that encode the outcome of the kernel function into Hilbert space [20]–[22]. A new set of metrics and a protocol have also been proposed to determine the possibility of quantum advantage for a given pair of data set and quantum feature map [23].

In this article, we focus on one kernel-based method that uses the quantum support vector machine (QSVM), estimating the kernel with a quantum computer and feeding it back into a classical support vector machine (SVM) for classification [11]. We empirically study classical and quantum SVMs, allowing for a direct comparison between a quantum method and a corresponding state-of-the-art classical algorithm. Since it has been shown that for large data sets, there are no significant differences between classical SVMs and other classical approaches such as random forests [24], we conjecture that our approach is generalizable to other settings where different classical algorithms are used. However, the existence of empirical advantage depends on the specific application studied.

To the best of our knowledge, there have not yet been any systematic studies regarding the applicability of quantum kernels to EHRs. Here, we predict the six-month persistence of rheumatoid arthritis (RA) patients on biologic therapies. The central research questions investigated in this article are as follows.

- 1) Can we enhance the prediction of medication persistence by applying quantum kernels to real-world EHRs?
- 2) Can we systematically identify problem instances (number of features and number of samples) where quantum computing may have an advantage for such real-world data sets?

The methods developed in this article are general and can be applied for a wide variety of problems with different-size data sets in machine learning and optimization. In this article, nevertheless, we focus on small data sets, particularly those where the ratio of the number of features to the number of samples is relatively large, which typically engenders hard classification problems. Such data sets are important in a range of medical settings, for instance in clinical trials, studies of very specific cohorts, and translational medicine.

Moreover, small data sets are naturally suited to near-term quantum computers. Predicting medication persistence is one of the more difficult problems when applying machine learning to EHRs; we deliberately selected such a challenging setting to explore quantum computing for research questions where classical solutions are often insufficient.

The concepts “quantum supremacy” [25] and “quantum advantage” [26] have been around for a while and refer to asymptotic performance comparisons between a quantum approach and the best classical approach. Complementing these two foundational concepts, in this article, we introduce a related concept called empirical quantum advantage (EQA). We define EQA as the incremental gain of using a specific quantum approach over a specific classical approach for a given problem. Once this heuristic measure is calculated, it is meaningful only in the context of three elements—the problem as well as the classical and quantum approaches used. It may not give any general asymptotic indications about “supremacy” or “advantage” for a family of problems. However, as in the field of practical classical algorithms [27], practitioners may use EQA to observe trends in empirical data. This is a key in biology and medicine where both theoretical and operational factors must be considered, in general, when exploring the benefits of quantum algorithms for a given application [28].

When making measurements of EQA, multiple metrics were considered, with a final choice of three key metrics—F1 score and balanced accuracy at the configuration space coordinate level as well as the phase space terrain ruggedness index (PTRI) at the configuration space landscape level. PTRI is thus a global metric, fully described in Section V. The reasoning behind choosing these metrics was as follows.

Both F1 score and balanced accuracy are commonly used in machine learning; they measure the performance of a given model. On the other hand, PTRI captures the hardness of the configuration space for a given set of machine learning problems. The typical coordinate structure to explore that space of problems consists of the number of features and the number of samples, whereas the final data set used for binary classification is quite balanced (52% to 48%), more imbalanced cases were also considered. F1 score and balanced accuracy are thus readily generalizable to more imbalanced settings in future research. Furthermore, we chose to present the F1 score because, although it equally weights false negatives and false positives, we do not have the exact cost of either of those. In other words, the relative cost of recall and precision are different in specific model deployments.

II. RESULTS

We started by evaluating multiple two-dimensional landscapes in the classical domain with the number of topmost important features ranging from 1 to 20 in increments of 1 and the training set sizes ranging from 50 to 600 samples in increments of 50. The topmost features were determined using the SHapley Additive exPlanations (SHAP) method [29] (see Section V). For each landscape coordinate, 200 random

train/test subsets were created out of the available data. Since the most resource- and time-demanding part of the process is the calculation of the custom kernels in quantum simulations (classical hardware simulating the behavior of a quantum computer) and on a real quantum processing unit (QPU), a small number of these data sets were chosen to evaluate the required total processing time. Custom kernels were calculated using both quantum simulations and QPUs. With the obtained runtimes, a realistic number of data sets that could be executed were calculated. As a result, the configuration landscape was reduced to four feature number values [5], [10], [15], [20] and three training set sizes [200, 250, 300], yielding a 12-point configuration space. For each configuration space coordinate, two data sets, each containing a train and a test subset, were selected out of the 200 train/test subsets. This was achieved by calculating balanced accuracy in the classical domain and then selecting two sets with the balanced accuracy close to the mean of the balanced accuracies of the full 200-sample set. The term subpoint is used to denote each individual data set within the given coordinate.

This yielded a total of 24 subpoints across the 12-coordinate configuration space. We found this to be feasible for quantum simulation and, critically, also for QPU execution. Since the analytical study of the hardness of such a large practical problem is extremely difficult, these types of large-scale simulation and hardware experiments across a broad configuration space are the most pragmatic way to identify trends and outliers. As a parallel outcome, this work hence represents one of the largest quantum machine learning experiments to date. The feature size component of the coordinates dictates the number of qubits for the QSVM. It was obtained by taking the most important features from the classical models built on the same data using the full-size data sets (see Section V).

We used the predict method from the svm.SVC class within scikit-learn [30] as the main method for comparing quantum and classical SVM performance. In addition, the predict_proba method was used to obtain estimates of probabilities. Thresholds were varied in the range of 0 to 1 in small increments and applied to the probabilities to generate the optimal split between the two class labels. The plots presented focus on the predict method; a detailed discussion of the probability-based results can be found in Section V.

Presented in Fig. 1 are the comparative three-dimensional plots of F1 score and balanced accuracy metrics with the orange surface presenting points of the averaged metric for classical computing and the blue surface for quantum computing (QPU). All QPU experiments presented in this article were run on ibmq_dublin (see Section V). Each point on the configuration space coordinate was averaged from two selected data sets for that coordinate; thus, each plotted configuration space has 12 points in total. The z -axis is the metric and the x - and y -axis are the number of features and training samples, respectively.

The PTRI was calculated for the full configuration space both for balanced accuracy and F1 score metrics and plotted

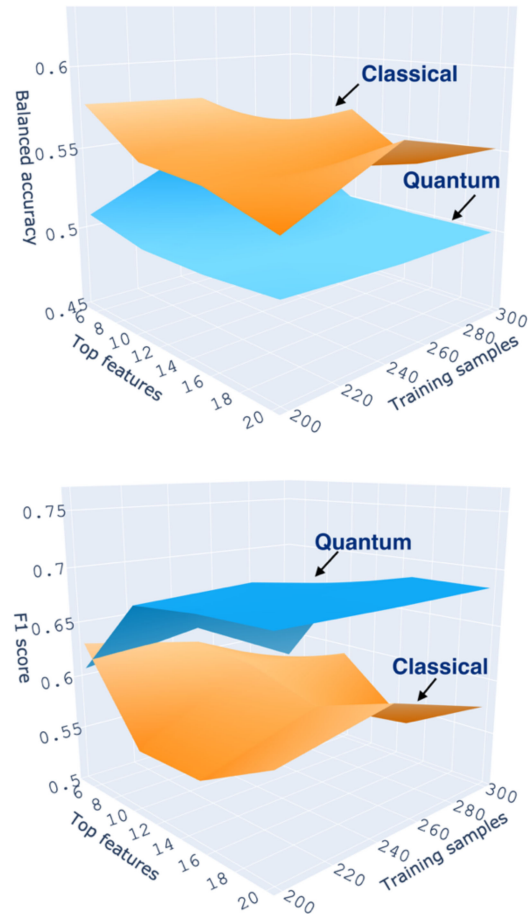


FIGURE 1. (Top) Balanced accuracy landscape. (Bottom) F1 score landscape. The classical domain is shown in orange and the quantum domain is in blue.

in Fig. 2 using the same approach as previously presented for the other metrics.

For each point in the configuration space and its corresponding two subpoints, we calculated the geometric difference defined in [23] between radial basis function and quantum kernels and averaged the corresponding two values at each coordinate, as shown in Fig. 3.

The plots in Fig. 4 illustrate the results obtained with QPU and classical processing for different metrics. The plots show the metric value achieved by QPU minus the metric value calculated by classical SVM. Since there are two random data sets for each coordinate, those are averaged and then the difference is calculated, resulting in 12 configuration space points. The points are plotted as three lines, one for each training set size, across four feature number values.

By comparing the position of the data points with the horizontal “zero-advantage line,” we observe EQA for a subset of problem instances in the configuration space: 0% and 92% of all instances for nonprobability-based balanced accuracy and F1 score, respectively, as well as 33% and 8% of all instances for probability-based balanced accuracy and F1 score, respectively, showed such quantum advantage.

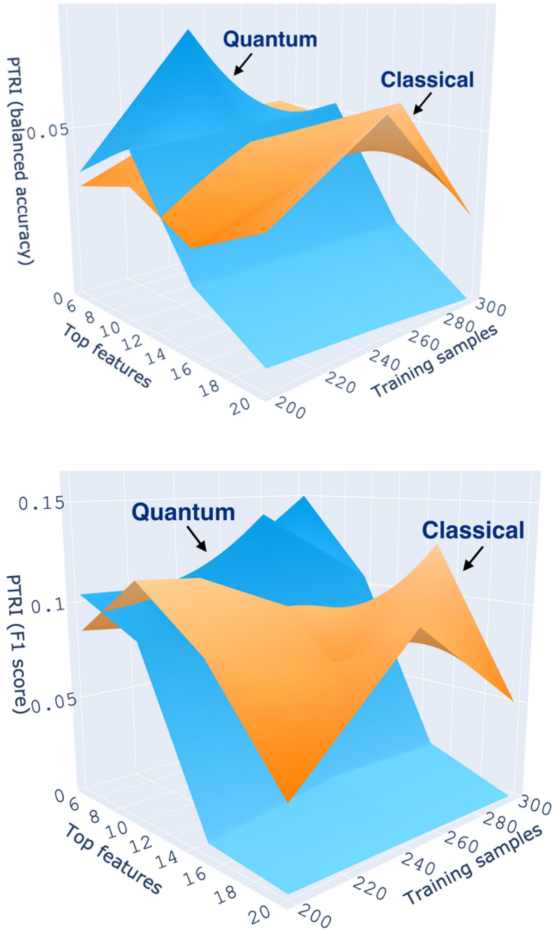


FIGURE 2. (Top) PTRI (balanced accuracy) landscape. (Bottom) PTRI (F1 score) landscape. The classical domain is shown in orange and the quantum domain is in blue.

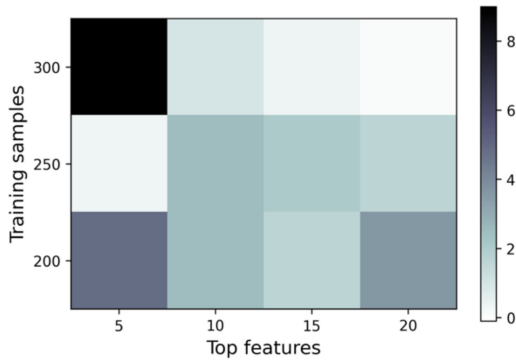



FIGURE 3. Geometric difference between classical and quantum kernels across the configuration space. Despite its name, the mathematical definition demonstrates that it is a ratio. The greater the geometric difference, the more potential for quantum advantage the given quantum kernel has compared to the given classical kernel.


III. DISCUSSION

In this article, we have considered a configuration space of classification problems with varying numbers of features and samples. On that manifold, we have observed EQA for 0% and 92% of classification problem instances for


Quantum-to-classical advantage balanced accuracy nonprobability-based



Quantum-to-classical advantage F1 nonprobability-based



Quantum-to-classical advantage balanced accuracy probability-based



Quantum-to-classical advantage F1 probability-based




FIGURE 4. QPU balanced accuracy and F1 score versus classical balanced accuracy and F1 score, nonprobability, and probability-based approach.

nonprobability-based balanced accuracy and F1 score, respectively, as well as 33% and 8% of classification problem instances for probability-based balanced accuracy and F1 score, respectively. This makes it apparent that EQA is something that must be evaluated on a case-by-case basis until clearer trends present themselves. Identification of hard instances through careful domain consideration has allowed us to observe such advantages with no circuit being executed more than 1024 times (i.e., a maximum of 1024 shots),

which is almost one order of magnitude less than previous large-scale quantum machine learning experiments [16], [17] (and therefore results in higher sampling noise). This indicates that, in the future, domain expertise in the hardness of practical problems is going to be crucial for the development and refinement of quantum algorithms. Our observation of these empirical trends reiterates the significance of developing such large-scale experiments to understand the trends and detect outliers. The considerable differences in EQA based on the choice of performance metric could suggest that practical quantum advantage is going to be highly domain-specific. Further work is needed to explore the applicability of different performance metrics for various domains. Given that there is a need to build robust machine learning models in medical settings where additional samples are costly or impossible to acquire, even a modest reduction in the number of samples required for training based on certain data distributions can yield considerable benefits for many predictions and inference problems in biology [28].

We also introduced a practical metric, i.e., PTRI, to quantify and thereby qualify the quantum advantage potential for a given problem. For any metric, PTRI helps identify the flattest and most rugged regions in configuration space. One could imagine that the flattest classical performance region is the configuration subspace where the performance of the classical techniques becomes stagnant and where a quantum algorithm should therefore be considered. In that case, computing the PTRI for the quantum approach over the given configuration space may give some insights into where quantum advantage is likeliest. This domain-agnostic metric is one of the first attempts for an operational tool that, in the future, quantum practitioners can use to determine when to use a quantum computer, a dynamic decision that may have to be taken very frequently, under severe timing constraints. Further study is needed to interpret the amount of correlation between the manifolds of classical and quantum performance metrics in terms of PTRI and related measures.

As a parallel result, we have also presented, to the best of our knowledge, the first independent application of the geometric difference, which we employed to determine the relative separation between classical and quantum feature maps. Further study is needed to understand how quantum practitioners may combine the concepts of PTRI and geometric difference to first identify the potential for quantum advantage in a configuration subspace and then estimate the potential of a specific quantum feature map in that subspace. We emphasize that there may be other relevant metrics worth exploring in the future when studying forms of quantum advantage, such as energy consumption.

It is also important to observe that we used the same kernel function and feature map for every classification problem. More studies are needed to determine appropriate combinations of kernel function and feature map that result in greater EQA. It may also be worthwhile to investigate whether there are kernel functions inspired by one-way [7], trapdoor, or learning with error [31] protocols that may not

only provide an advantage in prediction accuracy but also in time complexity.

Finally, we also note that being generated using near-term quantum computers, the kernels of interest in this work are noisy. It has been shown that the advantage of quantum kernels, expressed in terms of the generalization error, vanishes for large datasets, fewer measurements, and increased system noise [15]. However, we believe the setup used to understand the relationship between noise and quantum advantage may be incomplete. The quantum advantage would be better understood in the context of a phase transition in the computational hardness of SVMs. Instead of properties like dataset size, higher-level structural artifacts such as sparsity [32] and characteristics of the input distribution [33] have been used to study the phase transition in the computational hardness of SVMs. In a similar vein, the impact of system noise and finite measurement may have to be understood in terms of higher-level artifacts, for instance, the ratio between features and sample size instead of just sample size. However, it would be reasonable to conjecture that with higher system noise, the kernel will keep deviating. On the other hand, while in [15] it has been indicated that increased measurement will reduce the generalization error, one may also think the readout error during measurement may be counterproductive. Overall, we observe that the impact of noise is not yet well-understood and should be investigated further.

IV. CONCLUSION

Ultimately, we conducted the first systematic study of QSVM configuration space and quantum classification based on an EHR data set. We classified the persistence of RA patients on biologic therapies, predicting six-month persistence via binary classification. Furthermore, we proposed an end-to-end framework to study EQA that can be generalized for other machine learning and optimization problems and observed EQA for a subset of problem instances in the configuration space. Our framework represents progress toward *a priori* identification of data sets where quantum advantage could be achieved and underscored that even with current quantum computers it is possible to arrive at predictions that are at least as good as those obtained with classical computers. These results have implications for classification problems across industries, particularly for small data sets.

V. METHODS

A. QUANTUM FEATURE MAP

The feature map used in this article is known as the ZZFeatureMap, which gives rise to a feature space of 2^N dimensions where N is the number of qubits [11]. This family of circuits is believed to be hard to simulate classically [34].

B. IBM QUANTUM HARDWARE

ibmq_dublin is a 27-qubit superconducting qubit quantum computer available on the IBM Quantum Services. The qubit connectivity is shown in Fig. 5. For qubits, lighter color




FIGURE 5. Qubit connectivity of `ibmq_dublin`. For qubits, lighter color means higher T2 time; for couplings, lighter color means lower fidelity.




FIGURE 6. Cropped circuit for quantum kernel calculations with full subcircuit feature for one inner product calculation on the top two qubits for a five-feature instance. The full circuit repeats a similar pattern across different qubit pairs. Qubits and interactions were mapped based on the connectivity of `ibmq_dublin`. q_i is the i th algorithmic qubit and the double-digit index after the arrow sign is the physical qubit index on `ibmq_dublin`.

means higher T2 time, and for couplings, lighter color means lower fidelity.

The average CNOT error rate and average readout error rate, at the time of authoring this manuscript, were 1.097%, and 3.585%, respectively. The average T1 and T2 times were 107.03 μ s and 114.53 μ s, respectively. The average gate time was 473.397 ns. More details may be accessed in real time [35]. For every quantum circuit, 1024 shots were run. The circuits were always maximally optimized using application programming interface (API) calls before the runs. A sample circuit computing the feature map of a five-feature data set is given in Fig. 6.

C. QUANTUM SIMULATOR

The quantum simulations were run without noise models on the `qasm_simulator`, available on the IBM Quantum Services. Each circuit was run with 1024 shots and the circuits were always maximally optimized before each run. The simulations supported the experimental design and results.

D. EHR DATA

In this article, there have been two main challenges related to making predictions based on the EHR data. First, the problem of binary classification in patient persistence depends on the quality of the main classification label – that is, whether or not a given patient is persistent on the medication. This is derived from prescription, and there are known challenges in determining patient persistence from prescriptions [36], [37], whereas additional claims-based data sets could be used in

conjunction with EHR data to improve the certainty of prescription patterns, which was not an available option during this work.

In addition, we used imputation to fill in missing data points. For example, not all laboratory results are equally present in each of the patients, and more specifically, not in the period of time covered by the data used for model training and testing. As detailed in the data section, the data were chosen such that for the first 10 features there is no missing data for any selected patient. Above the 10th feature, the data are sparser. Thus, for the features, 11–20 missing data were imputed using the mean of the present data, whereas the impact of that imputation is minimal given that the majority of the model accuracy comes from the top ten features, applying different imputation techniques could be explored in the future.

E. COHORT RESTRICTIONS

The models were built using data from the Optum EHR data set, which includes deidentified and aggregated clinical and medical administrative data from over 100 million longitudinal EHR lives. Fields that were used included demographics, laboratory tests, observations, prescriptions, visits, and selected subsets of extracts from physicians’ notes pre processed using natural language processing methods. A list of RA International Classification of Diseases diagnosis codes was used to select a first set of patients, further narrowed down by the given biologic’s National Drug Codes. The persistence for a given patient was defined as the length of time from initiation to discontinuation of the biologic therapy.

The therapy start date, i.e., the index date, was set to be the start date of the first biologic prescription. Inclusion criteria required at least one year of data before the index date; any data prior to one year before the index date was truncated, thus guaranteeing the same interval length for all patients. Additional inclusion criteria applied were a minimum of 6 months of data after the index date with stable payer insurance and the patient had to be at least 18 years of age and be in an integrated delivery network (which was indicated by the flag in the data set). Exclusion criteria were more than one diagnosis of systemic lupus erythematosus or psoriatic arthritis, combined with prior use of targeted disease-modifying antirheumatic drugs including biologics and Janus kinase inhibitors. The inclusion/exclusion process is illustrated in Fig. 7.

F. CREATION OF TRAIN AND TEST SUBSETS

The full pipeline was developed on AWS/Databricks using PySpark, Python, and SQL. At the point where all preprocessing, inclusion and exclusion criteria were completed and the model training started, the data set size was reduced to 16 000 samples, with a relatively balanced target class (52% of patients persisted, whereas 48% did not). The number of features in that model exceeded 500, with the top 10 features accounting for more than 90% of the achieved accuracy of 0.64. The variability of the model metrics within the set of

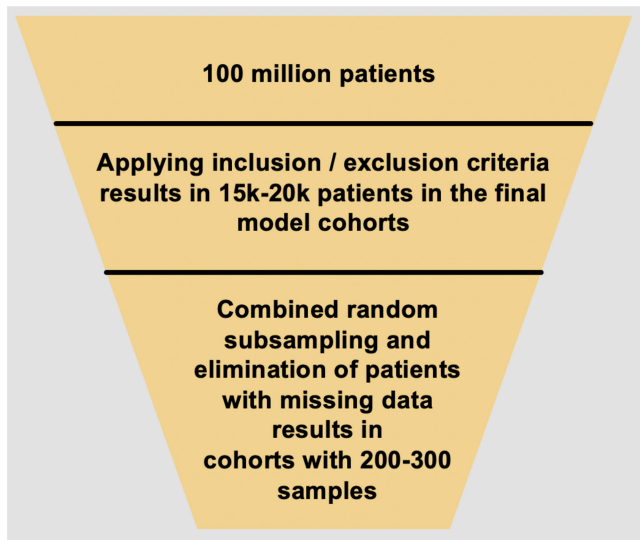


FIGURE 7. Cohort restrictions applied to the EHR data.

10 different train/test splits was under 4%. The train/test split ratio was set at 80/20.

In the preliminary experiments with classical SVMs and quantum simulations, the range of the explored landscape was 50–600 in training set size in increments of 50 and the range in the number of features was from 1 to 20 in increments of 1. Since it was known that the top 20 features carry >95% of accuracy, this was deemed sufficient and within the reach of QSVM, where each feature maps to one qubit, thus leading to 20-qubit QPU experiments.

To reduce the training set size further from 12 800 (80% of the original 16 000), an additional step was applied to the 16 000-sample model data. First, the samples with no missing data in the top 10 features were selected, leaving the features from 11 to 20 with some missing data. That reduction yielded a data set of 1300 samples. This resulted in data sets with minimum missing data while preserving the top 20 features required for the experiments. Since not all patients have all the selected laboratory measurements or other features collected during the year before the start of the medication, we used imputation to fill in those values. Using a longer period of 2–3 years of training data before the index date significantly increases the chance of a patient having at least some value for the given features but reduces the overall number of patients in the cohort; therefore, that approach was not utilized in the final model.

The size of the training set was narrowed down to three different values—200, 250, and 300. The final choice was to use 5, 10, 15, and 20 features, that, when combined with 200, 250, and 300 training set sizes, yielded a 12-point configuration space. From the 1300-sample data set, random sampling was used to create training data sets of 200, 250, and 300 samples. The test set size was kept at 150 to balance the constraints of reasonable runtimes for simulations and QPU

while achieving the best stratification of samples under the circumstances.

The downside of the training (and test) data set reduction to a few hundred samples is the reduction in predictive accuracy, originally at 0.64 with 16 000 samples. This decision was made to explore the very difficult cases where it is hard to get predictions better than random guesses, whereas we could have chosen starting models with 15 000–20 000 samples and accuracies above 0.75, which were available with somewhat different patient cohort structures, our goal was to tackle the most difficult problems. This careful consideration in selecting harder instances has allowed us to observe the empirical advantage of quantum kernels over classical kernels even though none of the circuits was run with more than 1024 shots. In future work, and as quantum hardware and software scales further, we would like to explore train/test data sets with 500–1000 samples, which would allow for reduced loss of accuracy and smaller variability due to reduction in a set size.

Knowing that there will be significant variability in the performance of different train/test data sets for each coordinate due to the small sample size, the maximum possible number of sets was evaluated. Given the preliminary runtimes for quantum simulations and QPU, the decision was made to use two random train/test splits for each configuration space coordinate, resulting in a total of 24 data sets to be run on QPU, whereas two random splits for each coordinate do not fully account for the variability resulting from such small data sets, this had to be limited due to QPU availability and simulation runtimes. Future work could be done to increase the number of data sets for each coordinate from two to ten or more.

During subsampling, it was ensured that the target classification label proportion was kept in the original proportion within each train and test subset. Different models and class imbalance ratios ranging from 1:1 to 1:5 were evaluated and the final decision was to use the aforementioned (almost) balanced class to reduce the impact of the small data set size. It was our judgment that more imbalanced cases would be better addressed in subsequent research.

One of the main challenges with the small data sets is that when splitting train and test sets and training models on multiple splits, the resulting model metrics vary widely, especially with the models where predictive accuracy is not very high. To explore that, 200 random splits at each of the 12 coordinates were made. We calculated classical SVM balanced accuracies for each of the splits, a process that was executed in less than 1 h.

For each of the 24 subpoints, having already calculated classical SVM metrics, quantum simulations and QPU runs were executed, both using the Qiskit framework. In total, 1024 shots were used as that allowed the execution for all 24 points within the time and resources available. The simulations were run with callback specified to provide additional insight during the running processes. The optimization level was set at three, the feature map to the ZZFeatureMap,




FIGURE 8. Distribution of balanced accuracy values for the classical models for different train/test splits for the point with 300 samples and 10 features.

the feature dimension equal to the number of features for the specific subpoint, the number of reps equal to two, and the entanglement to linear.

Quantum simulation and QPU processing were used to calculate custom kernel matrices for the given train and test set. The train kernel was then saved and used in scikit-learn on the classical computer to train the SVM model using a precomputed option. The test kernel was passed to the model’s predict method to make predictions. This way, we generated predictions for quantum simulations and QPU runs. The classical predictions were generated using kernels in SVMs. The models were trained for 18 different values of the regularization parameter C, ranging from 0.006 to 1024, and the best case was used from each model for the classical to quantum comparison. Every model was regularized separately for each of the two metrics; thus, the value of the optimal C parameter for a given set’s balanced accuracy is generally different than the value of C for the F1 score. All three predictions were created for each of the 24 sub-points, whereas the developed framework supports allocating an independent validation data set for the final model metric assessment, a single validation set is unlikely to provide useful insights due to the variability in the metrics for the small data sets in question. Allocating multiple validation data sets and running them through QPU was not feasible with the available time and resources, however; therefore, such a validation step was not included.

G. VARIABILITY AND ERRORS

For the classical models, Fig. 8. illustrates the distribution of balanced accuracy values for 200 different train/test splits for the configuration space coordinate with 300 samples and 10 features.

The scope of the SVM modeling was predicting the labels in the binary classification. We used scikit-learn and its predict method to predict labels directly and the predict_proba method for predicting approximate probabilities. Probabilities are estimated inside the predict_proba method using five-way cross-validation. As such, the process is subject to variability depending on the random seed that is provided to the svm.SVC call. One random value was used for calculations and comparisons to support reproducibility. In addition, a range of different random seed values was used in selected cases to obtain the variability of the predictions. It showed a 1.5% standard deviation in the predictions obtained with the predict_proba method, both in the classical and quantum case, with less distinct values in the quantum domain.

H. RUNTIMES

The QPU runtimes were between 12 and 24 h. During the QPU runs, both CPU and memory utilization was very low in the same Linux server as the main processing was executed on the QPU instance. The QPU processing was sequential.

The QPU processing was executed from a c5.18xlarge Ubuntu 18.04 cloud instance on Amazon Web Services (AWS), with 72 vCPUs and 144 GB RAM. We have not used Amazon Braket; instead, to support managing the whole process, we developed a Python package that has a code generation layer to simplify execution of different configuration spaces and management of the results. From within the package, the Qiskit API calls are made to simulators or QPUs. We used virtual environments that, over the course of a year, allowed us to effectively manage different versions of Qiskit and related software components.

I. SHAPLEY ADDITIVE EXPLANATIONS (SHAP)

SHAP is a game-theoretic approach to explain the output of any machine learning model. It connects optimal allocation with local explanations using the classic Shapley values from game theory and their related extensions. SHAP measures the impact of variables by considering the interaction with other variables. Shapley values calculate the importance of a feature by comparing what a model predicts with and without the feature (variable) [38], [39].

A starting list of the topmost 20 features was obtained from the original 16 000-sample data sets by training machine learning models and then using SHAP to obtain the top 20 features. Those 20 features were then used in the aforementioned analysis with reduced size subsets in the order of SHAP relevance. Each of the points for 5, 10, 15, and 20 features on the configuration space was obtained by taking that many top features from the full 20-feature list, preserving the order.

J. PTRI—SYSTEMATIC IDENTIFICATION OF PROBLEMS WHERE QUANTUM KERNELS MAY HAVE EMPIRICAL ADVANTAGES

Consider a set of classification problems where the number of features is between 1 and M and the number of samples between 1 and N . There are thus $M \times N$ classification problems. To help address the question which subset of these problems should be solved with a quantum kernel, we created a geophysics-inspired approach to identify regions of potential EQA in data sets. One way to select a suitable subset of problems involves studying the ruggedness of the manifold via PTRI, a metric we have adapted from [40] and defined as follows. We are considering the F1 score only as an instance of a metric in the formula. In the $M \times N$ configuration space defined before, each point is surrounded by eight other points except for the boundary points. For the boundary points, the performance result (F1 score in this case) of the adjacent points that are beyond the boundary is assumed to be 0. For the (i,j) th interior point, the local PTRI $_{i,j}$ is calculated according to

$$\begin{aligned} \text{PTRI}_{i,j} = & \left[(F_{l_{i,j}} - F_{l_{i-1,j-1}})^2 + (F_{l_{i,j}} - F_{l_{i-1,j}})^2 \right. \\ & + (F_{l_{i,j}} - F_{l_{i-1,j+1}})^2 + (F_{l_{i,j}} - F_{l_{i,j-1}})^2 \\ & + (F_{l_{i,j}} - F_{l_{i,j+1}})^2 \times (F_{l_{i,j}} - F_{l_{i+1,j-1}})^2 \\ & \left. + (F_{l_{i,j}} - F_{l_{i+1,j}})^2 + (F_{l_{i,j}} - F_{l_{i+1,j+1}})^2 \right]^{1/2}. \end{aligned} \quad (1)$$

To determine the PTRI of the full configuration space, we average across the PTRI $_{i,j}$ values.

VI. DATA AVAILABILITY

The data used in this work are commercially available via the Optum de-identified EHR data set.

VII. CODE AVAILABILITY

The code was developed with Python 3 and Qiskit. Code sections may be obtained on request from corresponding author Z.K.

VIII. COMPETING INTERESTS

The authors declare no competing interests.

IX. AUTHOR CONTRIBUTIONS

Zoran Kronic implemented the end-to-end computational pipeline, including the preprocessing of EHR data and the development of the quantum-classical machine learning models. Frederik F. Flöther contributed to the data preprocessing and machine learning model design. Nathan D. Earnest-Noble supported the optimization and execution of the experiments run on QPUs. Omar Shehab led the theoretical quantum algorithm analysis and helped implement the quantum-classical models. All authors contributed to the

design of the experiments, analyzed the results, and wrote the final manuscript.

ACKNOWLEDGMENT

The authors would like to thank Travis L. Scholten for helpful discussions and Winona Murphy, Matthew Stypulkoski, and Sarah Sheldon for supporting the dedicated access to ibmq_dublin.

REFERENCES

- [1] R. E. Sherman et al., "Real-world evidence – What is it and what can it tell us," *New England J. Med.*, vol. 375, no. 23, pp. 2293–2297, 2016, doi: [10.1056/NEJMs1609216](https://doi.org/10.1056/NEJMs1609216).
- [2] R. Miotto, L. Li, B. A. Kidd, and J. T. Dudley, "Deep patient: An unsupervised representation to predict the future of patients from the electronic health records," *Sci. Rep.*, vol. 6, no. 1, 2016, Art. no. 26094, doi: [10.1038/srep26094](https://doi.org/10.1038/srep26094).
- [3] A. Rajkomar et al., "Scalable and accurate deep learning with electronic health records," *npj Digit. Med.*, vol. 1, no. 1, 2018, Art. no. 18, doi: [10.1038/s41746-018-0029-1](https://doi.org/10.1038/s41746-018-0029-1).
- [4] C. Xiao, E. Choi, and J. Sun, "Opportunities and challenges in developing deep learning models using electronic health records data: A systematic review," *J. Amer. Med. Inform. Assoc.*, vol. 25, no. 10, pp. 1419–1428, 2018, doi: [10.1093/jamia/ocy068](https://doi.org/10.1093/jamia/ocy068).
- [5] S. Ravizza et al., "Predicting the early risk of chronic kidney disease in patients with diabetes using real-world data," *Nature Med.*, vol. 25, no. 1, pp. 57–59, 2019, doi: [10.1038/s41591-018-0239-8](https://doi.org/10.1038/s41591-018-0239-8).
- [6] U. M. Marigorta, J. A. Rodríguez, G. Gibson, and A. Navarro, "Replicability and prediction: Lessons and challenges from GWAS," *Trends Genet.*, vol. 34, no. 7, pp. 504–517, 2018, doi: [10.1016/j.tig.2018.03.005](https://doi.org/10.1016/j.tig.2018.03.005).
- [7] Y. Liu, S. Arunachalam, and K. Temme, "A rigorous and robust quantum speed-up in supervised machine learning," *Nature Phys.*, vol. 17, pp. 1013–1017, 2021, doi: [10.1038/s41567-021-01287-z](https://doi.org/10.1038/s41567-021-01287-z).
- [8] G. Hripesak, D. J. Albers, and A. Perotte, "Exploiting time in electronic health record correlations," *J. Amer. Med. Inform. Assoc.*, vol. 18, no. Supplement_1, pp. i109–i115, 2011, doi: [10.1136/amiajnl-2011-000463](https://doi.org/10.1136/amiajnl-2011-000463).
- [9] M. Schuld and F. Petruccione, *Supervised Learning With Quantum Computers*. Berlin, Germany: Springer, vol. 17, 2018, doi: [10.1007/978-3-319-96424-9](https://doi.org/10.1007/978-3-319-96424-9).
- [10] M. Schuld and N. Killoran, "Quantum machine learning in feature Hilbert spaces," *Phys. Rev. Lett.*, vol. 122, no. 4, 2019, Art. no. 040504, doi: [10.1103/PhysRevLett.122.040504](https://doi.org/10.1103/PhysRevLett.122.040504).
- [11] V. Havlíček et al., "Supervised learning with quantum-enhanced feature spaces," *Nature*, vol. 567, no. 7747, pp. 209–212, 2019, doi: [10.1038/s41586-019-0980-2](https://doi.org/10.1038/s41586-019-0980-2).
- [12] M. Schuld, A. Bocharov, K. M. Svore, and N. Wiebe, "Circuit-centric quantum classifiers," *Phys. Rev. A*, vol. 101, no. 3, 2020, Art. no. 032308, doi: [10.1103/PhysRevA.101.032308](https://doi.org/10.1103/PhysRevA.101.032308).
- [13] N. Mathur et al., "Medical image classification via quantum neural networks," 2021, *arXiv:2109.01831*, doi: [10.48550/arXiv.2109.01831](https://doi.org/10.48550/arXiv.2109.01831).
- [14] J. Heredge, C. Hill, L. Hollenberg, and M. Sevir, "Quantum support vector machines for continuum suppression in B meson decays," *Comput. Softw. Big Sci.*, vol. 5, no. 1, pp. 1–9, 2021, doi: [10.1007/s41781-021-00075-x](https://doi.org/10.1007/s41781-021-00075-x).
- [15] X. Wang, Y. Du, Y. Luo, and D. Tao, "Towards understanding the power of quantum kernels in the NISQ era," *Quantum*, vol. 5, pp. 531, 2021, doi: [10.22331/q-2021-08-30-531](https://doi.org/10.22331/q-2021-08-30-531).
- [16] S. L. Wu et al., "Application of quantum machine learning using the quantum kernel algorithm on high energy physics analysis at the LHC," *Phys. Rev. Res.*, vol. 3, 2021, Art. no. 033221, doi: [10.1103/PhysRevResearch.3.033221](https://doi.org/10.1103/PhysRevResearch.3.033221).
- [17] E. Peters et al., "Machine learning of high dimensional data on a noisy quantum processor," *npj Quantum Inf.*, vol. 7, no. 1, 2021, Art. no. 161, doi: [10.1038/s41534-021-00498-9](https://doi.org/10.1038/s41534-021-00498-9).
- [18] M. Grossi, "W Boson polarization studies for vector boson scattering at LHC: From classical approaches to Quantum computing," Ph.D. thesis, Grad. School Phys., Univ. Pavia, Pavia, Italy, 2021. [Online]. Available: <https://inspirehep.net/literature/1861661>, Accessed: Jun. 21, 2021.

- [19] W. Li and D. L. Deng, "Recent advances for quantum classifiers," *Sci. China Phys., Mech. Astron.*, vol. 65, no. 2, pp. 1–23, 2022, doi: [10.1007/s11433-021-1793-6](https://doi.org/10.1007/s11433-021-1793-6).
- [20] C. Blank, D. K. Park, J. K. K. Rhee, and F. Petruccione, "Quantum classifier with tailored quantum kernel," *npj Quantum Inf.*, vol. 6, no. 1, 2020, Art. no. 41, doi: [10.1038/s41534-020-0272-6](https://doi.org/10.1038/s41534-020-0272-6).
- [21] S. Jerbi, L. J. Fiderer, H. P. Nautrup, J. M. Kübler, H. J. Briegel, and V. Dunjko, "Quantum machine learning beyond kernel methods," 2021, *arXiv:2110.13162*, doi: [10.48550/arXiv.2110.13162](https://doi.org/10.48550/arXiv.2110.13162).
- [22] V. Heyraud, Z. Li, Z. Denis, A. L. Boité, and C. Ciuti, "Noisy quantum kernel machines," 2022, *arXiv:2204.12192*, doi: [10.48550/arXiv.2204.12192](https://doi.org/10.48550/arXiv.2204.12192).
- [23] H. Y. Huang *et al.*, "Power of data in quantum machine learning," *Nature Commun.*, vol. 12, no. 1, 2021, Art. no. 2631, doi: [10.1038/s41467-021-22539-9](https://doi.org/10.1038/s41467-021-22539-9).
- [24] M. Fernández-Delgado, E. Cernadas, S. Barro, and D. Amorim, "Do we need hundreds of classifiers to solve real world classification problems?," *J. Mach. Learn. Res.*, vol. 15, no. 1, pp. 3133–3181, 2014, doi: [10.1038/s41467-021-22539-9](https://doi.org/10.1038/s41467-021-22539-9).
- [25] J. Preskill, "Quantum computing and the entanglement frontier," 2012, *arXiv:1203.5813*, doi: [10.48550/arXiv.1203.5813](https://doi.org/10.48550/arXiv.1203.5813).
- [26] D. Ristè *et al.*, "Demonstration of quantum advantage in machine learning," *npj Quantum Inf.*, vol. 3, no. 1, 2017, Art. no. 16, doi: [10.1038/s41534-017-0017-3](https://doi.org/10.1038/s41534-017-0017-3).
- [27] C. C. McGeoch, *A Guide to Experimental Algorithmics*. Cambridge, U.K.: Cambridge Univ. Press, 2012, doi: [10.1017/CBO9780511843747](https://doi.org/10.1017/CBO9780511843747).
- [28] B. A. Cordier, N. P. Sawaya, G. G. Guerreschi, and S. K. McWeeney, "Biology and medicine in the landscape of quantum advantages," 2021, *arXiv:2112.00760*, doi: [10.48550/arXiv.2112.00760](https://doi.org/10.48550/arXiv.2112.00760).
- [29] GitHub, "SHAP," 2021. [Online]. Available: <https://github.com/slundberg/shap>. Accessed: Nov. 18, 2021.
- [30] Scikit-Learn, "Machine learning in Python," 2021. [Online]. Available: <https://scikit-learn.org>. Accessed: Nov. 18, 2021.
- [31] O. Regev, "On lattices, learning with errors, random linear codes, and cryptography," *J. ACM*, vol. 56, no. 6, 2009, Art. no. 34, doi: [10.1145/1568318.1568324](https://doi.org/10.1145/1568318.1568324).
- [32] D. Bertsimas, J. Pauphilet, and B. Van Parys, "Sparse classification: A scalable discrete optimization perspective," *Mach. Learn.*, vol. 110, no. 11, pp. 3177–3209, 2021, doi: [10.1007/s10994-021-06085-5](https://doi.org/10.1007/s10994-021-06085-5).
- [33] H. Sifaou, A. Kammoun, and M. S. Alouini, "Phase transition in the hard-margin support vector machines," in *Proc. IEEE 8th Int. Workshop Comput. Adv. Multi-Sensor Adaptive Process*, 2019, pp. 415–419, doi: [10.1109/CAMSAP45676.2019.9022461](https://doi.org/10.1109/CAMSAP45676.2019.9022461).
- [34] S. Aaronson and A. Ambainis, "Forrelation: A problem that optimally separates quantum from classical computing," *SIAM J. Comput.*, vol. 47, no. 3, pp. 982–1038, 2018, doi: [10.1145/2746539.2746547](https://doi.org/10.1145/2746539.2746547).
- [35] IBM, "ibmq_dublin," 2021. [Online]. Available: https://quantum-computing.ibm.com/services?services=systems&system=ibmq_dublin. Accessed: Oct. 28, 2021.
- [36] V. Ehrenstein, H. Kharrazi, H. Lehmann, and C. O. Taylor, "Chapter 4 – Obtaining data from electronic health records," in *Tools and Technologies for Registry Interoperability, Registries for Evaluating Patient Outcomes: A User's Guide, 3rd Edition, Addendum 2 [Internet]*, R. E. Gliklich, M. B. Leavy, and N. A. Dreyer, Eds. Rockville, MD, USA: Agency Healthcare Res. Quality, 2019.
- [37] L. Bastarache *et al.*, "Developing real-world evidence from real-world data: Transforming raw data into analytical datasets," *Learn Health Syst.*, vol. 6, 2021, Art. no. e10293, doi: [10.1002/lrh2.10293](https://doi.org/10.1002/lrh2.10293).
- [38] S. M. Lundberg and S. I. Lee, "A unified approach to interpreting model predictions," in *Proc. 31st Int. Conf. Neural Inf. Process. Syst.*, 2017, pp. 4768–4777. [Online]. Available: <https://proceedings.neurips.cc/paper/2017/hash/8a20a8621978632d76c43dfd28b67767-Abstract.html>
- [39] S. M. Lundberg *et al.*, "From local explanations to global understanding with explainable AI for trees," *Nature Mach. Intell.*, vol. 2, pp. 56–67, 2020, doi: [10.1038/s42256-019-0138-9](https://doi.org/10.1038/s42256-019-0138-9).
- [40] S. J. Riley, S. D. DeGloria, and R. Elliot, "A terrain ruggedness index that quantifies topographic heterogeneity," *Intermountain J. Sci.*, vol. 5, no. 1/4, pp. 23–27, 1999. [Online]. Available: http://download.osgeo.org/gis/doc/reference-docs/Terrain_Ruggedness_Index.pdf



Zoran Krunic received the M.S. degree in electrical engineering and computer science from the University of Belgrade, Belgrade, Serbia.

He worked in data engineering, data warehousing, and real-time systems development as a consultant for a number of large companies – Capital Group, ARCO Petroleum, Southern California Gas company, and PacifiCare. He then spent several years at Optum, where he worked on creating machine learning applications to predict hardware and software failures in enterprise architecture. He joined Amgen R&D in 2018 and since then has worked on applications of machine learning to modeling different patient outcomes and clinical trial enrollment based on electronic health records and clinical trial data sets. In the last two years, he led the quantum machine learning initiative in cooperation with the IBM Quantum research team, evaluating machine learning models in the quantum domain on the Qiskit platform and comparing them to classical machine learning models.



Frederik F. Flöther received the B.A. and M.Sc. degrees in natural sciences in 2011, the M.A. degree in natural sciences in 2014, and the Ph.D. degree in physics (photonic quantum computing) in 2015, all from the University of Cambridge, Cambridge, U.K., as well as a professional certificate in quantum computing applications from Massachusetts Institute of Technology, Cambridge, MA, USA, in 2019.


From 2015 to 2018, he was a Consultant and Data Scientist with IBM. Since 2018, he has been an industry consultant and computational scientist with IBM Quantum, Zurich, Switzerland, leading consulting and research partnerships with a focus on the healthcare and life sciences sectors. In total, he is the author of more than 40 inventions, peer-reviewed publications, and white papers. His research interests include the intersection of quantum computing, machine learning, and industry applications.

Dr. Flöther was elected to the IBM Academy of Technology and is an IBM Master Inventor and Qiskit Advocate.



George Seegan received the B.Sc. degree in chemistry and the Ph.D. degree in biochemistry from the University of California at Los Angeles, Los Angeles, CA, USA, in 1975 and 1981, respectively.

He joined Amgen in 2002. From 2004 to 2014, he led a team in the Research Informatics group at Amgen, which supported the translational sciences functional area in research and development. He has broad experience implementing technical software development solutions to advance scientific information management, scientific and clinical data analysis, and high-performance computing to support modeling and simulation. His experience includes extensive development and GxP validation of computer systems used in regulated environments. In 2014, he transitioned to the first Data Scientist role at Amgen with dual responsibilities for the Predictive Analytics group in the Center for Observational Research and the Digital Health group in Strategy and Innovation where he develops machine learning implementations focused on clinical outcome prediction, patient identification and stratification, and patient experience quantification.



Nathan D. Earnest-Noble received the B.S. degree in physics and applied mathematics from the University of California, Berkeley, Berkeley, CA, USA, in 2008, and the M.S. and Ph.D. degrees in physics from the University of Chicago, Chicago, IL, USA, in 2013 and 2018, respectively.

During his graduate studies, he worked on designing potential next generation superconducting qubits, which are robust to environmental noise, such as the Heavy Fluxonium or 0-Pi qubits. He joined IBM Quantum in 2019. He is currently a Quantum Software Integration Lead and Quantum Computing Application Researcher. He has worked on developing quantum algorithms for near term applications, leads efforts to integrate IBM's quantum software across the stack with external partners, collaborates with quantum computing researchers to leverage knowledge of hardware for the benefit of quantum algorithms and designs, and teaches courses on quantum algorithms.



Omar Shehab received the Ph.D. degree in computer science from the University of Maryland Baltimore County, Baltimore, MD, USA, in 2016. His dissertation was on quantum complexity theory.

From 2016 to 2017, he was a Visiting Assistant Professor at the University of Maryland Baltimore County. From 2017 to 2020, he was a Quantum Theorist at IonQ, Inc. Since 2020, he has a Quantum Computing Applications Researcher at IBM Quantum, Yorktown Heights, NY, USA.

Dr. Shehab was the recipient of funds by NASA, DoE, and DARPA.